



NTNU – Trondheim
Norwegian University of
Science and Technology

MASTER'S THESIS

MANAGERIAL ECONOMICS AND OPERATIONS RESEARCH

TIØ4905

Optimisation-Based Nurse Scheduling for Real-Life Instances

Authors:

Francesco Rugiano BECKMANN
Kjartan Kastet KLYVE

Supervisors:

Prof. Marielle CHRISTIANSEN
Prof. Henrik ANDERSSON
Post doc. Anders GULLHAV

November, 2016, Trondheim

DEPARTMENT OF INDUSTRIAL ECONOMICS
AND TECHNOLOGY MANAGEMENT

FACULTY OF SOCIAL SCIENCES
AND TECHNOLOGY MANAGEMENT

Problem statement in master's thesis contract

This thesis aims to develop a tool for optimisation-based scheduling of one of two maternity wards at St. Olavs Hospital in Trondheim. The ward has approximately 65 employees, with different skills, contracted work times, preferences etc. Furthermore, there is a certain demand for the different kinds of employees at different times of the week. This makes the task of creating their work schedules a very complex combinatorial problem. Today, the problem is solved manually, by a ward manager, who works primarily with solving the scheduling problem. Furthermore, other employees in the ward take part in the process to a varying degree, making the total time spent on the scheduling problem very long. We wish to develop an optimisation-based tool for performing the scheduling that will serve as decision support for the ward manager, creating schedules that can either be used directly or serve as drafts for the scheduling problem, depending on the wish of the ward manager. The tool will create schedules that respect all relevant rules and regulations as well as taking into account costs, employee preferences, fairness etc.

Preface

This master's thesis concludes our Master of Science at the Norwegian University of Science and Technology (NTNU) with specialisation in Applied Economics and Optimisation at the Department of Industrial Economics and Technology Management.

The master's thesis is written in collaboration with St. Olavs Hospital - Trondheim University Hospital.

We would like to thank our academic supervisors Professor Marielle Christiansen, Professor Henrik Andersson and Post Doctoral researcher Anders Gullhav at the Industrial Economics and Technology Management faculty at NTNU for their invaluable guidance and contributions during our work. We would also like to thank Ward Manager at Maternity Ward West at St. Olavs Hospital, Vigdis Myhren, for her continuous feedback and contributions to the development of the scheduling model. Furthermore, we want to thank Runa Heimstad, Aud Hiller, Valborg Sund, Johan Fredrik Skomsvoll, Jo-Åsmund Lund, Unn Elin Dahlberg, Per Olav Østbyhaug, Birger Henning Endreseth and Dr. Lars Johan Vatten, the board members of the Regional Centre of Health Care Development, for providing us with the mandate to collaborate with St. Olavs Hospital in this master's thesis.

Francesco Rugiano Beckmann and Kjartan Kastet Klyve

Trondheim, November, 2016

Abstract

In this master's thesis we present a planning problem at Maternity Ward West (MWW) at St. Olavs Hospital, concerning the scheduling of 69 employees for a planning horizon of 27 weeks. The scheduling problem involves covering demand for health workers of different skill categories, while respecting employees' preferences as much as possible and ensuring fairness. The goal of the master's thesis is to create a decision support tool that solves the scheduling problem by producing schedules for MWW of such quality that they are preferable to the manually made schedules produced in the current planning process at MWW. Furthermore, we discuss related literature and provide a theoretical context for our work. This leads to the formulation of the problem scope and the problem description.

Subsequently we formulate a general mathematical integer linear programming model and develop it using the commercial optimisation software FICO[®] Xpress Optimisation Suite 7.8. We refer to it as the MWW scheduling model. It runs successfully, creating a schedule with real-life data for all the employees at MWW. The same data has been used to create a schedule using manual techniques at MWW, making it possible to compare the results of the techniques. Relevant data is presented in Table 1.

The ward manager at MWW states that the MWW scheduling model guarantees the employees' influence on the the schedules, because it prioritises employees' preferences and lets the ward manager and scheduling group make changes to the produced schedules if needed. The MWW scheduling model respects preferences as it produces schedules allocating shifts in

Table 1: Comparison of the manually made schedule and the optimisation-based schedule.

Key information	Manual schedule	Scheduling model
Employee influence ensured	Yes	Yes
%Requests respected	85.3	89.8
Max over-coverage respected	No	Yes
All scheduling rules respected	No	Yes
Unbiased shift-allocation	No	Yes
Unfair bartering	Yes	No
Time to create schedule	6 weeks	5 min - 24 h

accordance with 89.8% of the employees' requested shifts. Furthermore, the MWW scheduling model distributes over-coverage more evenly than the manually made schedules, thus securing a more robust schedule. Also, the MWW scheduling model guarantees that schedules always abide by all scheduling rules. Another strength the MWW scheduling model possesses is its lack of bias when allocating shifts, making schedules fair to all employees. The current manual method for scheduling, as opposed to the automatic scheduling provided by the MWW scheduling model, includes a bartering process which the ward manager states is unfair. Removing the need for a bartering process with optimisation-based scheduling is perhaps the single most efficient measure to make the current planning process at MWW more fair. Lastly, the MWW scheduling model creates good schedules very fast. Although the scheduling model likely needs minor adjustments before creating schedules for new planning horizons, optimisation-based scheduling is still remarkably faster than the current scheduling system at MWW. The MWW scheduling model finds good integer solutions within few minutes and is close to reaching optimality within 24 hours, with an optimality gap of 0.036% for the full real-life instance.

We also use the MWW scheduling model to perform analyses. We perform technical analyses that show how the model is very scalable for different planning horizons and different staff levels. Furthermore, we have shown

that the MWW scheduling model can be used as a management tool for tactical and strategic decisions, by implementing changes in policies and staff levels for different instances and producing feasible schedules for these instances. These policy changes have been chosen after discussions with the ward manager and after receiving input from the board of the Regional Centre of Health Care Development. Most notable is that it seems realistic to open an extra bed unit during weekends without increasing the current staff level, by implementing a policy change that trades extra weekend shifts for extra off-days. Our analyses also show that the MWW can meet coverage requirements with reduced staff levels. Lastly, we perform an analysis that shows that employees can be scheduled to work less than contracted, creating extra off-shifts that serve as a buffer for tackling staffing shortages due to sudden long-term sickness. This policy proves complicated, and the approach needs further development and greater insights into the online operational planning level to be efficient.

We have succeeded in reaching our goal of creating a decision support tool that solves the scheduling problem by producing schedules for MWW of such quality that they are preferable to the manually made schedules, and the best testimony to this is that the ward manager at MWW states that she wants to use our model in her work and that she wants a similar model developed for other wards.

Sammendrag

I denne masteroppgaven presenterer vi et planleggingsproblem på Fødeavdeling vest ved St. Olavs hospital, som omhandler turnusplanlegging av 69 ansatte for en planleggingshorisont på 27 uker. I turnusplanleggingsproblemet må flere ulike etterpørselstyper dekkes, ansattes preferanser må respekteres i størst mulig grad og turnusplanen må være rettferdig. Målet i masteroppgaven er å lage et verktøy for beslutningsstøtte som generer turnusplaner for Fødeavdeling vest med en så høy kvalitet at de er å foretrekke fremfor deres manuelt produserte turnusplaner. I avhandlingen vår presenterer vi den nåværende manuelle planleggingsprosessen ved fødeavdelingen og diskuterer relevant teori for å sette arbeidet vårt inn i en teoretisk kontekst. Dette leder opp til problemavgrænsningen og problembeskrivelsen for denne masteroppgaven.

Videre formulerer vi en matematisk linær heltallsprogrammeringsmodel og utvikler den med den kommersielle optimeringsprogramvaren FICO[®] Xpress Optimisation Suite 7.8. I masteroppgaven kaller vi denne modellen 'MWW scheduling model', men i dette sammendraget refererer vi til den som turnusmodellen. Turnusmodellen evner å produsere turnusplaner ved hjelp av reell informasjon om de ansatte på Fødeavdeling vest. Den samme informasjonen har blitt brukt til å manuelt produsere en turnusplan, som gjør det mulig å sammenlikne resultater for de to metodene. En oppsummering av disse resultatene finnes i Tabell 2.

Avdelingslederen på Fødeavdeling vest slår fast at turnusmodellen garan-

Table 2: Sammenlikning av den manuelt produserte turnusplanen og den optimeringsbaserte turnusplanen.

Nøkkelinformasjon	Manuell turnus	Optimeringsbasert turnus
Ansattes påvirkningskraft bevart	Ja	Ja
%Ønsker oppfylt	85.3	89.8
Maks overbemanning respektert	Nei	Ja
Alle turnusregler respektert	Nei	Ja
Objektiv skiftallokering	Nei	Ja
Urettferdig jenking	Ja	Nei
Tid for a lage turnus	6 uker	5 min - 24 t

terer for de ansattes påvirkning på turnusplanen, ettersom den prioriterer de ansattes preferanser og tillater avdelingslederen og turnurplanleggingsgruppen å gjøre endringer i den planen ved behov. Utsagnet støttes av at modellen evner å allokere vakter som sammenfaller med 89.9% av de ansattes ønskede vakter. Utover dette evner modellen å utjevne overbemanningen bedre enn hva fødeavdelingen ellers evner manuelt, og dette medfører at turnusplanen blir mer robust. Turnusmodellen vil også alltid garantere at alle turnusregler overholdes, mens dette ikke er tilfellet for de manuelt produserte turnusplanene. En annen styrke ved modellen er at den er upartisk og dernest behandler alle de ansatte like rettferdig. Den manuelle planleggingen innebærer derimot en såkalt jenkeprosess, og avdelingslederen slår fast at den er urettferdig. Å bruke optimeringsbasert turnusplanlegging til å fjerne behovet for jenkeprosessen er muligens det mest effektive tiltaket for å gjøre planleggingsprosessen mer rettferdig. Videre lager turnusmodellen gode turnusplaner raskt. Selv om turnusmodellen trolig må justeres noe når den skal produsere turnusplaner for nye planleggingsperioder, er optimeringsbasert turnusplanlegging betraktelig raskere enn den manuelle turnusplanleggingen på Fødeavdeling vest. Turnusmodellen finner gode heltallsløsninger etter noen minutter og finner løsninger nært optimal løsning på 24 timer, med et optimalitetsgap på 0.036% for den fulle instansen.

Vi bruker også turnusmodellen til å gjøre analyser. Vi utfører en teknisk analyse som viser at turnusmodellen er veldig skalerbar både for ulike planleggingshorisonter og ulike bemanningsnivå. Vi har også vist at turnusmodellen kan brukes som et styringsverktøy for taktiske og strategiske beslutninger. Dette er gjort ved å implementere endringer i styringsprinsipper og bemanningsnivå i ulike instanser og vist at gyldige turnusplaner kan genereres. Endringene i styringsprinsipper er utformet basert på samtaler med avdelingslederen på fødeavdelingen og etter tilbakemeldinger fra styret i Regionalt senter for helseutvikling. Det mest oppsiktsvekkende resultatet er trolig at er realistisk å åpne et ekstra sengetun i helgene uten å øke bemanningen ved fødeavdelingen, ved å benytte en ordning der ansatte kan bytte til seg ekstra frivakter mot å jobbe ekstra helger. I tillegg viser analysene våre at Fødeavdeling vest kan dekke sitt nåværende bemanningsnivå med en redusert bemanning. Vi har også utført en analyse av et styringsprinsipp som viser at ansatte kan allokere et redusert antall arbeidstimer for å kunne dekke vakter med frafall grunnet plutselig langtidssykdom. Dette styringsprinsippet er derimot komplisert, og tilnærmingen krever bearbeiding og mer detaljert innsikt i det online operasjonelle planleggingsnivået.

Vi har nådd målet vårt om å lage et verktøy for beslutningsstøtte som generer turnusplaner for Fødeavdeling vest med en så høy kvalitet at de foretrukne fremfor de manuelt produserte turnusplanene. Det beste vitnesbyrdet for dette er at avdelingslederen på Fødeavdeling vest ønsker å bruke turnusmodellen i arbeidet sitt og at hun ønsker liknende modeller utviklet også for flere avdelinger.

Contents

Problem statement in master's thesis contract	iii
Preface	v
Abstract	vii
Sammendrag	xi
1 Introduction	1
2 Background	3
2.1 Terminology	3
2.2 Health care industry	4
2.3 St. Olavs Hospital	6
2.4 Maternity Ward West	7
2.4.1 Brief overview	7
2.4.2 Employees	9
2.4.3 Days	12
2.4.4 Shifts	12
2.4.5 Demand type	12
2.4.6 Scheduling at the Maternity Ward	13
2.5 The planning process	14
	xv

3	Literature	21
3.1	Nurse scheduling and OR Terminology	22
3.1.1	Administrative modes	22
3.1.2	Useful terms	23
3.2	Positioning in the literature	24
3.2.1	A fitting framework	24
3.2.2	Managerial areas	25
3.2.3	Hierarchical decomposition	26
3.2.4	Personnel scheduling	27
3.2.5	Positioning the planning process at MWW	28
3.3	Key aspects within nurse scheduling	29
3.3.1	Preferences	29
3.3.2	Fairness	35
3.3.3	Costs	38
3.4	Optimisation approaches	39
3.4.1	Exact methods	40
3.4.2	Heuristic approaches	41
3.4.3	Implementation and optimality	42
3.4.4	Optimisation at MWW	43
4	Scope and problem description	45
4.1	Problem Scope	45
4.1.1	Key aspects	46
4.1.2	Hierarchical limits	46
4.1.3	The MWW scheduling model in the planning process	47
4.2	Problem description	52
4.2.1	Employees	52
4.2.2	Shifts	52
4.2.3	Demand types	53
4.2.4	List of scheduling rules	53
4.2.5	Objectives	55

5	Model	57
5.1	Modelling principles	57
5.1.1	Model design	57
5.1.2	Mathematical formulation	58
5.2	Definitions	58
5.2.1	Indices	59
5.2.2	Sets	59
5.2.3	Parameters	60
5.2.4	Variables	64
5.3	Objective function	65
5.4	Constraints	65
5.4.1	Covering demand	65
5.4.2	Work hours	66
5.4.3	Required rest	67
5.4.4	Weekends	68
5.4.5	Shift-patterns	68
5.4.6	Related to fairness	75
5.4.7	Variable declarations and fixations	75
6	Computational Study	77
6.1	Instances	78
6.2	Technical analyses	81
6.2.1	Length of planning horizon	82
6.2.2	Number of employees	85
6.2.3	Combining planning horizon and employees	87
6.3	Staffing analyses	89
6.3.1	Open extra bed unit during weekends	90
6.3.2	Reduced staff level	97
6.3.3	Open hours	102
6.4	Comparing schedules	107
7	Concluding remarks	113
8	Future Work	115

9	Compressed model	117
9.1	Definitions	117
9.1.1	Indices	117
9.1.2	Sets	117
9.1.3	Parameters	119
9.1.4	Variables	121
9.2	Objective function	121
9.3	Constraints	121
9.3.1	Covering demand	121
9.3.2	Work hours	122
9.3.3	Required rest	122
9.3.4	Weekends	123
9.3.5	Shift-patterns	123
9.3.6	Related to fairness	124
9.3.7	Variable declarations and fixations	124

List of Figures

2.1	Rough presentation of the current planning process, that visualises the information flow that directly results in the final schedules. In reality, much information also flows through other channels than depicted.	19
3.1	Managerial areas vs hierarchical decomposition	25
3.2	Managerial areas vs hierarchical decomposition. Personnel scheduling, staffing and scheduling positioned.	28
4.1	Rough presentation of the new planning process including the MWW scheduling model, that visualises the information flow that directly results in the final schedules. In reality, much information also flows through other channels.	51
6.1	Excerpt showing two weeks of the schedule for the midwives chosen to be allocated open hours in the in the '9E' instance. Shifts 'F1' and 'F2' are off-shifts and are highlighted in the figure.	104
6.2	During each day in the planning horizon, midwives with open hours are allocated to a number of off-shifts. The sum of these off-shifts vary on each day, and the columns show the number of times different sums occur in the planning horizon.	105
6.3	Graphic presentation of running the full instance of 69 employees for a planning horizon of 27 weeks.	108

List of Figures

- 6.4 A small excerpt from an anonymous version of the output. We chose 14 random employees during 2 random weeks, to help visualise the output of the MWW scheduling model. The rows containing "P:" shows the planned schedule while rows containing "R:" shows the employee's requests. . . . 109

List of Tables

1	Comparison of the manually made schedule and the optimisation-based schedule.	viii
2	Sammenlikning av den manuelt produserte turnusplanen og den optimeringsbaserte turnusplanen.	xii
2.1	Examples of some of the most usual skill categories and demand types. An 'x' denotes eligibility to cover a demand type. Note how midwives are eligible to pose as assistant nurses, while assistant nurses may act as ward assistants.	13
5.6	Illegal shift-pattern p_1 , indicating it is illegal to work 'N-OFF-N'	63
5.7	Illegal shift-pattern p_2 , indicating it is illegal to work 'D-OFF-N-E'	63
5.11	Illegal shift-pattern p_A , indicating it is illegal to work 'D-E-OFF-N-E'	71
5.12	Shift-pattern p_1 , 'N-N-N', which must be followed by p_2 if p_1 is allocated	73
5.13	Shift-pattern p_2 , 'OFF-OFF', which must succeed p_1 when p_1 is allocated	73
6.1	Key data to understand the size of the scheduling problem at MWW today.	78

List of Tables

6.2	Legal ranges of demand for all the 7 relevant demand types.	79
6.3	Amount of work during weekends and nights for standard midwives and assistant nurses throughout a 27 week planning horizon.	80
6.4	Objective function weighting parameters and the proportion of the contracted work each employee must work.	81
6.5	Run times and gaps for the three instances of different planning horizons, spanning 27, 12 and 6 weeks for all 69 employees.	82
6.6	Number of constraints and variables created for the different instances of weeks, with all employees.	84
6.7	Run times and gaps for the three instances of different numbers of employees, with the full 27 week planning horizon. .	86
6.8	Number of constraints and variables created for the different instances of employees, for 27 weeks.	86
6.9	Run times and gaps for all combinations of the instances of different time horizons and employees.	88
6.10	Number of constraints and variables created for all combinations of the instances of different time horizons and employees.	89
6.11	The increase in demand for different days when we open the extra bed unit. Note that night shifts on Mondays begin Sunday 22:15.	91
6.12	Legal ranges of demand when we open an extra bed.	91
6.13	Key data for different instances that address MWW's problem of opening the extra bed unit during the weekends. %Utilised work hours is defined as scheduled work, hours counted for vacations and extra off-days divided by the total contracted work.	95
6.14	Table presenting key data from the most interesting instances of reduced staff levels.	99
6.15	Data for different instances open hours is implemented for.	103
6.16	Key data for the full instance.	107
6.17	Information for comparing the manual schedule to the one created by the MWW scheduling model.	110

Chapter 1

Introduction

Health care organisations are expensive to service, and over the last decades public health care expenditure has seen a steady rise, spurring the need for better resource utilisation (Hans et al., 2011). Hospitals encompass a great variety of health care services, making hospital management a highly complex task. Furthermore, combining the variety of services with cost efficiency is challenging, given the high service standards required in the health care industry.

Many hospitals employ thousands of people and salaries make up large portions of hospitals' overall costs, St. Olavs Hospital being no exception. Shift work is used to staff the wards, which in many cases are open around the clock. This requires schedules in which employees are assigned with shifts and working times throughout some planning period. Solving the scheduling problem, i.e. making sure all shifts are covered with sufficient staff, while respecting numerous employee preferences and complying with work place regulations, is no trivial task. The scheduling problem is simply too complex for any good schedule to be identified easily. Nevertheless, the majority of hospitals still create schedules through manual methods, although these have proven both time-consuming and expensive (Kellogg and Walczak, 2007).

The goal of this master's thesis is to create a decision support tool that solves the scheduling problem by producing schedules for Maternity Ward West (MWW) of such quality that they are preferable to manually made schedules. To reach this goal we will attain a thorough understanding of the planning process at the ward and develop an integer programming model, referred to as the MWW scheduling model. The scheduling problem at MWW is best described as a nurse scheduling problem, even though several employees are not nurses. Our contribution to the maternity ward is an objective and fair decision support tool drastically lowering the time spent on scheduling, while also increasing the resource utilisation of human capital. Also we provide insights into the scheduling problem from an operations research point of view, aiding hospital management in cost considerations related to staffing and scheduling. Our academic contribution is to present a success story of utilising existing theory in developing a useful scheduling model for a highly detailed real-life scheduling problem. Norwegian health care industry is characterised by strong trade unions, as well as a culture of considering employee preferences greatly, making the scheduling problem at MWW more detailed than what is common in literature in the field of nurse scheduling.

This thesis begins by providing insights into the health care industry, St. Olavs Hospital and its MWW, before explaining the planning process at the ward in Chapter 2. Subsequently, the scheduling problem at MWW is given theoretical context, and useful literature is presented in Chapter 3. In Chapter 4 the scope of our thesis and the problem description is presented, before our model is described in Chapter 5. In Chapter 6, we perform a computational study. Lastly, the concluding remarks and future work is presented in Chapters 7 and 8, respectively.

Chapter 2

Background

The purpose of this chapter is to present background information relevant for the master's thesis. The information about Maternity Ward West is collected through close co-operation with the ward manager, involving several meetings and continuous e-mail correspondence. We begin by presenting important terminology in Section 2.1, useful for understanding scheduling in the health care industry. Subsequently we introduce the health care industry and present key figures and problem areas to put this thesis into a wider context in Section 2.2. Furthermore, we narrow in on St. Olavs Hospital and present information relevant to understand characteristics typical of the Norwegian health care industry in Section 2.3. Lastly we present Maternity Ward West in Section 2.4 and the planning process related to producing schedules at MWW in Section 2.5.

2.1 Terminology

In this section we present some useful terms for understanding the background information of this thesis. The terms are written in italic style the first time they occur later in the text.

- *Individual schedules* - the part of a schedule that only includes one person's shifts throughout the planning period.

- *Planning horizon* - the time period planned for in the scheduling problem, i.e. the time period of a schedule.
- *Resource Management System* - software used to register employees' requests for shifts, create schedules, calculate worked hours and salaries etc. It also warns the user if she registers schedules that break governmental rules. The Resource Management System is created for St. Olavs Hospital by an external firm and is based on Microsoft Excel.
- *Red days* - days that are either a holiday or a Sunday.
- *Scheduling* - the act of creating schedules.
- *Scheduling rules* - An umbrella term encompassing governmental work regulations, ward policies, preferred practises and agreements with trade unions concerning schedules.
- *Self-scheduling* - a system in which nurses create their own schedules, then barter amongst themselves until all requirements regarding staff coverage is fulfilled. This approach is popular amongst nurses, as it lets them influence the schedules with their preferences greatly. The process is however very time consuming and in some cases the negotiations cause conflicts amongst the nurses.
- *Shift-pattern* - the combination of shifts a nurse has over a set of consecutive days.
- *Work regulations* - contractual agreements between nurses and the hospital or between the hospital and the government. E.g mandatory rest between shifts.

2.2 Health care industry

Population growth and the rise in human life expectancy continuously increase the need for and expenditure of medical services. There is therefore

a need to cope with the overall increase in demand and costs for medical services, through the employment of more healthcare personnel and better resource utilisation.

In a 2005-report from the Organisation for Economic Co-operation and Development (OECD) it was indicated that all countries, with the exception of a few OECD-countries, face shortages in health personnel, and that this shortage would persist or even increase in the future unless efforts were made to increase employment or productivity (Simoens et al., 2005). An analysis conducted by the World Health Organization (WHO) and Global Health Workforce Alliance (GHWA) confirmed that the shortage persisted, when a global shortage of 7.2 million professional health workers was estimated in 2012. Alliance (2012) predicts that this number will rise beyond 12 million over the next decades. Shortage of health workers is a problem not only relevant for under-developed regions of the world. Countries of all socio-economic development face challenges in guaranteeing their citizens universal health care (Alliance, 2012).

Large health care organisations like hospitals are very complex when considering the scope of the services they provide. Hospitals often employ thousands of health care workers and service several wards around the clock. Logistically it is difficult to manage all these services and employees while ensuring cost efficiency and high quality health care. Hans et al. (2011) point to several reasons for why health care organisations differ from organisations in other industries. Large health care providers like hospitals are typically made up of autonomously managed departments, and their managers tend not to properly consider their department as part of a greater interconnected planning environment. This makes planning and control more fragmented than what is typically the case in e.g. manufacturing, where the entire supply chain must be considered for maximising profits. Even though health care managers are generally committed to providing the best possible services, they typically come from a background within health care, and lack sufficient training and knowledge to make optimal use of scarce resources in complex planning environments. An example of

this is that ward managers are typically attained by promoting health care workers with long tenure, as opposed to hiring someone with a background in management and planning.

2.3 St. Olavs Hospital

St. Olavs Hospital is amongst Norway's five largest hospitals and employs roughly 10,000 people. In the period 2009 to 2015 operating expenses increased by approximately 58% and amounted to 9,931,922,000 NOK in 2015, indicating a need for increased cost efficiency. The amount spent on wages accounted for approximately 63% in 2009 and 65% in 2015 of total operating expenses (Helse-Midt-Norge, 2009, 2015).

Several wards at St. Olavs hospital are open around the clock and shift work is used to ensure that all wards are staffed sufficiently. There is no consensus amongst the wards on how planning the shift work should be done, so *scheduling* is conducted as the ward managers see fit, without significant involvement from the upper administration other than deciding the staff levels and budgets. As is typical in the health care industry, the ward managers usually come from a background as nurses or doctors, without any particular training or education in management and planning.

It is common knowledge that trade unions are quite influential in Norway, and they have especially high impact in organisations using shift work. The trade unions negotiate with the hospitals, and together they decide the terms for how employees can be scheduled to work. St. Olavs hospital also carries traditions for involving the employees in the creation of *schedules*, so that they themselves can influence their own work day. When schedules are created, this is done by hand, which has proven to be very time consuming. Because of this, some wards have found it preferable to create schedules with far longer *planning horizons* than what is commonly discussed in academic literature, making the planning process less frequent. Some schedules span half a year, but as new information arises, the actual

work carried out increasingly deviates from the schedule, putting the effectiveness of the planning process and long schedules into question. A very recent example pointing towards the need for innovating scheduling in health care organisations in Norway is the physician strike in 2016, which included physicians at St. Olavs Hospital. The key conflict resulting in the strike was physicians being dissatisfied with current scheduling policies (Dagsavisen, 2016). The conflict had to be resolved through compulsory arbitration due to the strike seriously affecting the availability of medical services. However, the physicians are not content with the solution and some signal they are willing to leave the public health care sector due to their issues with the scheduling policies (NRK, 2016). As such, working to innovate scheduling is very relevant at St. Olavs Hospital as well as in the rest of Norway.

In 2016 we established a cooperation with St. Olavs hospital looking at how scheduling can be done more efficiently and effective through methods of operations research, focusing on Maternity Ward West.

2.4 Maternity Ward West

In this section we present background information meant to provide insights to Maternity Ward West. We begin this section by briefly presenting the maternity ward and its main functions and challenges. Subsequently we explain how the employees, shifts and demand types are organised to carry out the ward's services. Lastly we present important aspects influencing how the schedules at MWW are created.

2.4.1 Brief overview

The maternity ward's main functions are to help patients at the St. Olavs Hospital with child birth and assist patients and their babies in the aftermath of child birth. The ward also assists rural areas around Trondheim with these functions by stationing employees in Orkdal. The maternity

ward is divided in two neighbouring wards, Maternity Ward East and Maternity Ward West. Together the two wards annually deliver approximately 4000 newborns. The eastern and western maternity wards carry out the same functions and are managed closely together by the same ward manager. Our collaboration with the ward manager is however focused on developing a scheduling model for Maternity Ward West. Therefore the remainder of this section focuses on MWW and describes the current planning environment found there.

Childbirth is inherently difficult to plan for, and the planning environment at MWW differs from those at many other wards at St. Olavs Hospital. The ward manager describes several reasons why the MWW is particularly challenging to create schedules for. The main reasons are presented below:

- Patient arrival is typically highest during day time for most wards, but this is not the case at MWW, where 43% of the patients arrive at night. Women in labour must be treated immediately, so MWW have high staffing requirements also at night time.
- Patient arrival is normally similar during weekends and during weekdays. Thus, the need for staff is also similar, as most tasks at the ward cannot be postponed. Working during weekends is not popular among employees and it is restricted by work regulations. Therefore it can be challenging to staff sufficiently in weekends.
- Night shifts at the MWW are arduous compared with night shifts at wards where patients sleep during the night. The women who give birth need treatment and newborn babies are in need of constant supervision and help, also during nights. An example of this is that all babies must be fed every three hours. Newborn babies are typically not cooperative, which makes this work tedious and employees working nights get little rest during their shifts.
- The strain caused by the arduous shifts may harm employees' health. Therefore it is very important to carefully manage how such shifts are covered.

2.4.2 Employees

MWW is open and staffed all hours all days. Providing health care at the ward requires sufficient staffing by qualified employees at all hours. Also, the hospital has a responsibility to take the employees' needs and preferences into account. The Maternity Ward uses a method close to that of *self-scheduling* when creating the schedules. Self-scheduling puts emphasis on involving the employees in the planning process. This entails that employees' requests, registered as a part of the self-scheduling, should be respected as often as possible. For each employee at MWW there is information essential to the scheduling problem. The most noteworthy information is:

- Contracted Work
- Skill category
- Personal inclinations
- Requests

Contracted work

Contracted work is the number of hours an employee should work during one planning horizon. All employees are scheduled to work their contracted work, but small deviations are allowed. Overtime pay is calculated based on the employee's average weekly workload over the planning horizon. This allows for employees to be scheduled with varying weekly workloads without incurring overtime pay. As such, no employees are scheduled to work overtime. Overtime work only occurs if some disruption forces an employee to work more than scheduled. These scenarios, however, occur strictly when executing the schedule, not when creating it. Employees have varying contracted work and MWW employs a mix of part time workers and full time workers.

Certain employees are affected by a policy called 3-part-average, which involves recalculating their work hours to compensate for arduous *individual schedules*. Eligibility for the 3-part-average requires that an employee:

Background

- works shifts during days, evenings and nights
- works a minimum of 25% of total work during days and evenings
- works one in three *red days* on average

Fulfilling these criteria leads to a recalculation of hours from nightly work and work on red days. Each hour worked during nights is recalculated into 1 hour and 15 minutes and each hour worked on a red day is counted as 1 hour and 10 minutes. Employees with 3-part-average do not receive extra pay due to the recalculation, but reach their contracted work hours quicker, and thus work comparably fewer hours than those without 3-part-average. In practise, it is decided before the schedules are created who is eligible for the 3-part-average.

Skill category

The employees are divided into categories, based on their qualifications, called skill categories. The skill categories and their related main tasks are listed below:

- Assistant managers - tasked with supervising daily activities and contingency planning
- Shift coordinators - tasked with coordinating daily activities across both maternity wards and assisting in functions with high demand.
- Midwives - tasked with child delivery, routine controls for newborns and education for parents.
- Assistant nurses - tasked with packing and disinfecting important equipment, tending to and feeding newborns
- Ward assistants - tasked with cooking food for patients and next of kin.
- Secretaries - tasked with enlisting and discharging patients

At MWW each skill category has a set of different tasks to handle. However, shift coordinators can also cover the tasks for midwives, midwives may cover tasks for assistant nurses and assistant nurses may cover tasks for ward assistants. MWW can generally not hire external employees if they are under-staffed, but they can schedule employees from the neighbour ward to cover certain shifts on certain days and ask employees to cover extra shifts.

Personal inclinations

The personal inclinations can take many forms. Some employees have health issues and cannot work during certain times of the day, e.g. preventing them for covering night shifts. This must be respected. In other cases employees can work during all times of the day, but health issues constrain the amount of work they should do during certain times of the day. There are also employees with an extra job, preventing them from servicing the maternity ward certain days. If the second job is also held at St. Olavs Hospital, consecutive workdays must be counted for both jobs, so as to ensure legal amounts of rest for these employees. Employees commuting from far away are granted certain *shift-patterns* allowing them to sleep at the hospital certain days before working an early shift the next day. There are many more personal inclinations, but the aforementioned examples serve well to illustrate the needs and wishes that affect the scheduling. Personal inclinations will be discussed further in Section 3.3.

Requests

At the maternity ward all employees make one request per day in the planning horizon. The requests are either for a desired work shift or an off-shift. It may occur that employees request shift-patterns that are usually not allocated due to preferred practises, governmental regulations etc. In these cases, the requested shift-patterns can be allocated if they do not contradict governmental *work regulations*. E.g. the shift-pattern Night-Off-Night is generally not allocated to any employees, but if an employee requests a night-shift day 9, an off-shift day 10 and a night-shift day 11, the shift-pattern could be allocated to that employee during those days. Requests

will be discussed further in Section 3.3.

2.4.3 Days

Days are the preferred time increment to describe the scheduling problem at MWW. Some days are far more unpopular to work than others, including weekends, some shifts during red days etc. Therefore, all work that can be expedited or postponed is usually performed during the daytime of regular weekdays. Thus, the demand for midwives and assistant nurses are generally set to be lower during the weekends, while some supporting functions are not in demand during weekends. Furthermore, the policies regarding work during weekends, vacation and unpopular days are quite specific, to ensure that schedules are perceived as fair.

2.4.4 Shifts

Each employee is scheduled to have one shift every day in the planning horizon, and all shifts are either an off-shift, where the employee is off duty, or a work shift. The work shifts take place during three time periods of the day; day time, evening time and night time. Together these three time periods span the entirety of one day, i.e. 24 hours, and the MWW is continuously staffed. The following information related to shifts is especially important:

- Skill categories eligible to work the shift
- When (during a day) the shift starts and ends
- Day of the shift

2.4.5 Demand type

Each work shift must satisfy a demand for different services, and these are categorised into demand types. Each demand type has corresponding skill categories, and only employees within these skill categories can service

Table 2.1: Examples of some of the most usual skill categories and demand types. An 'x' denotes eligibility to cover a demand type. Note how midwives are eligible to pose as assistant nurses, while assistant nurses may act as ward assistants.

		Demand type		
		Midwives' evening-shift	Ass. nurses' night-shift	Ward ass.' shift
Skill category	Midwives	x	x	
	Assistant nurses		x	x
	Ward assistants			x

the demand type. A demand type is considered serviced only if enough employees from the eligible skill categories are allocated to the shift facing the demand type. E.g. the demand for midwives during a day is covered if 6 or more midwives are allocated to the day-shift. However, assistant nurses working day-shifts the same day will not be eligible to cover this demand type, as they are not part of the right skill category. This is illustrated in 2.1. An employee cannot cover two different demand types simultaneously. E.g. a midwife covering an assistant nurse's shift cannot cover a midwife's shift at the same time. The following information related to demand types is especially important:

- Skill categories eligible to cover the demand type
- Demand for a demand type each day

2.4.6 Scheduling at the Maternity Ward

The Ward Manager makes schedules spanning 24 or 27 weeks two times annually for all employees at the ward. It is notable that most demand types have low coverage requirements, making them easy to cover. Many of these demand types are covered by preallocating certain employees to shifts, as described in Section 2.5, i.e. the shifts preallocated in step 1 in

the planning process. Midwives and assistant nurses are, however, needed at all times, due to the demand for their skills being very high. As such, the main scheduling problem solved is allocating midwives and assistant nurses to shifts throughout the planning horizon, while respecting *scheduling rules*.

As the ward uses self-scheduling, it is important to respect the employees' preferences, and it is a policy of high priority at MWW to do so. Policies are guidelines for the scheduling. Preferences cannot always be met and employees accept this. It is however important for the employees to feel that they are being treated fairly. It is therefore a policy at the ward that burdensome tasks and violations of preferences are distributed evenly between the employees. Additionally, it is policy that employees cannot be scheduled to work overtime.

The scheduling rules are based on governmental work regulations, ward policies, preferred practises and agreements with trade unions. When schedules deviate from standard governmental work regulations, this is due to specifications of work regulations between St. Olavs Hospital and trade unions. These specifications relate to cases where the work regulations would prohibit the supply of proficient patient care or contradict employees' interests. Preferred practises encompass rules with the goal of respecting preferences considered universal for the employees at MWW. Many scheduling rules are also adapted to meet individual inclinations employees have. Therefore the scheduling rules may vary quite substantially for different employees.

2.5 The planning process

Maternity Ward West has a planning process similar to self-scheduling, which is described in more detail in Section 3.1. Schedules are *acyclical* with a planning horizon of 24 to 27 weeks. The time from the planning process starts until the schedule is published is approximately 11 weeks. The 11 weeks consist of 3 weeks of collecting information, 6 weeks of scheduling and

2 weeks where trade unions evaluate the schedule. At the time the schedule is published, the schedule will take effect in approximately 4 weeks. During the time spent on the planning process all entities (employees, groups, unions, software) are involved in the scheduling in various manners, before the ward manager finalises the schedules. A rough description of this planning process is presented in the paragraphs below and illustrated in Figure 2.1. It should be noted that in reality, the information flows much more freely between entities than described below. However, we focus on describing the information flow that directly results in the final schedules at MWW.

Step 1) Registering preallocated shifts:

11 weeks before the schedule is published

The planning process begins with planning vacations and the variety of shifts that are to be preallocated when solving the scheduling problem. These are delivered by multiple sources, including policy makers, the ward manager etc. The preallocated shifts are shifts such as single days when employees attend classes (competence days), coordinating shifts, etc.

Step 2) Registration of requests:

11 weeks before the schedule is published

Employees register their preferred individual schedules for the entire planning horizon in a simple database-system, primarily based on Microsoft Excel, called the *Resource Management System*(RMS). Their requested individual schedules may consist of requests for any existing kind of shifts (including off-shifts), as well as comments elaborating on their choices of requests. Any employee may register a request for every day in the planning horizon.

Step 3) Retrieval of requests:

8 weeks before the schedule is published

After the employees have registered all their requests and other personal information, the ward manager retrieves the information from the RMS.

Step 4) Sending draft schedule:

8 weeks before the schedule is published

The ward manager checks the validity of the requests at the internal RMS and change potential false registrations of requests, e.g. registering for a competence class on a day without such classes being held. The ward manager then aggregate the individual schedules into a draft schedule containing all employees' requests. The ward manager also decides how to allocate vacations and extra off-days due to seniority. The ward manager then removes some work shifts on days with overcapacity for employees who have requested more work than their contracts suggest they should have. The draft schedule will not at this point cover the demand for all shifts throughout the planning horizon, and must thus be processed significantly in order to become feasible.

Step 5) Delivering processed draft schedule:

7.5 weeks before the scheduling is complete

A scheduling group consisting of employees from both maternity wards at St. Olavs Hospital help the ward manager by processing the draft schedule. The ward manager is only part of this step as a supervisor. The scheduling group processes the draft schedule by changing and shuffling shifts in the aggregated schedule, trying to create a feasible

schedule while ensuring employees are awarded as many requests as possible and ensuring some perception of fairness for all employees. The goal is not reaching a final result, but to improve the draft schedule and make it ready for the subsequent bartering process.

Step 6) Bartering process:

7 weeks before the schedule is published

The employees begin a process of bartering with each other with the goal of making the schedule processed in step 5 comply with daily needed staffing levels and other scheduling rules, i.e. creating a feasible solution to the scheduling problem. To do this, employees negotiate informally to distribute uncovered shifts while trying to satisfy their own personal needs and preferences.

Step 7) Sending bartered schedule:

4 weeks before the schedule is published

When the negotiations in the bartering process stagnate, the bartered schedule is sent back to the scheduling group, which must process the schedule further.

Step 8) Delivering finalised schedule:

2 weeks before the schedule is published

The scheduling group makes final changes to the schedule to ensure it is feasible and legal to use. Often this involves contacting employees and negotiating or informing them about decisions that are made in the scheduling group. This can be an iterative process between the scheduling group and some employees. When the schedule is finalised,

it is sent to the trade union representative for assessment. If any errors are found, the schedule is corrected by the ward manager and re-assessed by the trade union representative.

Step 9) Publishing schedule:

The schedule is published and takes effect in 4 weeks

Once the final schedule is verified by the trade union representative it is made public four weeks ahead of taking effect. When the schedule is in effect, shifts are swapped internally by employees. Also, the assistant manager finds employees to work shifts when staff is absent, e.g. due to sudden sickness.

In Figure 2.1 we illustrate the planning process described above in a simplified manner. In reality the ward manager, scheduling group and employees communicate more than we depict in Figure 2.1. We only present the main flow of information between all the entities to make the planning process visually more understandable. In Figure 2.1 these entities are shown as blue geometries while arrows labelled with numbers represent the steps in the planning process. Arrows point towards the main entity receiving new information when a step is completed, and point from the main source of this information. The arrow labelled 1 points from no source, because the information attained in step 1 is generated by several entities, many of which are not involved in the planning process in any other way. The arrow labelled 9 points towards no entity, because step 9 involves publishing the finalised schedule, so the information is given to everyone.




Figure 2.1: Rough presentation of the current planning process, that visualises the information flow that directly results in the final schedules. In reality, much information also flows through other channels than depicted.

Chapter 3

Literature

In this chapter, relevant literature is provided to the reader and discussed in relation to the scheduling problem at Maternity Ward West. The discussed literature is chosen because of its relevance to provide the reader with a theoretical context for the planning process and to serve as useful information before designing a mathematical model to solve the scheduling problem at MWW. Furthermore, some of the work in this chapter is based on the work done in our specialisation project 'Telling a success story of nurse scheduling - Model development and analysis', Beckmann and Klyve (2016). Thus parts of this chapter reproduces some of the work presented there. Also, please note that it is assumed that the reader is familiar with basic principles from operations research and optimisation.

Our selection of literature was decided upon by reading articles recommended to us by our supervisors and further obtaining referenced articles that seemed relevant as well as searching for other articles related to nurse scheduling using Google Scholar. The articles Bergh et al. (2013) and Burke et al. (2004) stood out as inspiring sources for locating relevant literature.

Firstly in this chapter we explain some useful concepts in Section 3.1. In Section 3.2 we position the scheduling problem in related literature and

in Section 3.3 we present some key aspects, to provide the reader with a theoretical context for the planning process. Lastly we discuss the most relevant optimisation approaches for our model in Section 3.4.

3.1 Nurse scheduling and OR Terminology

In related literature many different terms are discussed. We find it useful to define these to ease the understanding of discussions carried out later in this thesis. The terms are written in italic style the first time they occur later in the text.

3.1.1 Administrative modes

Burke et al. (2004) describes different administrative modes, where the schedules of nurses are created from a bottom-up approach to a top-down approach. We define three different levels of administrative modes, where the definition of self-scheduling from Section 2.1 is repeated for the reader's convenience:

- *Centralised scheduling* describes the situation where one administrative department in a hospital carries out all the personnel scheduling (Burke et al., 2004). It represents a top-down approach to scheduling. The advantage of this is that health workers can focus on health care, instead of using valuable time on scheduling and costs can be controlled from the top. An inconvenience is that nurses' preferences and requests are very hard to take into account.
- *Unit scheduling* is when a head nurse or unit manager is responsible for creating the schedules locally for their ward. This is more time consuming than centralised scheduling and has a limited capacity with regards to capturing and respecting all the requests made by co-workers.
- *Self-scheduling* is a system in which nurses create their own schedules, then barter amongst themselves until all requirements regarding staff

coverage is fulfilled. This approach is popular amongst nurses, as it lets them influence the schedules with their own preferences greatly. However, the self-scheduling process is very time consuming, and in some cases the negotiations cause conflicts amongst the nurses.

Traditionally, there is a trade-off between time spent on creating schedules and the ability to comply with nurse preferences, i.e. much time is needed to manually create schedules of high quality. However, using the MWW scheduling model makes it possible to combine some of the strengths of the different administrative modes. The individual influence of self-scheduling is preserved in the MWW scheduling model, while the efficiency is increased by producing the schedule at a unit level. Furthermore, the introduction of the model makes data related to the scheduling problem much more available to the Hospital Management, facilitating informed tactical and strategic decisions at a centralised level.

3.1.2 Useful terms

Below are some terms that are useful when presenting related literature. We sort them alphabetically and define them below. Note that these expressions appear in *italic* type the first time used in the text.

- *Consumable resources* - non-durable medical supplies that cannot withstand repeated use by more than one individual, e.g. bandages.
- *Continuity in schedules* - refers to unbroken and consistent work for a nurse over a period of time, e.g. a nurse working four consecutive day shifts.
- *Hard constraints* - constraints which must be upheld, e.g. governmental work regulations.
- *Opportunity costs* - the loss of potential gain from other alternatives when one alternative is chosen.
- *Renewable equipment* - equipment that does not perish after use, e.g. beds.

- *Shift-transition* - occurs when a nurse works a shift during other hours than the day before, e.g. working a day shift then an evening shift the next day would incur a shift transition.
- *Soft constraints* - constraints which are desirable to uphold, but may be broken for a penalty if needed. Constraints regarding nurses' preferences are typical examples of soft constraints.
- *Sunk costs* - costs that have already incurred and thus cannot be recovered.
- *Triaging* - deciding the order of patient treatment.

3.2 Positioning in the literature

In this section we introduce a useful framework and position relevant theory in it. We also bring up some important terms from related literature, to provide context for the planning process.

3.2.1 A fitting framework

To position the planning process at Maternity Ward West in the literature, Figure 3.1 from Hans et al. (2011) is presented. It describes the different hierarchical planning levels and managerial areas of organisations. There exists several frameworks with significant relevance to the planning process at MWW, and hierarchical decomposition is very common in the field of manufacturing planning and control. Examples of such frameworks are the ones presented in Zijm (2000) and Vissers et al. (2001). The framework developed and used in Hans et al. (2011) is an expansion of typical frameworks used in manufacturing planning and control, adapted for health care organisations. The different levels on the two axes will be explained below, before we position the planning process at MWW into the framework.

	Medical planning	Resource capacity planning	Materials planning	Financial planning	
Strategic	Research, development of medical protocols	Case mix planning, capacity dimensioning, workforce planning	Supply chain and warehouse design	Investment plans, contracting with insurance companies	↑ hierarchical decomposition ↓
Tactical	Treatment selection, protocol selection	Block planning, staffing, admission planning	Supplier selection, tendering	Budget and cost allocation	
Offline operational	Diagnosis and planning of an individual treatment	Appointment scheduling, workforce scheduling	Materials purchasing, determining order sizes	DRG billing, cash flow analysis	
Online operational	Triage, diagnosing emergencies and complications	Monitoring, emergency coordination	Rush ordering, inventory replenishing	Billing complications and changes	
	← managerial areas →				

Figure 3.1: Managerial areas vs hierarchical decomposition

3.2.2 Managerial areas

The different managerial areas of an organisation can be regarded as the different functions of the organisation. They can be divided into different areas. However, Hans et al. (2011) argues that the division shown below is favourable as they consider them to be relevant for all research projects they have done on optimisation of health care operations.

The managerial areas are divided into the four categories below:

- *Medical planning.* Medical Planning includes decisions regarding e.g. medical protocols, treatments, diagnoses and *triaging*. More autonomy is needed for the health workers the more complex and unpredictable the health care process is.
- *Resource capacity planning.* Resource capacity planning includes dimensioning, planning, scheduling, monitoring and control of *renewable equipment* and facilities, like instruments, operating rooms and staff.
- *Materials planning.* Materials planning includes acquisition, storage, distribution and retrieval of *consumable resources*, like prostheses, blood, bandages, food, etc.

- *Financial planning.* Financial planning deals with how organisations should manage cost and revenues, with the regards to its objectives.

3.2.3 Hierarchical decomposition

Hierarchical decomposition is done to show how "decision making disaggregates as time progresses and information gradually becomes available" (Hans et al., 2011). Classically the decomposition comprises the strategic, tactical and operational levels. Here, however, the operational level is further decomposed into offline operational and online operational to better categorise how operational planning is done in real life scenarios. We make use of the definitions provided in Hans et al. (2011) to explain the hierarchical levels.

- The strategic level has a long planning horizon and revolves around the structure of an organisation. It involves defining the organisation's missions, and making decisions to translate this mission into the design, dimensioning, and development of the health care delivery process. Examples of such decision areas are developing and implementing new medical protocols and mergers of nursing homes.
- The tactical level addresses the organisation of the operations/execution of the health care delivery process. In this way, it is similar to operational planning. However, decisions are made with a longer planning horizon. Examples of decision areas are deciding staffing levels at wards.
- The operational level involves the short-term decision making related to the execution of the health care delivery process. The flexibility on this planning level is low as the higher levels has already set the scope for operational decision making. Furthermore, the operational level is divided into two categories:
 - Offline operational planning. All offline operational planning can be planned for ahead of incidents occurring. Examples are

treatment selection, nurse scheduling and inventory replenishment ordering.

- Online operational planning. Online operational planning is done after sudden changes in circumstances and encompass triaging, finding substitute workers and replenishing depleted supplies etc.

3.2.4 Personnel scheduling

As the term personnel scheduling encompasses many kinds of scheduling problems, we wish to place it in the framework provided by Hans et al. (2011). Personnel scheduling seeks to create schedules for staff so that an organisation can satisfy the demand for its products or services. This firstly entails determining the number of staff, work policies and required skills needed to comply with wanted service levels (Ernst et al., 2004). Individuals are then allocated work to satisfy the required staffing levels at different times. Personnel scheduling thus spans operational and tactical planning, but also some parts of strategic planning, as policies regarding the workforce are a part of personnel scheduling. Online operational planning is also a part of personnel scheduling, as dealing with contingencies like sudden illness must be considered. Personnel scheduling is best described by the column *resource capacity planning* in Figure 3.1, keeping in mind that the workforce planning is considered in the strategic level. Our categorisations are illustrated in Figure 3.2.

The offline operational planning of personnel scheduling is the most studied aspect in the discipline of nurse scheduling, and this aspect is what is usually referred to as nurse scheduling. In most nurse scheduling problems, staffing is assumed to be fixed by tactical and strategic decisions, so that creating schedules is the main focus. With this in mind, nurse scheduling, as we use the term, fits best into the intersection between offline operational planning and resource capacity planning as seen in Figure 3.2.




Figure 3.2: Managerial areas vs hierarchical decomposition. Personnel scheduling, staffing and scheduling positioned.

3.2.5 Positioning the planning process at MWW

In Chapter 1 of this thesis, we implied that the most fitting theoretical characterisation of the planning process at MWW was nurse scheduling, also commonly referred to as nurse rostering. This is despite the fact that employees at the maternity ward prefers being titled by their professions, e.g. midwives, due to a difference in education from nurses. However, the description of the planning process at MWW resembles the description of nurse scheduling greatly, with staffing assumed fixed and focus on the offline operational planning. MWW’s planning process thus falls into the same block as most nurse scheduling-problems. However, it is useful to explore limits of the framework provided by Hans et al. (2011).

Hans et al. (2011) argues that many organisations fail to properly consider how planning in the different hierarchical levels affect each other. This is relevant for the planning process at MWW as well. Even though it is best positioned with scheduling as seen in Figure 3.2, it is important to understand that both decisions made in the strategic and tactical levels affect how the operational planning is done. Examples of this is the number of nurses required to cover each shift. It can be considered a tactical decision to decide whether or not to use excess staffing during night shifts and a strategic decision to create policies that prioritise employees' health before costs. Furthermore, online operational planning will in reality affect the real life planning process. E.g. sudden sickness may occur and reduce the availability of the staff, or surprisingly many women could give birth at the same time and increase demand for medical service beyond what was expected. This can to some extent be accounted for, e.g. by over-staffing. Thus in reality, the planning process at MWW cannot be regarded as exclusively belonging to the offline operational hierarchical level, although the general characteristics of the process fits that level well.

3.3 Key aspects within nurse scheduling

Some aspects of nurse scheduling are especially important in order to create the best possible schedules. We therefore study literature on three aspects that have shown to be key in earlier work presented in literature and in our specialisation project Beckmann and Klyve (2016). These aspects are employees' preferences, fairness of schedules and costs inflicted by the schedules. We study the key aspects and discuss the extent to which they are accounted for in our selection of related literature. The study is done to create a basis for our understanding of the scheduling problem at MWW.

3.3.1 Preferences

When discussing the planning problem, we learnt that it was practical to categorise the different preferences the employees have towards schedules.

We have therefore divided the preference-term into three categories:

- *Preferred practises* are rules concerning employees' preferences towards working certain shifts, days or combinations of such. An example is that policy dictates the number of night shifts each employee should work given contracted working hours and position, if no personal reason contradicts the policy. Another example is that employees wish to work either the entire weekend or have the entire weekend off.
- *Personal inclinations* are wishes and needs specific for each employee, that remain the same throughout planning horizons. An example is a preference against working Wednesday evenings, due to some personal reason. Another example could be preferring not to work night shifts due to health reasons.
- *Individual requests* regarding specific shifts and off-days. These preferences are both employee and time specific in that an employee makes a request for a given day, e.g. asking for an off-day. These individual requests constitute an employee's original individual schedule in classic self-scheduling.

Preferred practises

Whether modelled as hard or soft, many constraints undoubtedly deal with what we have categorised as preferred practises in related literature. E.g. by ensuring employees have enough off-days to rest. Examples of constraints dealing with a maximum number of allowed consecutive workdays can be found in Rönnerberg and Larsson (2010), Ruzzakiah et al. (2011), Rönnerberg et al. (2012), Azaiez and Sharif (2005), Bard and Purnomo (2005), Vaz and Moz (2007) and Bester et al. (2007). Similar constraints also exist for specific shifts, like night shifts. Ruzzakiah et al. (2011), Rönnerberg et al. (2012) and Azaiez and Sharif (2005) deal with this. This was also the case in the model we analysed in Beckmann and Klyve (2016), where both a maximum number of consecutive shifts and consecutive night

shifts were constrained.

In related literature we find many examples of constraints regarding patterns of shifts. Many shifts and shift-patterns are modelled as unpopular, i.e. they are constrained with either hard or *soft constraints* for all employees throughout the planning horizon. An example of an unpopular shift-pattern is having a night shift followed by an early shift the next day, i.e. the 'night-day' shift-pattern. In Bard and Purnomo (2005), Rönnberg et al. (2012) and Vaz and Moz (2007) the 'night-day' shift-pattern is illegal, due to a mandatory waiting period between shifts, and is thus modelled with *hard constraints*. In Ruzzakiah et al. (2011) the 'night-day' shift-pattern is legal, but avoided by penalising the violation of soft constraints. In Beckmann and Klyve (2016) there is a large number of hard constraints dictating both illegal and sometimes mandatory work-patterns. There are also some soft constraints penalising the use of certain shift-patterns.

We find that several models include constraints specific to the weekends, and the authors often argue that policies regarding the weekends are important to the employees. In Rönnberg and Larsson (2010) employees can either work the entire weekend or be off duty. If the employees are working the weekend, the employee must also work Friday evening or Monday morning. Similar constraints can also be found in Bester et al. (2007), Beckmann and Klyve (2016) and Rönnberg et al. (2012). The latter also includes constraints so that only employees working the full weekend can cover the preceding Friday evening shift. This ensures that once a weekend off is given, the employee is also off duty early on Friday, so that more consecutive leisure time is ensured.

Employees often desire *continuity in schedules*. As such, *shift-transitions* are considered undesirable, but not illegal, in many articles. An example of this is the 'day-night' shift-pattern, which is avoided in Ruzzakiah et al. (2011) and Azaiez and Sharif (2005) through soft constraints. To ensure compliance with preferences on continuity, Bard and Purnomo (2005) sets an upper limit to the number of weekly shift-transitions for each employee's

personal schedule, while Beckmann and Klyve (2016) penalises working day, evening and night during the same week. Other typical policies involving continuity, deal with the allocation of off-days. An isolated workday between two off-days, i.e. the shift-pattern 'off-on-off', and conversely the 'on-off-on' shift-pattern is avoided respectively in Ruzzakiah et al. (2011), Beckmann and Klyve (2016) and Rönnerberg et al. (2012). Rönnerberg et al. (2012) also avoids the 'evening-off' shift-pattern, while Ruzzakiah et al. (2011) constrains a pattern to ensure three off-days are given to an employee after consecutively working three night shifts. In Azaiez and Sharif (2005) and Bard and Purnomo (2005) both the 'off-on-off' shift-pattern and 'on-off-on' shift-patterns are avoided. Additionally, Azaiez and Sharif (2005) avoids the 'day-night-off' shift-pattern. These constraints reflect how employees tend to prefer continuity in both their workdays and their off-days.

Night shifts are often considered unpopular as it affects employees' circadian rhythms. In Azaiez and Sharif (2005) moderating this is dealt with by trying to ensure that more day shifts than night shifts are given to each employee, in a schedule.

When discussing preferred practises with the ward manager at MWW, our impression was that many of the above mentioned types of constraints are present in their real-life scheduling problem. When to work night shifts and weekends were indeed of great importance to the employees, as the literature implied. However, although the ward manager at MWW affirmed that continuity was an advantage for a schedule, she has a quite different view of what a continuous schedule was than the authors of most articles. Instead of avoiding shift-transitions or working at different hours throughout a week, she stressed the importance of shift-transitions that follow the circadian rhythm of the employees. E.g. 5 day shifts in a row was not necessarily seen as a better shift pattern than working a shift-pattern of 'day-day-evening-evening-night'. Furthermore, she explained that an employee always will strive to retain a normal circadian rhythm when having day, evening or off-shifts, making shift-transitions to and from night shifts

very important. Although Azaiez and Sharif (2005) also focused somewhat on night shifts' effect on the circadian rhythm, the approach to scheduling shift-transitions at MWW is different from what we have come across in related literature, and thus especially interesting. The different approach to scheduling night shifts at MWW could be due to the fact that there is a large demand for staff during the night relative to most other wards, as mentioned in Section 2.4. This finding is very useful when designing the MWW scheduling model in Chapter 5.

Personal inclinations

Vaz and Moz (2007) keeps track of which shifts the individual employee does not prefer, and include penalties in the objective function whenever an employee is given such a shift throughout the planning horizon. Akbari et al. (2012) lets employees rank 6 kinds of shifts as preferred, where the 6 preferences have different priorities, and use soft constraints to implement the preferences. Maenhout and Vanhoucke (2013) registers which wards each employee prefers working in, and restrains the model from allocating an employee to a different ward, using soft constraints. Conversely, awarding the allocation of preferred shifts or preferred job types can also be done, as in Rönnberg et al. (2012), Christiansen et al. (2015), Beckmann and Klyve (2016) and Vaz and Moz (2007), respectively. In Beckmann and Klyve (2016) employees may also decide for themselves if they want 8-hour shifts every second weekend or 12-hour shifts every third weekend.

When the Ward Manager creates schedules at the MWW, personal inclinations are of great focus. The Ward Manager knows each employee and their individual needs. E.g. each employee's health is considered very carefully before allocating shifts. This can translate to a variety of constraints when designing the MWW scheduling model in Chapter 5 and is likely to pose a challenge. Another interesting issue the ward manager raised was that some employees are commuters and thus like to work an 'evening-day' shift-pattern, spending the night at the hospital between the shifts.

Individual requests regarding specific shifts and days

Rönnerberg and Larsson (2010) deals with requests by introducing requests with different hierarchical priorities, namely vetoes, hard requests and soft requests, in descending order of priority. These can either be requests for working certain shifts or requests to avoid working certain shifts. The vetoes are modelled as hard constraints, and in cases where the model can't comply with a veto, the head employee decides what to do. The strong requests have heavier weights than the normal requests in the objective function, to reflect their higher priority. No limit is set on the number of requests. However the more requests that are made, the less relative weight they carry, lessening the probability of each individual request being granted. Other papers deal with only one level of priority for requests, such as Bard and Purnomo (2005), Maenhout and Vanhoucke (2013) and Bester et al. (2007). These authors include employees' requests to work certain shifts or have certain days off in their models, and penalties incur whenever a request is violated. The cyclic scheduling model in Vaz and Moz (2007) works in a similar fashion, but additionally allows for requesting sequences of shifts.

As Rönnerberg and Larsson (2010) deals with automating an existing self-scheduling system, it is obvious that a great deal of individual requests need to be respected to avoid radical changes in the employees' lives. Furthermore, as creating schedules using OR arguably is a top-down initiative, self-scheduling seems like somewhat of a contrast. Thus, it is only natural that there exists less literature about OR-models that deal with individual requests that are unique for certain shifts and certain days. However, it is reasonable to assume that the MWW scheduling model, which will substitute a self-scheduling system, should indeed take into account employees' requests. However, with the existing RMS, the importance of each request can generally not be determined. This makes it hard to decide on the relative importance of registered requests.

General notions

It is notable how all articles we have read that present mathematical models for nurse scheduling, take into account preferred practises related to shifts and days in some way or another, while fewer deal with personal inclinations and individual requests for specific shifts and days. However, it is not very surprising, as recovering data about all employees can be costly and time-consuming. Furthermore, to what degree a model takes preferences into account may very well reflect the administrative modes that are most prominent at a hospital. If OR-based nurse scheduling is implemented at a hospital with a high degree of *centralised scheduling*, it would seem natural to continue creating schedules without taking into account individual preferences. Conversely, implementing a system to replace self-scheduling without taking individual preferences into account would involve very big changes to employees' everyday life.

As MWW is currently practising a version of self-scheduling, the MWW scheduling model should have a high focus on employees' preferences, given the information found in related literature as well as our experiences from our specialisation project. This is also consistent with the fact that trade unions have a large influence on the way the ward is run. It seems that focusing on preferences when creating schedules is a *hygiene factor* for an acceptable schedule.

3.3.2 Fairness

Fairness is of high importance when creating a schedule. As put by Bard and Purnomo (2005) "A critical measure of success in the use of preference scheduling is the perceived fairness or balance in the posted rosters." Fairness is, however, a somewhat ambiguous term. It is thus useful to explore the term itself shortly and then how related literature deals with the aspect in different problems.

The fairness-term

The aspect of fairness is wide and complex. However, most models include constraints that are obviously concerned with it. Most literature regard fairness as the effort of complying with preferences on an equal level for all the nurses. In Warner (1976), an early article on the subject of nurse scheduling, it is stated that rotations of shift-patterns and the time between undesirable shifts should be small for each nurse, so that such preferences are divided fairly. Warner (1976) also make sure there is a rotation between working and having off-days during weekends, in order to secure fairness. Evening out the allocation of different desirable and burdensome shifts has become quite standard over the years, and can be seen in a lot of related literature. Nonetheless, there are exceptions to this view, as there is no objective way to say state what is truly a fair schedule. It may depend on several aspects, such as the culture of the workplace. In Akbari et al. (2012) it is stated that "Our model objective tries to maximise preference of part time workers by a minimisation objective while considering seniority, availability, and priority of employees." Akbari et al. (2012) thus considers fairness in a slightly different matter, where rank and seniority is included in the considerations of who's requests to prioritise.

Fairness in related literature

More recent literature typically utilises constraints that deal with fairness in several ways. One popular technique is to enforce a maximum allowed difference between the highest and the lowest number of preferences granted each nurse. Ruzzakiah et al. (2011) and Azaiez and Sharif (2005) makes sure each nurse's total workload contains more than a certain share of unpopular shifts. This evens out the distribution of these shifts and makes sure no one is allocated many unpopular shifts, while others get few or none.

In Beckmann and Klyve (2016), hard constraints exist to ensure a maximum allowed difference in the number of unpopular shifts each nurse is given. Furthermore, in Rönnerberg and Larsson (2010) and Rönnerberg et al.

(2012) a preference score is calculated for each nurse, based on the number of respected preferences. The lowest amongst all the scores is maximised in the objective function, improving the quality of the the worst personal schedule, and thus evening out the compliance rate of preferences to increase fairness. Additionally, their model makes sure all nurses have the same number of unpopular shifts, scaled by the nurses' contracted working hours. Such a scaling of unpopular shifts relative to contracted working hours is also done in Azaiez and Sharif (2005).

Maenhout and Vanhoucke (2013) makes sure that distribution of the workload is divided evenly with respect to the number of working hours, duties, assignments per shift type and weekends. This is ensured using preference thresholds, that penalise the objective function if a nurse has fewer or more preferences granted than the interval allows. Bard and Purnomo (2005) models fairness as how closely a nurse's schedule matches the nurse's contractual agreements, and by the severity of the preference violations in the derived schedule. Bard and Purnomo (2005) furthermore creates maximum and minimum-limits for the standard deviation of a variable including those aspects. Bester et al. (2007) minimises the nurses' accumulated dissatisfaction, accounting for both current and previous shift assignments, to ensure fairness in both the short and long run. Azaiez and Sharif (2005) sets a minimum number of night shifts for all nurses, to ensure they are evenly distributed. In Beckmann and Klyve (2016) ungranted individual requests are penalised increasingly if fewer requests are granted. Also, the number of personal inclinations fulfilled, for the employee with the fewest personal inclinations fulfilled out of all employees, are penalised. This evens out the distribution of respected personal inclinations between employees.

Fairness at MWW

When discussing fairness with the ward manager at MWW, she lays forward some important principles for creating schedules she perceives as fair. Firstly and most importantly, weekend shifts and night shifts should be distributed in a fair way, as these shifts are highly unpopular. In general

this means dividing those shifts equally between all employees, according to the ward manager. At MWW there is a mutual agreement between employees, trade unions and the management stating how many weekend shifts and night shifts each employee should work. Each employee works approximately equal numbers of weekend shifts, although minor variations may occur due to differences in kinds of positions etc, which corresponds to the standard view on fairness in related literature. The differences in night shifts worked are notably larger, but this is primarily due to health issues. The ward manager states that all employees' preferences are equally important to them, but states that fairness implies equality in respect and influence, not necessarily equality in the individual schedules, as people have different health issues and needs. Although these principles seem to be upheld well when creating manual schedules at MWW, she admits the current bartering process, denoted step 6 in Section 2.5, in practise creates winners and losers. Some employees seem to always get their way when bartering, while others give in to co-workers much more easily and the ward manager has explicitly stated that the bartering process is unfair.

3.3.3 Costs

Costs seems like the most unmistakable of the three aspects focused on in this section. However, the costs are not as easily represented in a mathematical model as one might think. Most nurse scheduling models aim to reduce labour costs, but they will typically also focus on respecting regulations and nurse preferences. Therefore, although nurse scheduling models to some extent model the costs, the objective function tend to be adjusted greatly to take into account other priorities, making costs hard to identify in models. Furthermore, *sunk costs* are usually not included in models, as including constant terms in the objective function does not affect any variables. Instead *opportunity costs* are represented. Furthermore, respecting preferences is standard, and may contribute to a more effective workforce, and thus indirect cost savings (Burke et al., 2004). Furthermore, one can argue that scheduling models focusing on the offline operational level can contribute to indirect cost savings in other hierarchical levels as well. E.g.

it is reasonable to assume that if employees experience significantly better schedules, this can reduce sickness leave and increase efficiency. Also, using scheduling models may contribute greatly to understanding the scheduling problem at a ward, giving the managers better information as to what is an optimal staff level and the implications of changes in policies etc.

At MWW, the staffing level is set to always be sufficient for covering demand at all times, and they do not accept scheduling more work than contracted for any employee, as this triggers overtime pay. Thus, in an offline operational perspective, all labour costs are sunk, as employees are entitled to their contracted salaries, never anything more or less. When they create schedules at MWW today, they generally only consider costs in an offline perspective, making preferences and fairness the major quality aspects of their scheduling problem. Staffing costs seem to be a given entity when solving the problem, while respecting preferences and ensuring a sense of fairness is what decides the quality of the schedules. However, if we include the staffing decision in the problem, the salaries of the nurses are among the decisions for the MWW scheduling model to make. In this case, salaries are among the direct costs of the scheduling problem. Similarly, if the scheduling model includes an online operational focus, factors such as distributing excess staff in an optimal way will reduce costs inflicted by sudden illness or a sudden and large increase in demand for employees. When creating the MWW scheduling model, the impacts of the different hierarchical levels should be considered carefully.

3.4 Optimisation approaches

In broad terms, solution approaches for nurse scheduling can be divided in two categories: manual and automated processes. The manual processes are still widely popular (Rönnerberg et al., 2012), and the one most discussed in academic literature is self-scheduling. There exists computer tools assisting this manual process, but these are mostly commercial software, with little coverage by academic papers, at least to the best of our knowledge.

Aside from manual scheduling, there are several automatic solution methods discussed in literature on OR. These methods can be divided in two broad categories: exact and heuristic approaches (Burke et al., 2004). Exact approaches seek to find optimal solutions given some chosen goals, e.g. optimal with regards to costs while respecting other necessary constraints. Heuristic approaches can generate, but do not guarantee, optimal solutions. Rather, they are concerned with quickly generating good solutions. Heuristic methods have become increasingly popular in later years as finding optimal solutions to real life nurse scheduling problems has proven to be extremely difficult (Ernst et al., 2004). Burke et al. (2004) discuss several solution methods to nurse rostering and state that with regards to implementability and capability of solving real life problems, heuristic approaches outperform exact methods, like optimisation with linear programming. In this section we will briefly discuss solution approaches typically found in Operations Research-literature on nurse scheduling.

3.4.1 Exact methods

Integer linear programming (ILP) is suitable for finding optimal solutions, but is limited by the huge and complex search spaces found in real life nurse scheduling (Burke et al., 2004). Furthermore, ILP is typically solved by finding the linear programming (LP) solution before branching on variables to find integer solutions. This can be very time-consuming. "These approaches [mathematical programming approaches] are more commonly applied to simplified versions of the real world rostering problem or where there are few complications in the original problem." (Ernst et al., 2004). However, such simplifications tend to make the models less able to tackle the complexity of real life scheduling. This trade-off is essential when creating an ILP nurse scheduling model. Some researchers aim to reduce run time using decomposition and column generation. Typically, each nurse's schedule is created in a sub-problem, while a master problem combines the best proposed individual schedules (Ernst et al., 2004). This can be both an exact method and a heuristic method depending on how the sub-problem and master problem is solved. In Bard and Purnomo (2005) nurses' individ-

ual schedules are created using a heuristic. The main advantage of column generation for nurse scheduling is that it utilises the fact that creating optimal individual schedules is closely related to creating optimal schedules for all nurses. By exploiting this nice structure, the run time could potentially be reduced significantly. On the other hand, the model quickly becomes less flexible to changes in the scheduling environment, compared to models that are not decomposed (Range, 2016). Another technique, called goal-programming, has a particular focus on being able to prioritise several aspects in one model. It defines a target level for each goal, together with an associated priority. The aim of the method is to find a solution as close as possible to each goal, based on the goal's relative priority (Burke et al., 2004).

3.4.2 Heuristic approaches

The size of the nurse scheduling problem, together with the lack of knowledge regarding the problem structure, has hindered the applicability of exact methods. In their place, many heuristics have been utilised and perform well with regards to schedule quality and run time. One popular heuristic method is metaheuristics. Burke et al. (2004) believes metaheuristics are generally better suited than most other approaches for generating an acceptable solution in cases where the constraint load is extremely high and indeed in cases where feasible solutions are very difficult to find. "The practical advantage of metaheuristics lies in both their effectiveness and their general applicability. The effectiveness lies in the production of reasonably good feasible solutions within a limited amount of running time, whereas mathematical programming techniques run the risk of not returning a feasible solution for a long time. However, using metaheuristics also results in a number of drawbacks, since they cannot demonstrably produce optimal solutions nor can they demonstrably reduce the search space" (Bergh et al., 2013). Simulated annealing, tabu-search and genetic algorithms are amongst the most commonly discussed metaheuristics in OR-papers. Furthermore, Burke et al. (2004) argues that hybrid approaches combining metaheuristics with other solution approaches are very promising.

3.4.3 Implementation and optimality

Even though several automatic methods exist and a plenitude of research has been done on nurse scheduling, these methods and applications are rarely found implemented and used in real life nurse scheduling. Kellogg and Walczak (2007) discusses several potential reasons for this gap between academia and real-life application. One being that researchers historically has oversimplified the problem, e.g. by assuming the nurses in a ward can be seen as homogeneous, failing to model how highly constrained real world nurse rostering is. The resulting scheduling models have therefore had a too narrow scope for them to be practical to implement. Another reason is that automating the scheduling has been met with scepticism by nurses fearing to lose their impact on the planning process. Several research papers talk about the importance of closing the gap between theory and practice, including Burke et al. (2004) and Kellogg and Walczak (2007), suggesting new directions research could take to become more relevant for real life application. Amongst their suggestions are abandoning the search for "optimal" solutions and rather pursue heuristic approaches.

"For most real problems, the goal of finding the 'optimal' solution is not only completely infeasible, it is also largely meaningless. Hospital administrators want to quickly generate a high quality schedule that satisfies all hard constraints and as many of a wide range of soft constraints as possible." (Burke et al., 2004)

Creating optimal schedules is of limited use for hospitals, or at least the nurses, if model implementation is impossible. Rather than focusing on finding optimal solutions, research can focus on creating methods which minimise the efforts of creating the schedules and that prioritise improving the job satisfaction of the employees. This would not only make the schedules more desirable in the eyes of the nurses, but also makes it easier to create methods with better applicability in real world nurse scheduling. Even if optimal solutions were found, these are heavily reliant on higher-level decisions (i.e. the strategical and tactical levels). E.g. the relative

importance of cost versus granting certain personnel requests makes the optimal schedule rather ambiguous.

3.4.4 Optimisation at MWW

Two above mentioned factors that seem very likely to affect the MWW scheduling model when choosing between optimisation approaches are the run time and the flexibility of the model. As they usually plan for 24 or 27 weeks at MWW, the amount of data is huge, making the process of solving the scheduling problem very challenging. Thus, the run time should be considered before deciding on optimisation method. Furthermore, the flexibility to make continuous changes to the MWW scheduling model when needed seems vital to us. This is both due to our experiences during the specialisation project and how we observe that information continuously changes and evolves at the ward.

Furthermore, the notion of abandoning the search for theoretically optimal solutions and rather pursue feasible high quality solutions, fits our impression of the needs at MWW. The ward manager believes that schedules using the current planning process are of quite good quality, but also believes that the planning process is very time consuming, that respecting more requests would prove useful and that the current bartering process is unfair. This suggests that the MWW scheduling model should prioritise creating practically useful schedules with a focus on preferences and fairness.

Chapter 4

Scope and problem description

In this chapter we state the scope of our master's thesis in Section 4.1. We describe the planning process that involves creating schedules for the employees at Maternity Ward West in Section 4.2.

4.1 Problem Scope

As stated in Chapter 1, the goal of this master's thesis is to create a decision support tool that solves the scheduling problem at Maternity Ward West by producing schedules of such quality that they are preferable to manually made schedules. To reach this goal we develop the MWW scheduling model. The problem scope should thus be formulated so that the MWW scheduling model tackles the specific scheduling problem the ward manager faces at MWW, as it is described in Chapter 2. Furthermore, the problem scope should be based on the understanding of how different authors model similar problems, acquired by studying related literature in Chapter 3. The previous chapters thus lead to the formulation of this problem scope, providing what to include in the problem description in Section 4.2 and the model formulation in Chapter 5.

4.1.1 Key aspects

Firstly, the MWW scheduling model should be able to take into account the most important considerations of the existing scheduling problem. In Section 2.4 we pointed out how scheduling rules are based on governmental work regulations, ward policies, preferred practises and agreements with trade unions. These scheduling rules are vital when creating the MWW scheduling model. Furthermore, both employee preferences and fairness were highlighted as important aspects in Sections 2.4 and 3.3. Thus, we conclude we should focus on these aspects when designing the MWW scheduling model as well. The MWW scheduling model should account for employees' preferences in a very detailed manner, by implementing all preference measures the ward manager describes if possible. Furthermore, we will do our best to create schedules the ward manager perceives as fair. To do this, we will model all the preferred practises concerning amounts of night shifts, weekends and vacations as they are practised today. Furthermore, we will let the ward manager evaluate the schedule we produce and evaluate the fairness. Lastly, we will remove the need of the unfair bartering in the current planning process, as will be described later in Section 4.1. Costs are, however, a more ambiguous aspect when designing the MWW scheduling model. As all costs are sunk in the offline operational hierarchical level, they are constant and thus unnecessary to model. The modelling of costs therefore rely on the hierarchical limits of our scope.

4.1.2 Hierarchical limits

We must include the offline operational hierarchical level in our problem scope to produce useful schedules at MWW. However, some decisions from the tactical/strategical levels could be useful to include, especially those concerning staffing, as mentioned in Section 3.2. However, as the ward manager does not have the power to decide the staffing levels at the ward, we will not include the staffing levels as a decision for the model. Instead, we perform some analyses with different levels of available staff in Chapter 6 when the model is functional. Furthermore, it would be useful to take

into account some decisions from the online operational level, as mentioned in Section 3.2. However, due to the time spent on the planning process and the long time horizon of the scheduling problem, we will not be able to evaluate the real-life data of how schedules work when they are in effect until this master's thesis is long overdue. Thus, we do not focus on the online operational perspective when developing the MWW scheduling model, with the exception of evening out the over-coverage between days and shifts to create more robust schedules. We can thus conclude that no costs are modelled explicitly in the model.

4.1.3 The MWW scheduling model in the planning process

Because we wish to create a model that makes scheduling with our model preferable to manual scheduling, our model should simplify the planning process significantly and produce high quality schedules. It is also important that the employees experience that they have a significant influence on the final schedule, to avoid them feeling sidelined by some computer program that works in mysterious ways. However, we assume that the employees will experience an acceptable level of influence as long as the model grants the majority of their requests and respects their personal inclinations. We therefore want the mathematical scheduling model to replace the very comprehensive steps 5, 6 and 7 in the 9-step model in Section 2.5. Thus, the 6 weeks long scheduling in the current planning process is removed while providing the scheduling group with an optimised schedule. The schedule should still be evaluated and adjusted by the ward manager and planning group manually, in case the model fails to capture some real life considerations. If the scheduling group and ward manager are not satisfied with the optimised schedule, the model can be adjusted accordingly and run again with little effort. The new optimised schedule is then evaluated by the ward manager and scheduling group, resulting in an iterative process of improving the scheduling model.

We wish to maintain the earliest steps of the planning process in Section 2.5 as these produce and validate input data essential for the MWW scheduling

model. Below we stepwise describe our proposed planning process where the MWW scheduling model is integrated and illustrate this process in Figure 4.1. The presentation of the steps below is a rough description showing the most important information flow between entities. In reality, there is also a lot of informal exchange of information, ensuring the employees' influence on the process. The reader should note that the times of the different steps are estimated, based on the current time used in each step. This is, however, a very conservative estimate for some steps. E.g. it seems realistic that the time used for the trade union's representatives to assess the schedule can be reduced when the feasibility of the schedules are ensured using hard constraints.

Step 1) Registering preallocated shifts:

7 weeks before the schedule is published

The planning process begins with planning vacations and the variety of shifts that are to be preallocated when solving the scheduling problem. These are delivered by multiple sources, including policy makers, the ward manager etc. The preallocated shifts are shifts such as single days when employees attend classes (competence days), co-ordinating shifts, etc.

Step 2) Registration of requests:

7 weeks before the schedule is published

Employees register their preferred individual schedules for the entire planning period in the RMS. Their requested individual schedules may consist of requests for any existing kind of shift (including off-shifts) or simply a blank space for a day, as well as comments elaborating on their choices of requests. Any employee may register a request for every day in the planning period.

Step 3) Retrieval of requests:

4 weeks before the schedule is published

After the employees have registered all their requests and other personal information, the ward manager retrieves the information from the RMS.

Step 4) Sending data:

3.5 weeks before the schedule is published

The ward manager checks the validity of the requests in the RMS, change potential false registrations of requests and schedule the pre-allocated shifts. The requests are then sent to the MWW scheduling model as input data, with data concerning potential changes in employment or scheduling rules.

Step 5) Presenting optimised schedule:

3 weeks before the schedule is published

The MWW scheduling model is run taking into account all employees' requests, personal inclinations and other scheduling rules. The model produces an optimised schedule prioritising heavily the fulfilment of employees' requests. The optimised schedule is then sent to the scheduling group and ward manager. This may be an iterative process if ward manager and scheduling group are dissatisfied with the optimised schedule.

Step 6) Finalising schedule:

2 weeks before the schedule is published

The ward manager and scheduling group finalises the schedule, which is feasible. Some adjustments are made manually if some important employee preferences have been omitted. The final schedule is then sent to the trade union representative for assessment.

Step 7) Publishing the schedule

4 weeks before the schedule is executed

Once the final schedule is verified by the trade union representative it is given to the employees four weeks ahead of taking effect. When the schedule is in effect, shifts are swapped internally by employees and the ward's assistant manager finds employees to cover shifts that suddenly are without cover due to sickness etc. The cover of weekend shifts is especially challenging due to lack of over-coverage during weekends, thus the weekend shifts are evaluated weekly.

In Figure 4.1 we illustrate how information flows between the different entities in our proposed planning process. The arrows are labelled with numbers representing the steps explained above, and the other geometries represent the different entities in the planning process. Note that in Figure 4.1 we group the ward manager and scheduling group together as one entity since they in reality work closely together during the planning process.

It should be noted that removing both the processing in step 5 and the bartering in step 6 with a mathematical model could potentially produce far higher quality schedules. Employees are only able to perform changes in the schedule by swapping and trading some shifts. It is impossible for any human to solve the large scheduling problem optimally without the help of some OR-techniques. Thus, our model should provide a superior way of respecting employees' preferences in a fair way. The challenge is to develop a mathematical model that indeed reflects the real-life planning problem and creates useful schedules. We will use the planning horizon that




Figure 4.1: Rough presentation of the new planning process including the MWW scheduling model, that visualises the information flow that directly results in the final schedules. In reality, much information also flows through other channels.

they currently use at MWW, although it is very long for such a scheduling problem. This is because we want our model to fit into the existing planning process as seamlessly as possible.

4.2 Problem description

In this section we provide the reader with a description of the scheduling problem. The problem is very detailed in the case of MWW, and to provide a full problem description much information regarding every employee is required. We will, however, present a more general problem description, e.g. avoiding specific details only valid for one employee. We therefore begin by introducing the general characteristics of the scheduling problem before presenting scheduling rules and the aspects of the objective function.

4.2.1 Employees

In every planning horizon, employees must be scheduled to work the number of hours stated in their contracted work. However, an employee's scheduled hours of work may have small deviations from its contracted work. Any deviation will be accounted for in the subsequent planning horizon. The employees are divided into skill categories, which determine the shifts they are qualified for. Each employee reports requests regarding their personal schedule in every planning period. Also, employees have personal inclinations affecting the shifts they can and should be allocated to.

4.2.2 Shifts

Each employee is scheduled to have one shift every day in the planning horizon, and all shifts are either an off-shift or a work shift. Work shifts take place either during daytime, evening time or night time. Employees must rest a minimum number of hours between work shifts. The work shifts start and end with some overlap so that MWW can be continuously staffed on all days.

4.2.3 Demand types

The demand for all demand types must be covered on all shifts in the planning horizon. Most shifts have one demand type while some have two. A demand type is covered only when the targeted number of employees within the right skill category are allocated to a shift servicing that demand type. Some skill categories of employees are eligible to cover more than one demand type.

4.2.4 List of scheduling rules

We present all scheduling rules in the following list, briefly explaining the most important implications they have. The schedules at MWW are acyclical. Some scheduling rules were introduced in the previous paragraphs, but we include them here to create a complete list. We note that many scheduling rules are adapted for employees' personal inclinations. This either leads to a scheduling rule being adjusted for an employee, or that an employee is disregarded entirely for certain scheduling rules. If an employee requests to break a scheduling rule, this is generally accepted, unless the rule is vital in securing demand coverage or in securing the employee's health.

- Every day all employees are allocated to either one work shift or one off-shift.
- The demand for all demand types must be covered on all days.
- The demand for a demand type cannot be covered by only newly hired employees.
- Certain employees cannot work together on the same shift.
- Over-coverage of shifts is allowed, but under-coverage is not. There is a limit to how much over-coverage that is acceptable for different demand types and days.

Scope and problem description

- There is a minimum number of hours each employee must work during the planning horizon and a maximum number of hours each employee can work during the planning horizon.
- Over any span of seven days, there is a maximum number of hours employees can work.
- Between two work shifts an employee must rest no less than a minimum number of hours. This makes it illegal to allocate employees to certain successions of shifts.
- Employees cannot work more than a maximum number of consecutive work shifts.
- Some shifts have a limit to the maximum number of consecutive days an employee can work that shift.
- Every employee must have one protected off-day each week. The protected off-day must be placed on Sunday if the employee has the weekend off duty, otherwise it can be placed on any day within the given week.
- During the protected off-day, employees must rest no less than a minimum number of hours. As such, some combinations of shifts cannot be allocated to the day before and after the protected off-day.
- An employee either has the weekend off duty or works both Saturday and Sunday.
- Each employee must work at least a minimum number of weekends each planning horizon, but not more than a maximum number of weekends.
- An employee must have the successive weekend off-duty after working a weekend, unless the employee requested working two weekends in a row.

- Under certain circumstances some shift-patterns are illegal to allocate to an employee.
- Under certain circumstances some shift-patterns must be allocated to an employee.
- Each employee must work at least a minimum number of certain shifts in each planning horizon, but not more than a maximum number of that shift.
- Some employees should always have their requests granted.
- All preallocated shifts must be allocated, including vacations, particular kinds of shifts and agreements made between the ward manager and employees.
- Some employees cannot work certain shifts due to health issues.
- Shifts considered unpopular should be distributed fairly.
- Schedules should be co-ordinated with the employees' other workplaces if they work two jobs.

4.2.5 Objectives

The objectives in the scheduling problem are to maximise the fulfilment of employees' requests and to maximise the allocation of shift-patterns supporting employees' personal inclinations. The reason for these goals is that employees must feel a sense of involvement in the planning process to accept the schedule. The following list contains the goals of the scheduling:

- Maximise the number of granted requests
- Maximise the number of allocated shift-patterns desired through personal inclinations

Chapter 5

Model

In this chapter, the MWW scheduling problem is formulated as a general mathematical ILP model. Section 5.1 explains principles for how we construct the scheduling model and how we present the scheduling model mathematically. Section 5.2 defines indices, sets and parameters which encompass all input data required to solve the scheduling problem. Variables representing the decisions made by the scheduling model are also presented in Section 5.2. Subsequently, we present the objective function containing the goals the model seeks to fulfil in Section 5.3. Finally scheduling rules are formulated mathematically as constraints in Section 5.4.

5.1 Modelling principles

In Section 5.1 we elaborate on choices for how the MWW scheduling model is designed and how we wish to present the scheduling model as a general mathematical formulation.

5.1.1 Model design

We model the scheduling problem at MWW as an ILP model. ILP or MIP approaches have been used in similar models like in Rönnerberg and Larsson

(2010), Maenhout and Vanhoucke (2013), Christiansen et al. (2015), Beckmann and Klyve (2016) etc. ILP and MIP models are flexible and easy to adapt, which is advantageous since the MWW scheduling model will have to be updated several times based on the feedback from the ward manager. Also, the flexibility allows the model to be easily updated in line with future changes to the scheduling problem at MWW. In our ILP formulation, the most noteworthy decision variables are binary and defined for each nurse, each shift on each day. In an alternate formulation-technique, like column generation, the decision variables are defined for employees and sequences of shifts spanning the planning horizon. Although such formulations typically outperform ILP models with respect to run time, they are also less flexible, and in our opinion less suitable for tackling the real-life scheduling problem at MWW.

5.1.2 Mathematical formulation

Capturing all the relevant data for the scheduling problem at MWW requires a highly detailed model, including a large number of parameters and sets as well as constraints describing very specific situations, sometimes occurring only once or twice in the planning horizon. To present the model, we have in several cases made slight simplifications to generalise the formulation. Thus, some specific and similar constraints are generalised to be presented as one. We select some of the instance-specific cases as examples and formulate them as constraints containing real data, though excluding employees' names. These examples are meant to convey that many of the general aspects we present also have instance-specific variations which must be accounted for in real life scheduling.

5.2 Definitions

We name sets using upper case calligraphic letters, parameters using upper case letters, and variables using lower case letters. We use subscripts to denote valid indices and superscripts containing capital letters or numbers

to specify what some parameters and sets represent. Parameters also have overlines or underlines when the parameters represent upper or lower limits respectively. To generalise certain constraints we make use of generic sets denoted by 'GEN' in their superscript. Whenever generic sets are used they are elaborated on in the explanation of the constraints.

5.2.1 Indices

n	-	employee
s	-	shift
t	-	day
u	-	demand type
k	-	week
p	-	shift-pattern

5.2.2 Sets

\mathcal{N}	-	Set of employees.
\mathcal{N}_u^U	-	Set of employees in the skill category reciprocating demand type u , $\bigcup_{u \in \mathcal{U}} \mathcal{N}_u^U = \mathcal{N}$.
\mathcal{N}^{DES}	-	Set of employees who desire shift-patterns rewarded in the objective function.
\mathcal{N}^{GEN}	-	A generic set of employees explained further whenever used.
\mathcal{S}	-	Set of shifts.
\mathcal{S}^W	-	Set of shifts which are work-shifts, $\mathcal{S}^W \subset \mathcal{S}$.
\mathcal{S}^O	-	Set of shifts representing all off-days, $\mathcal{S}^O \subset \mathcal{S}$. $\mathcal{S}^O \cup \mathcal{S}^W = \mathcal{S}$.
\mathcal{S}_u^U	-	Sets of shifts covering demand type u , $\bigcup_{u \in \mathcal{U}} \mathcal{S}_u^U \cup \mathcal{S}^O = \mathcal{S}$.
\mathcal{S}^{GEN}	-	A generic set of shifts explained further whenever used.
\mathcal{K}	-	Set of weeks in the planning horizon.
\mathcal{T}	-	Set of days in the current planning horizon.
\mathcal{T}^{SUN}	-	Set of Sundays in the planning horizon.
\mathcal{T}^{MON}	-	Set of Mondays in the planning horizon.

- \mathcal{T}^C - Set of days during Christmas.
- \mathcal{T}^P - Set of days in the current and previous planning horizon. \mathcal{T}^P includes positive and negative values of t where values of $t < 1$ are days in the previous planning horizon and $t = 0$ is the last day in the previous planning horizon. $\mathcal{T} \subset \mathcal{T}^P$.
- \mathcal{T}_k - Set of the days in week k , $\bigcup_{k \in \mathcal{K}} \mathcal{T}_k = \mathcal{T}$.
- \mathcal{T}^{GEN} - Generic set of days explained further whenever used.
- \mathcal{U} - Set of demand types.
- \mathcal{U}^{GEN} - Generic set of demand types explained further whenever used.
- \mathcal{P}_n^{ILL} - Set of shift-patterns which are illegal to allocate to employee n .
- \mathcal{P}_n^{MAN} - Set of shift-patterns which are mandatory to allocate to employee n , under certain circumstances.
- \mathcal{P}_n^{AFT} - Set of shift-patterns that force other shift-patterns to succeed it and the succeeding shift-patterns, for employee n . This is explained thoroughly in example B in Section 5.4.
- \mathcal{P}_n^{DES} - Set of desirable shift-patterns which should be rewarded in the objective function when allocated to employee n .

5.2.3 Parameters

Weighting parameters

The weighting parameters are coefficients that assign relative weights to the different terms in the objective function. All weighting parameters are non-negative.

- W^R - Weight parameter rewarding requests.
- W^D - Weight parameter rewarding desirable shift-patterns.

Limit parameters

Limit parameters set limits for constraints. When a parameter limits constraints on the maximum or minimum number of shifts or days over some

time horizon we use the letter 'M'. 'H' is used for limits on hours. 'D' is used for limits regarding demand.

- \underline{D}_{tu} - Minimum number of employees needed to cover the demand for demand type u on day t .
- \overline{D}_{tu}^{OC} - Maximum number of employees allowed to over-cover demand type u on day t .
- \overline{M}_n^{CW} - Maximum number of consecutive work shifts for employee n .
- \overline{M}_{ns}^{CS} - Maximum number of consecutive shifts s for employee n .
- \overline{M}_{ns}^{TS} - Maximum number of shifts s employee n can work in the planning horizon.
- \underline{M}_{ns}^{TS} - Minimum number of shifts s employee n must work during the planning horizon.
- \overline{M}_n^{WE} - Maximum number of weekends employee n can work during the planning horizon.
- \underline{M}_n^{WE} - Minimum number of weekends employee n must work during the planning horizon.
- \overline{M}_n^{NW} - Employee n works one in every \overline{M}_n^{NW} weekends.
- \overline{M}_s^{FT} - Maximum deviation in number of unpopular shifts s allocated to each employee within a certain skill category.
- \overline{H}_n^{7D} - Maximum number of hours employee n can work during any 7-day interval.
- \overline{H}_n^{CW} - Maximum number of hours employee n can work during the planning horizon.
- \underline{H}_n^{CW} - Minimum number of hours employee n must work during the planning horizon.

Indicating parameters

Indicating parameters are binary, and all indicating parameters are denoted by P .

- P_{nst}^{PA} - 1 if employee n should have shift s preallocated on day t , 0 else.
 P_{nst}^{R} - 1 if employee n requests shift s on day t , 0 else.
 P_n^{DE} - 1 if employee n should always be allocated its requests, 0 else.
 $P_{s_1s_2}^A$ - 1 if there is sufficient time between shifts s_1 and s_2 on days $t - 1$ and t respectively for an employee to work them both, 0 else.
 $P_{s_1s_2}^{F1}$ - 1 if there is sufficient time between shifts s_1 and s_2 on days $t - 2$ and t respectively for an employee to be allocated to an 'F1' off-day on day $t - 1$, 0 else.
 $P_{n_1n_2}^{TOG}$ - 1 if employees n_1 and n_2 can work together during the same shift.
 $P_{u_1u_2}^U$ - 1 if $\mathcal{S}_{u_1}^U = \mathcal{S}_{u_2}^U$, i.e., 1 if the two demand types, u_1 and u_2 , are serviced by the same set of shifts, 0 else.
 P_{ns}^S - 1 if employee n can work shift s , 0 else.
 $P_{np_1p_2}^{AFT}$ - 1 if shift-pattern p_2 must be allocated to employee n immediately after shift-pattern p_1 is allocated, 0 else.

Dynamic parameters

Dynamic parameters are parameters with varying column length. There are several types of shift-patterns dealt with at MWW. Certain shift-patterns are illegal to allocate to an employee while others are regarded as preferable if allocated to certain employees. To represent the shift-patterns in a general way, we present the dynamic parameters as arrays, like the ones shown in Tables 5.6 and 5.7. Each array represents a shift-pattern p , within a group of shift-patterns, e.g. illegal shift-patterns are in the set \mathcal{P}_n^{ILL} . The different columns represent different shifts.

We create an example where s_1 , s_2 , s_3 and s_4 represent the day shift, evening shift, night shift and off-shift, respectively. The rows represent the different days in shift-pattern p , and the length of the columns give the number of days the shift-pattern spans, $\bar{\tau}_p$. As p_1 and p_2 can span different numbers of days, a parameter containing both shift-patterns becomes dy-

namic, as the dimension representing days can vary with p . The parameter S_{stp}^{ILL} is 1 when shift s on day t of shift-pattern p is included in an illegal shift-pattern. Looking at Table 5.6 we see this illegal shift-pattern: night shift on day 1, off-shift on day 2 and night shift on day 3. As such, the shift-pattern 'day-evening-night' is illegal to allocate to employee n during any succession of three days within the planning horizon, if that shift-pattern is included in the set of illegal shift-patterns \mathcal{P}_n^{ILL} . In Table 5.7 the illegal shift-pattern 'day-off-night-evening' is presented. In this case, the shift-pattern spans four days, rather than three as in Table 5.6. The other forms of shift-patterns are defined exactly like the illegal shift-patterns, but their related scheduling rules are formulated in separate constraints. Constraints regarding some of the different types of shift-patterns are given as examples in Section 5.4 alongside tables defining the shift-patterns.

Table 5.6: Illegal shift-pattern p_1 , indicating it is illegal to work 'N-OFF-N'

	Shift, s			
	s_1	s_2	s_3	s_4
Day in shift-	1	0	0	1
pattern p_1	2	0	0	0
	3	0	0	1

Table 5.7: Illegal shift-pattern p_2 , indicating it is illegal to work 'D-OFF-N-E'

	Shift, s			
	s_1	s_2	s_3	s_4
Day in shift-	1	1	0	0
pattern p_2	2	0	0	0
	3	0	0	1
	4	0	1	0

- S_{stp}^{ILL} - Illegal shift-patterns. $S_{stp}^{ILL} = 1$ if shift-pattern p includes shift s on day t of the shift-pattern, and 0 else.
 S_{stp}^{AFT} - Shift-patterns that must be allocated after certain shift-pattern is allocated. $S_{stp}^{AFT} = 1$ if shift-pattern p includes shift s on day t of the shift-pattern, and 0 else.
 S_{stp}^{OBJ} - Shift-patterns that are rewarded in the objective function. $S_{stp}^{OBJ} = 1$ if shift-pattern p includes shift s on day t of the shift-pattern, and 0 else.
 S_{ntp}^{MAN} - Shift-pattern that must be allocated to an employee if the employee is allocated any of the shifts included in the shift-pattern. $S_{stp}^{MAN} = 1$ if shift-pattern p includes shift s on day t of the shift-pattern, and 0 else.

General parameters

The remaining parameters are called general parameters.

- H_{nst} - Calculated duration of shift s , in hours, for employees n on day t .
 V_n^H - Due to vacations, the number of hours employee n should work is reduced by V_n^H .
 V_{ns}^S - Due to vacations, the number of shifts s employee n should work is reduced by V_{ns}^S .
 $\bar{\tau}_p$ - Number of days spanned by shift-pattern p .

5.2.4 Variables

For all negative values t and $t = 0$, x_{nst} represent the schedule in the previous planning horizon. This allows us to constrain the problem using information from the last planning horizon, as seen in e.g. constraints (5.9).

- x_{nst} - 1 if employee n works shift s on day t , 0 else.
 q_{ntp} - 1 if employee n works a desirable shift-pattern p ending on day t , 0 else.

y_{tu} - Integer variable taking value when a demand type u is covered on day t by an employee not in \mathcal{N}_u^U . This is relevant for midwives and assistant nurses.

5.3 Objective function

Here we present the objective function in MWW scheduling model containing the goals we seek to achieve for the schedules. We have made efforts into reducing the overall number of soft constraints in the scheduling model, resulting in a lean and understandable objective function.

$$\max Z = W^R \sum_{n \in \mathcal{N}} \sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{T}} P_{nst}^R x_{nst} + W^D \sum_{n \in \mathcal{N}^{DES}} \sum_{t \in \mathcal{T}} \sum_{p \in \mathcal{P}_n^{DES}} q_{ntp} \quad (5.1)$$

In the objective function, the first term is the total amount of respected requests and the second is the number of desirable shift-patterns allocated to employees in the set \mathcal{N}^{DES} .

5.4 Constraints

Here we present all constraints in the model. These constitute the scheduling rules the schedule must abide by. The constraints are categorised together based on their functionality. E.g. constraints related to weekend-work are presented in a subsection together.

5.4.1 Covering demand

$$\sum_{s \in \mathcal{S}} x_{nst} = 1, \quad n \in \mathcal{N}, t \in \mathcal{T} \quad (5.2)$$

$$y_{tu_1} + y_{tu_2} = 0, \quad t \in \mathcal{T}, u_1, u_2 \in \mathcal{U} \mid P_{u_1 u_2}^U = 1 \quad (5.3)$$

$$\underline{D}_{tu} \leq \sum_{n \in \mathcal{N}_u^U} \sum_{s \in \mathcal{S}_u^U} x_{nst} + y_{tu} \leq \underline{D}_{tu} + \overline{D}_{tu}^{OC}, \quad t \in \mathcal{T}, u \in \mathcal{U} \quad (5.4)$$

$$\sum_{n \in \mathcal{N}_u^U \cap \mathcal{N}^{GEN}} \sum_{s \in \mathcal{S}_u^U} x_{nst} \geq 1, \quad t \in \mathcal{T}, u \in \mathcal{U}^{GEN} \quad (5.5)$$

$$x_{n_1st} + x_{n_2st} \leq 1, \quad n_1, n_2 \in \mathcal{N}, s \in \mathcal{S}, t \in \mathcal{T} \mid P_{n_1n_2}^{TOG} = 0 \quad (5.6)$$

Constraints (9.2) make sure one shift is allocated to all employees on all days. Constraints (9.3) ensure midwives and assistant nurses can work the same shifts while covering different demand types. When midwives cover the demand type for assistant nurses on day t , they will not contribute to covering the demand for midwives' demand types, since an employee can only service one demand type each day. This is ensured by constraints (9.3). Constraints (9.3) and (9.4) together ensure that all types of demand u are covered legally every day t . With this formulation the model does not specify which midwife that covers the demand type for an assistant nurse, and we thus avoid symmetric solutions. Furthermore, it should be noted that assistant nurses cannot cover the demand for midwives, as is assured by constraints (5.32) and (5.33). Constraints (9.4) ensure that types of demand u on days t are over-covered by a maximum of \overline{D}_{tu}^{OC} employees. Constraints (9.5) ensure that certain demand types are covered by at least one experienced employee. In constraints (9.5) the set \mathcal{N}^{GEN} include all employees regarded as experienced and the set \mathcal{U}^{GEN} contains all demand types that must be covered by at least one experienced employee. Constraints (9.6) ensure that certain employees n_1 and n_2 never work together.

5.4.2 Work hours

$$\underline{H}_n^{CW} \leq \sum_{s \in \mathcal{S}^W} \sum_{t \in \mathcal{T}} H_{nst} x_{nst} + V_n^H \leq \overline{H}_n^{CW}, \quad n \in \mathcal{N} \quad (5.7)$$

$$\sum_{s \in \mathcal{S}^W} \sum_{\tau=t-6}^t H_{nst} x_{ns\tau} \leq \overline{H}_n^{7D}, \quad n \in \mathcal{N}, t \in \mathcal{T} \quad (5.8)$$

Constraints (5.7) ensure that employee n works a correct number of hours during the planning horizon. V_n^H is included when counting work hours,

to account for the vacation employee n had during the planning horizon. At MWW, some employees have two different contracted works applying to each its part of the planning horizon. This happens e.g. when an employee returns from maternity leave in the midst of the planning horizon. When implementing the model, we handled the employees with changing contracted work by dividing the planning horizon into two parts, one for each contracted work. This allows us to appropriately constrain employees' work hours according to their two contracts. Here however we present only the general case, where the contracted work remains unchanged over the planning horizon. Constraints (5.8) ensure that no employee works more than \overline{H}_n^{7D} hours in any given 7-day period.

5.4.3 Required rest

$$x_{ns_1(t-1)} + \sum_{s_2 \in \mathcal{S} \mid P_{s_1 s_2}^A = 0} x_{ns_2 t} \leq 1, \quad n \in \mathcal{N}, s_1 \in \mathcal{S}, t \in \mathcal{T} \quad (5.9)$$

$$x_{ns_1(t-2)} + x_{n'F1'(t-1)} + \sum_{s_2 \in \mathcal{S} \mid P_{s_1 s_2}^{F1} = 0} x_{ns_2 t} \leq 2, \quad n \in \mathcal{N}, s_1 \in \mathcal{S}, t \in \mathcal{T} \quad (5.10)$$

$$\sum_{t \in \mathcal{T}_k} x_{n'F1't} = 1, \quad n \in \mathcal{N}, k \in \mathcal{K} \quad (5.11)$$

$$\sum_{s \in \mathcal{S}^W} \sum_{\tau=t-\overline{M}_n^{CW}}^t x_{ns\tau} \leq \overline{M}_n^{CW}, \quad n \in \mathcal{N}, t \in \mathcal{T} \quad (5.12)$$

$$\sum_{\tau=t-\overline{M}_{ns}^{CS}}^t x_{ns\tau} \leq \overline{M}_{ns}^{CS}, \quad n \in \mathcal{N}, s \in \mathcal{S}^W, t \in \mathcal{T} \quad (5.13)$$

Constraints (5.9) ensure that only some shifts may follow each other. This is due to regulations of how long a resting period between shifts should be. Constraints (5.10) make sure a similar regulation as in constraints (5.9) is

followed if a protected off-day (F1) is allocated between the two shifts s_1 and s_2 . Constraints (5.11) ensure that all employees have exactly one protected off-day in every week. Constraints (5.12) make sure that employee n works a maximum of \overline{M}_n^{CW} consecutive work shifts before being granted an off-shift. Constraints (5.13) ensure that employee n never works more than \overline{M}_{ns}^{CS} consecutive shifts s .

5.4.4 Weekends

$$\sum_{s \in \mathcal{S}^W} (x_{ns(t-1)} - x_{nst}) = 0, \quad n \in \mathcal{N}, t \in \mathcal{T}^{SUN} \setminus \mathcal{T}^C \quad (5.14)$$

$$\underline{M}_n^{WE} \leq \sum_{s \in \mathcal{S}^W} \sum_{t \in \mathcal{T}^{SUN} \setminus \mathcal{T}^C} x_{nst} \leq \overline{M}_n^{WE}, \quad n \in \mathcal{N} \quad (5.15)$$

$$\sum_{s \in \mathcal{S}^W} \sum_{\tau=0}^{\overline{M}_n^{NW}-1} x_{ns(t-7\tau)} \leq 1, \quad n \in \mathcal{N}, t \in \mathcal{T}^{GEN} \quad (5.16)$$

Constraints (5.14) enforce that all employees work both Saturday and Sunday if allocated to a work shift during the weekend. There are exceptions to this scheduling rule during Christmas, when it is allowed to work only one of the days. Constraints (5.15) ensure that employee n works an appropriate number of weekends, excluding Christmas, during the planning horizon. Constraints (5.16) ensure that employees cannot work during weekends more than once every \overline{M}_n^{NW} week. In constraints (5.16) the set \mathcal{T}^{GEN} only includes all Sundays where employee n has not registered requests contradicting this scheduling rule.

5.4.5 Shift-patterns

For these constraints, it is fundamental that only one shift is allocated to an employee on each day and that $\overline{\tau}_p$ is the number of days the shift-pattern p spans. The MWW scheduling model initialises several constraints for different shift-patterns affecting different instances of employees, shifts and

days. After some simplifications all these shift-patterns may be generalised into four different kinds of shift-patterns, namely shift-patterns rewarded in the objective function, shift-patterns illegal to allocate, shift-patterns mandatory to allocate and shift-patterns which must be allocated after an employee works a certain shift-pattern.

$$\sum_{s \in S} \sum_{\tau=t-\bar{\tau}_p+1}^t S_{s(\tau-t+\bar{\tau}_p)p}^{DES} x_{ns\tau} - \bar{\tau}_p q_{ntp} \geq 0, \quad n \in \mathcal{N}^{DES}, t \in \mathcal{T}, p \in \mathcal{P}_n^{DES} \quad (5.17)$$

$$\sum_{s \in S} \sum_{\tau=t-\bar{\tau}_p+1}^t S_{s(\tau-t+\bar{\tau}_p)p}^{ILL} x_{ns\tau} \leq \bar{\tau}_p - 1, \quad n \in \mathcal{N}, t \in \mathcal{T}^{GEN}, p \in \mathcal{P}_n^{ILL} \quad (5.18)$$

$$\sum_{s \in S} \sum_{\tau=t-\bar{\tau}_p+1}^t S_{s(\tau-t+\bar{\tau}_p)p}^{MAN} x_{ns\tau} - \bar{\tau}_p \sum_{s \in S} S_{s(1)p}^{MAN} x_{ns(t-\bar{\tau}_p+1)} = 0, \quad (5.19)$$

$$n \in \mathcal{N}, t \in \mathcal{T}^{GEN}, p \in \mathcal{P}_n^{MAN}$$

$$\sum_{s \in S} \sum_{\tau_1=t-\bar{\tau}_{p_1}-\bar{\tau}_{p_2}+1}^{t-\bar{\tau}_{p_2}} S_{s(\tau_1-t+\bar{\tau}_{p_1}+\bar{\tau}_{p_2})p_1}^{AFT} x_{ns\tau_1} - \frac{1}{\bar{\tau}_{p_2}} \sum_{\tau_2=t-\bar{\tau}_{p_2}+1}^t S_{s(\tau_2-t+\bar{\tau}_{p_2})p_2}^{AFT} x_{ns\tau_2} \leq \bar{\tau}_{p_1} - 1, \quad (5.20)$$

$$n \in \mathcal{N}, t \in \mathcal{T}^{GEN}, p_1, p_2 \in \mathcal{P}_n^{AFT} \mid P_{np_1 p_2}^{AFT} = 1$$

Constraints (5.17) assign value to q_{ntp} whenever a desirable shift-pattern $p \in \mathcal{P}_n^{DES}$ ending on day t is allocated to an employee in the set \mathcal{N}^{DES} , and q_{ntp} is then rewarded in the objective function. Certain shift-patterns are illegal to allocate to employee n on all days while some shift-patterns

are illegal only during certain days, e.g. in the weekend. Constraints (5.18) enforce that the illegal shift-patterns $p \in \mathcal{P}_n^{ILL}$ are never allocated to employee n during the days t ending in \mathcal{T}^{GEN} . Constraints (5.19) ensure that if employee n is given a shift in a mandatory shift-pattern $p \in \mathcal{P}_n^{MAN}$ ending on day t in \mathcal{T}^{GEN} , the entire mandatory shift-pattern must be allocated to employee n . Constraint (5.20) ensure that shift-pattern p_2 is allocated to employee n immediately after shift-patterns p_1 is allocated, provided that shift-pattern p_2 must immediately succeed shift-pattern p_1 when p_1 is allocated to employee n ($P_{np_1p_2}^{AFT} = 1$). An example of this is that employee n must be allocated two off-days immediately after working three consecutive night shifts. In constraints (5.18)-(5.20) the set \mathcal{T}^{GEN} contains only the days relevant for the different shift-patterns p , but does not include the days where an employee made requests which should be prioritised above respecting the scheduling rules regarding shift-patterns.

Example A: Illegal shift-patterns

In this example we illustrate one illegal shift-pattern by looking at a specific instance of shifts s . The shift-pattern 'D-E-OFF-N-E' is considered very unpopular and therefore made illegal to allocate to all employees on all days, except those days where an employee specifically requests to work that shift-pattern. We denote the illegal shift-pattern p_A and express it by the dynamic parameter $S_{stp_A}^{ILL}$ shown in Table 5.11. The span of shift-pattern p_A is 5 days, i.e $\bar{\tau}_{p_A} = 5$. In Table 5.11, s_1, s_2, s_3 and s_4 respectively represent the day shift 'D', evening shift 'E', night shift 'N' and off-shift 'OFF'. Using the aforementioned instances and parameters we formulate constraints (5.21), which are consistent with the general illegal shift-patterns in constraints (5.18).

$$\begin{aligned}
 & x_{n'D'(t-4)} + x_{n'E'(t-3)} + x_{n'OFF'(t-2)} + x_{n'N'(t-1)} + x_{n'E't} \leq 4, \quad (5.21) \\
 & n \in \mathcal{N}, t \in \mathcal{T} \\
 & | (P_n^{DE} = 0) \cap (P_{n'D'(t-4)}^R P_{n'E'(t-3)}^R P_{n'OFF'(t-2)}^R P_{n'N'(t-1)}^R P_{n'E't}^R = 0)
 \end{aligned}$$

Constraints (5.21) enforce that the illegal shift-pattern 'D-E-OFF-N-E' is never allocated to employees on any days, except those days ending on t where an employee n specifically requests to work that shift-pattern. Expressed mathematically, the shift-pattern is illegal for all employee and days given that:

$$n \in \mathcal{N}, t \in \mathcal{T} \mid (P_{n'D'(t-4)}^R P_{n'E'(t-3)}^R P_{n'OFF'(t-2)}^R P_{n'N'(t-1)}^R P_{n'E't}^R = 0).$$

All employees deciding their own schedule $P_n^{DE} = 1$ are also exempt from the constraints. In our opinion, formulating the conditions for constraints like in Example A does not increase the understanding of the mathematical structures of the MWW scheduling problem, and we prefer explaining verbally the elements general sets include for readability.

Table 5.11: Illegal shift-pattern p_A , indicating it is illegal to work 'D-E-OFF-N-E'

	Shift, s			
	s_1	s_2	s_3	s_4
Day in shift-	1	1	0	0
pattern p_A	2	0	1	0
	3	0	0	1
	4	0	0	1
	5	0	1	0

Example B: Shift-pattern that must be allocated immediately after another shift-pattern

We wish to illustrate scheduling rules regarding other forms of shift-patterns dealt with at MWW. In this case shift-patterns which must be allocated immediately after a certain shift-pattern is allocated an employee. To handle

such scheduling rules we use two different shift-patterns and the parameter $P_{np_1p_2}^{AFT}$. An example of this is that when employee n works three consecutive night shifts that employee must be allocated two off-days on the subsequent days, unless the employee requests one or two work shifts during the two subsequent days. In this example, shift-pattern p_1 is 'N-N-N' (three consecutive night shifts), shift-pattern p_2 is 'OFF-OFF' (two consecutive off-shifts), $P_{np_1p_2}^{AFT} = 1$ (employee n must be allocated to shift-pattern p_2 immediately after shift-pattern p_1 is allocated), $\tau_{p_1} = 3$ (shift-pattern p_1 spans three days) and $\tau_{p_2} = 2$ (shift-pattern p_2 spans two days). Shift-pattern p_1 and p_2 are expressed by the dynamic parameters $S_{stp_1}^{AFT}$ and $S_{stp_2}^{AFT}$, given in Tables 5.12 and 5.13. The shifts s_1, s_2, s_3 and s_4 denote the day shift 'D', evening shift 'E', night shift 'N' and off-shift 'OFF', respectively. Using the aforementioned instances and parameters we formulate constraints (5.22) and (5.23), which are consistent with the general constraints (5.20).

$$\sum_{s \in S} \sum_{\tau_1=t-3-2+1}^{t-2} S_{s(\tau_1-t+3+2)p_1}^{AFT} x_{ns\tau_1} - \frac{1}{2} \sum_{\tau_2=t-2+1}^t S_{s(\tau_2-t+2)p_2}^{AFT} x_{ns\tau_2} \leq 3 - 1, \quad (5.22)$$

$$n \in \mathcal{N}, t \in \mathcal{T}, p_1, p_2 \in \mathcal{P}_n^{AFT}$$

$$|(P_{np_1p_2}^{AFT} = 1) \cap (P_{n'N'(t-4)}^R P_{n'N'(t-3)}^R P_{n'N'(t-2)}^R = 0) \cap (\sum_{s \in S^W} \sum_{\tau_3=t-1}^t P_{ns\tau_3}^R = 0)|$$

Completing the summation over days and looking at a specific instance of t provides more insight to the functionality of constraints (5.22). We show this in constraints (5.23) for the day $t = 10$ and employee n_1 . In constraints (5.23) employee n_1 has not made requests which should be prioritised above fulfilling the scheduling rule discussed in Example B. Constraints (5.23) shows more clearly now, that shift-pattern p_2 must be allocated only when all shifts in shift-pattern p_1 are allocated to employee n_1 . Note that the different values of S_{ntp}^{AFT} can be found in Tables 5.12 and 5.13.

$$\sum_{s \in \mathcal{S}} (S_{s1p_1}^{AFT} x_{n_1s6} + S_{s2p_1}^{AFT} x_{n_1s7} + S_{s3p_1}^{AFT} x_{n_1s8}) - \frac{1}{2} \sum_{s \in \mathcal{S}} (S_{s1p_2}^{AFT} x_{n_1s9} + S_{s2p_2}^{AFT} x_{n_1s10}) \leq 2, \quad (5.23)$$

$$p_1, p_2 \in \mathcal{P}_{n_1}^{AFT} \mid P_{n_1p_1p_2}^{AFT} = 1$$

Table 5.12: Shift-pattern p_1 , 'N-N-N', which must be followed by p_2 if p_1 is allocated

	Shift, s			
	s_1	s_2	s_3	s_4
Day in shift-	1	0	0	1
pattern p_1	2	0	0	1
	3	0	0	1

Table 5.13: Shift-pattern p_2 , 'OFF-OFF', which must succeed p_1 when p_1 is allocated

	Shift, s			
	s_1	s_2	s_3	s_4
Day in shift-	1	0	0	0
pattern p_2	2	0	0	1

Example C: Mandatory shift-patterns and fairness

Some important aspects at MWW require specific instances of employees n , shifts s and days t to be presented. In this example, we wish to shed light on how adapting scheduling rules to employees' preferences give rise to additional scheduling rules, so that fairness is achieved. Firstly, during weekends there are shift-patterns considered very popular and these have been made mandatory, so that if an employee is allocated to one shift within the mandatory shift-pattern, the remaining shifts must be allocated as well. In

this example, we discuss the mandatory 'N-N-N' shift-pattern, mandatory only for Saturday-Monday. Employees with personal inclinations suggesting they can't work more than two consecutive night shifts must be exempt from constraints ensuring the 'N-N-N' shift-pattern. These employees will instead be allocated to the 'N-N' shift-pattern spanning Saturday-Sunday, if allocated to a night shift during weekends. However, working night shifts on Mondays is unpopular and employees should work approximately equally many night shifts during Mondays as Sundays (or Saturdays) to ensure the schedule is perceived as fair. Sometimes employees with preferences against working 'N-N-N' during weekends request working this pattern, though they generally don't want to work that shift-pattern. Staffing the MWW in line with employees' preferences during holidays is difficult and many scheduling rules are adjusted or disregarded for the Christmas period.

$$2x_{n'N't} = x_{n'N'(t-1)} + x_{n'N'(t-2)}, \quad (5.24)$$

$$n \in \mathcal{N}, t \in \mathcal{T}^{MON} \setminus \mathcal{T}^C \mid (\overline{M}_{n'N'}^{CS} \geq 3) \cap (P_n^{DE} = 0) \cap \left(\sum_{\tau=t-2}^t P_{n'N'\tau}^R \neq 2 \right)$$

$$x_{n'N't} = x_{n'N'(t-1)}, \quad n \in \mathcal{N}, t \in \mathcal{T}^{SUN} \setminus \mathcal{T}^C \mid (\overline{M}_{n'N'}^{CS} = 2) \cap (P_n^{DE} = 0) \\ \cap ((P_{n'N'(t-1)}^R(1 - P_{n'N'(t)}^R) + (1 - P_{n'N'(t-1)}^R)P_{n'N't}^R) = 0) \quad (5.25)$$

$$\sum_{t \in \mathcal{T}^{MON} \setminus \mathcal{T}^C} x_{n'N't} = \sum_{t \in \mathcal{T}^{SUN} \setminus \mathcal{T}^C} x_{n'N't}, \quad n \in \mathcal{N} \mid (\overline{M}_{n'N'}^{CS} = 2) \cap (P_n^{DE} = 0) \quad (5.26)$$

Constraints (5.24) enforce that employees work night shifts Saturday, Sunday and Monday if allocated to a night shift one of these days, unless the employee requests to only work two night shifts from Saturday to Monday. Similarly, constraints (5.25) ensure employees are allocated to night shifts Saturday and Sunday if working a night shift during the weekend, unless the

employee requests to work a night shift only one of the days in the weekend, but not the other. Constraints (5.26) ensure employees work equally many night shifts during Mondays and Sundays. In Constraints (5.24)-(5.26) only employees who don't self determine their own schedule ($P_n^{DE} = 0$) are included in the set of employees \mathcal{N} . In constraints (5.24) employees who can work three or more consecutive night shifts ($\overline{M}_{n'N'}^{CS} \geq 3$) are included in the set of employees \mathcal{N} , while in Constraints (5.25) and (5.26) employees inclined to work a maximum of two consecutive night shift ($\overline{M}_{n'N'}^{CS} = 2$) are included in the set \mathcal{N} .

5.4.6 Related to fairness

$$\sum_{t \in \mathcal{T}} x_{nst} + V_{ns}^S \geq \underline{M}_{ns}^{TS}, \quad n \in \mathcal{N}, s \in \mathcal{S} \quad (5.27)$$

$$\sum_{t \in \mathcal{T}} x_{nst} + V_{ns}^S \leq \overline{M}_{ns}^{TS}, \quad n \in \mathcal{N}, s \in \mathcal{S} \quad (5.28)$$

$$\sum_{s \in \mathcal{S}^{GEN}} \sum_{t \in \mathcal{T}} (x_{n_1st} - x_{n_2st}) \leq \overline{M}_s^{FT}, \quad n_1, n_2 \in \mathcal{N}^{GEN} \quad (5.29)$$

Constraints (5.27) ensure that employee n works no less than \underline{M}_{ns}^{TS} number of shifts s during the planning horizon. Constraints (5.28) ensure that employee n works no more than \overline{M}_{ns}^{TS} number of shifts s during the planning horizon. In constraints (5.29) \mathcal{S}^{GEN} include only the shifts which must be allocated approximately equally many times to employees in \mathcal{N}^{GEN} . The set \mathcal{N}^{GEN} include employees that can cover shifts regarded as unpopular, and constraints (5.29) ensure that these employees are allocated approximately equally many unpopular shifts. As such, constraints (5.29) even out the number of unpopular shifts allocated to employees.

5.4.7 Variable declarations and fixations

$$x_{nst} \in \{0, 1\}, \quad n \in \mathcal{N}, s \in \mathcal{S}, t \in \mathcal{T}^P \quad (5.30)$$

$$q_{ntp} \in \{0, 1\}, \quad n \in \mathcal{N}, t \in \mathcal{T}, p \in \mathcal{P}^{DES} \quad (5.31)$$

$$y_{tu} \in \{-\underline{D}_{tu}, \dots, 0\}, \text{integer} \quad t \in \mathcal{T}, u \in \mathcal{U}^{GEN} \quad (5.32)$$

$$y_{tu} \in \{0, \dots, \underline{D}_{tu}\}, \text{integer} \quad t \in \mathcal{T}, u \in \mathcal{U}^{GEN} \quad (5.33)$$

$$y_{tu} = 0, \quad t \in \mathcal{T}, u \in \mathcal{U}^{GEN} \quad (5.34)$$

$$x_{nst} = 1, \quad n \in \mathcal{N}, s \in \mathcal{S}, t \in \mathcal{T}^P \mid P_{nst}^{PA} = 1 \quad (5.35)$$

$$\sum_{s \in \mathcal{S}} P_{nst}^R x_{nst} = 1, \quad n \in \mathcal{N}, t \in \mathcal{T} \mid P_n^{DE} = 1 \quad (5.36)$$

$$x_{nst} = 0, \quad n \in \mathcal{N}, s \in \mathcal{S}, t \in \mathcal{T} \mid P_{ns}^S = 0 \quad (5.37)$$

Constraints (5.30) and (5.31) enforce binary restrictions on respectively x_{nst} and q_{ntp} . Constraints (5.32)-(5.34) ensure y_{tu} have proper boundaries. In constraints (5.32) the set \mathcal{U}^{GEN} contains only the demand type reciprocated by midwives, in constraints (5.33) the set \mathcal{U}^{GEN} contains only the demand type reciprocated by assistant nurses, while in constraints (5.34) the set \mathcal{U}^{GEN} contains all demand types except those reciprocated by midwives and assistant nurses. Constraints (5.35) ensures all preallocated shifts are allocated. This includes some particular kinds of shifts, vacations, off-days for employees with more than one employer, agreements made between the ward manager and employees and all shifts from the end of the previous planning period. Constraints (5.36) allocate all shifts on all days to employees deciding their own schedules. Constraints (5.37) force x_{nst} to 0 on all days for all shifts an employee cannot cover.

Chapter 6

Computational Study

The MWW scheduling model is implemented using the commercial optimisation software FICO[®] Xpress Optimisation Suite 7.8. The code is written in the Mosel-language and input-parameters are stored in Microsoft Excel- and .txt-files. Some parameters are pre-processed using Matlab, as they vary according to planning horizons and employees that exist in the instance. All instances in this chapter were solved on computers with 8 core Intel i7-3770 (3.40GHz) CPU and 16 GB RAM, running on Windows 7 Enterprise 64-bit Operating System. Output is presented in Excel-sheets that are processed to make the visual representation as understandable as possible. All computational studies in this chapter are created with real-life data that is used to create schedules for the period Desember 5th 2016 to June 11th 2017 (27 weeks) at MWW.

In Section 6.1, information regarding the instances is presented. In Section 6.2, two technical analyses are done, before two staffing analyses are performed and discussed in Section 6.3. Lastly, an excerpt from the output from running the full instance is presented and discussed and some key data is compared to the manually created schedule for the same period at MWW.

6.1 Instances

Our computational study was performed in parallel with the late stages of the current 9-step planning process happening at MWW, as it is presented in Section 2.5. We received all data necessary to create a schedule, including requests, vacations, etc. as soon as the information was available. Then, while the planning process progressed as usual at MWW, we created a full schedule for the 27 weeks and performed technical and economical analyses. All instances run in the computational study are based on the information for this period.

Some key data is presented in Tables 6.1-6.3 to provide the reader with insights to the size of the scheduling problem. We present the data for the instance that represents the full scheduling problem at MWW today.

Table 6.1: Key data to understand the size of the scheduling problem at MWW today.

	Data	Value
	Number of employees	69
	Number of weeks in planning horizon	27
	Number of different shifts	21

The total of 69 employees stated in Table 6.1 consists of 1 assistant ward manager, 48 midwives, 15 assistant nurses, 3 ward assistants and 2 secretaries. Many of them have different amounts of contracted work. There are 21 different types of shifts. They include 2 kinds of off-shifts and 19 different work shifts, of which only some are available for most employees. E.g. a typical midwife or assistant nurse is eligible for day shifts, evening shifts, night shifts and off-shifts, as well as two kinds of classes, but only on specific days. Thus a standard midwife is eligible to cover 7 types of shifts. A standard assistant nurse may also cover ward assistant shifts for ward assistants. In addition to these 7 very common types of shifts there are several preallocated types of shifts that a varying number of employees

are eligible to work. Examples are coordinating midwife shifts and shifts covering demand for midwives in rural areas around Trondheim. Furthermore, there are some shifts that are not preallocated, but only exist for very few employees.

In order to visualise the scheduling problem, it is useful to know the demand at MWW, and it is thus presented in Table 6.2. We only present the 7 most relevant demand types for the scheduling problem at MWW. There exists some other demand types, but the ones presented in Table 6.2 are the by far most relevant for the scheduling problem. Furthermore, there exists a maximum allowed over-coverage for each of the 7 demand types. The maximum allowed over-coverage is the ward manager's preferred limit for over-coverage, although it is not complied with for all days when schedules are created manually. The minimum and maximum demand for each demand type creates legal ranges for each demand type on all days. The legal ranges are presented in Table 6.2. Readers should note that a night shift covers the beginning of a day, thus a night shift on Monday begins at 22:15 at Sunday and ends at 07:45 on Monday.

Table 6.2: Legal ranges of demand for all the 7 relevant demand types.

Demand	Mon	Tue-Thu	Fri	Sat-Sun
Midwives day	6-8	6-8	6-8	4-6
Midwives evening	5-7	5-7	4-6	4-6
Midwives night	4-5	5-6	5-6	4-5
Assistant nurses day	3	3	3	2-3
Assistant nurses evening	3	3	2-3	2-3
Assistant nurses night	1	1	1	1
Ward assistant	1	1	1	1

It is also interesting to know the number of weekends and night shifts a midwife or assistant nurse is supposed to work throughout the planning horizon, depending on their contracts. In reality, however, the employees

work close to the right amount of weekends, but there are large variations in how many nights they work. This is due to the many personal inclinations, and is not considered unfair by the ward manager, but rather a symptom of different employees having different needs. The data is presented in Table 6.3.

Table 6.3: Amount of work during weekends and nights for standard midwives and assistant nurses throughout a 27 week planning horizon.

Contracted work (% of full-time)	Weekends	Nights
75-100	9	20
50-74	9	15
<50	9	10

The reader should also be familiar with the magnitude of the parameters in the objective function. They are the weighting parameters W^R and W^{DES} , that are listed in Table 6.4, that reward respecting requests and allocating evening-day shift-patterns to commuters, defined in Section 5.2. The value of the request reward W^R was chosen due to simplicity. After discussions with the ward manager, we agreed that the commuter reward was less than half as important as rewarding a request, resulting in the value 0.4 being seen as reasonable for W^{DES} .

Calibrating the proportion of the contracted work that each employee must work was more challenging. Forcing the proportion to be close to 1 seems to incur a higher run time and potentially leads to infeasibility. On the other hand, allowing the proportion to be far less than 1 produces schedules the ward manager sees as incomplete. Thus, we have agreed with the ward manager on the value seen in Table 6.4. Furthermore, another parameter stands out as key to calibrate when running the model. This is \underline{H}_n^{CW} , that represents the minimum hours each employee must work throughout the planning horizon. Some employees are not very flexible, making it preferable for the model to not allocate as much work to some employees

Table 6.4: Objective function weighting parameters and the proportion of the contracted work each employee must work.

Parameter	Value
Request reward W^R	1.00
Shift-pattern reward W^D	0.40
Proportion of contracted work	0.90

as it can. Thus, a minimum amount of work allocated to each employee is necessary. When designing the MWW scheduling model we have formulated \underline{H}_n^{CW} as a proportion of the contracted work for employee n . Thus, $\underline{H}_n^{CW} = \text{Contracted work} \cdot \text{proportion}$. Calibrating this proportion has been done after discussions with the ward manager. It should be noted that it is a policy at the ward for the staff manager or the scheduling group to allocate the remaining parts of an employee's contracted work in the end of the planning process if their individual schedule does not include all the hours they are paid to work. The ward manager has stated that they will continue to do so if they implement the MWW scheduling model, as it is deemed unfair if some employees receive paid off-days because of inflexible personal inclinations. In fact, some personal inclinations are too inflexible to make it feasible to create personal schedules with a proportion of contracted work close to 1.

6.2 Technical analyses

In this section, technical analyses are performed using the MWW scheduling model, in order to shed light on how changes in planning horizons and number of employees affect the model's run time. All adjustments that are made to the MWW scheduling model to create new instances are described in each subsection.

6.2.1 Length of planning horizon

How the planning horizon affects the MWW scheduling model’s run time is highly relevant, as the current planning horizons are set primarily because of the large amount of work associated with making a schedule. If the MWW scheduling model is implemented, there is a chance the ward manager would be interested in shortening the planning horizons. In that case, it would be advantageous to have insights to how this affects the run time of the MWW scheduling model. Creating satisfactory solutions quickly can facilitate a rapid iterative process that involves receiving feedback from the ward manager and the scheduling group and adjusting the model.

When we create instances of shorter planning horizons than that of the full 27 week instance, we create a schedule from the same starting date, but end it sooner. E.g. for a 12-week instance, we plan from December 5th 2016 to February 26th 2017. This also entails adjusting parameters that are connected to the planning horizon. Employees who begin working late in the planning horizon are removed from the instances if they cannot work at all or if they can only work a marginal part of the planning horizon, to avoid the instances becoming infeasible.

Table 6.5: Run times and gaps for the three instances of different planning horizons, spanning 27, 12 and 6 weeks for all 69 employees.

Instances (weeks)	Obj. func. value	Run time (s)	Optimality gap (%)	LP gap (%)
6	2652.8	3600	0.085	1.81
12	5365.0	3600	0.091	1.32
27	11896.8	3600	0.159	1.56

In Table 6.5 some key technical data is presented. The model is not allowed to run for longer than 3600 seconds, but stops running before if the optimal solution is found. In cases where the model does not find an optimal solution we look to the optimality gap to indicate how hard it is to solve an instance. We define the optimality gap as shown below.

$$\text{Optimality gap} = \frac{|\text{best bound} - \text{best solution objective value}|}{\text{best bound}}$$

Thus, the optimality gap describes the difference in our best solution and the best solution that potentially exists. As can be seen in Table 6.5, none of the three instances of different planning horizons reach optimality until they are stopped. The fact that the largest instances have higher optimality gaps seems very reasonable. The fact that the optimality gap is so similar for the different instances in Table 6.5 suggests that the MWW scheduling model scales very well for differences in planning horizons.

To get a perspective on how much the optimality gap affects the real life scheduling problem, we multiply the value of the optimality gap for the instance of 27 weeks in Table 6.5 with the value of the best bound for the same instance. Note that the best bound was found to be 11915.7 for this instance, although the value is not presented in Table 6.5. We get the value $0.00159 \cdot 11915.7 \approx 19$. In Section 6.1, the weighting parameter for one request, W^R , was set to 1. This means that the optimal schedule is in the region of respecting 0 to 19 requests more than the best solution we found by running the model for 1 hour. Furthermore, there is a total of 13230 requests for this instance. Thus, a difference of 0 to 19 requests would correspond to a change of 0 to 0.14% of the total respected requests. This is an amount of request fulfilment that would be extremely hard to attain through manual methods if the same scheduling rules were complied with.

Furthermore, the linear programming (LP) gap is defined as below.

$$\text{LP gap} = \frac{|\text{LP solution objective value} - \text{best solution objective value}|}{\text{LP solution objective value}}$$

The LP gap is interesting as it tells us how tight the model's LP solution space is relative to the integer programming (IP) solution space. Ideally, we want them to be as close as possible, which would imply less computer power is needed to find the optimal IP solution when the LP solution is

found. In the case of the real life scheduling problem, we see that the LP gap is 1.56%. The LP solution value is 12085.9, and the LP gap thus corresponds to approximately $0.0156 \cdot 12085.9 \approx 188.5$ requests, which seems to imply a fairly tight formulation. Furthermore, we see that the LP gap seems to stay quite stable for the different planning horizons in Table 6.5, implying that the MWW scheduling model is very scalable for different planning horizons.

Table 6.6: Number of constraints and variables created for the different instances of weeks, with all employees.

Instances (Weeks)	Number of constraints	Number of variables
6	29644	9158
12	61818	18521
27	143043	42186

In Table 6.6, we compare the number of constraints and variables produced in the three instances. As one would expect, the number of constraints and variables rises with the number of weeks in the planning horizon. The observed increase in optimality gaps in Table 6.5 is thus expected.

The fact that the MWW scheduling model seems to scale very well for differences in planning horizons could have several reasons, e.g. it could be related to the fact that there is extra labour available during most weekdays, possibly making it easier to perform changes in one part of the planning horizon without this affecting the schedule in other parts of the planning horizon. Furthermore, the scalability for different planning horizons suggests that the scheduling model's run time should not be a prominent factor when deciding on a ward's planning horizon, if it is to be implemented.

6.2.2 Number of employees

If the managers at St. Olavs Hospital decide that they want similar scheduling models as the MWW scheduling model at other wards, it is relevant to know how the sizes of wards affect how hard it is to solve the scheduling problems. Thus, it is interesting to test the scalability of the MWW scheduling model for different numbers of employees.

At MWW we have real-life information, making it very easy to scale down the number of employees, as opposed to creating fictional employees and adding them to the real-life data. Therefore we test instances of fewer employees than the full instance, while planning for the entire planning horizon of 27 weeks. We do as few changes to the instances as possible when we downscale the number of employees, but some are inevitable. For the instances of 48 employees, all employees other than the midwives are removed, but the average contracted work remains approximately equal to the case of 69 employees. Employees are removed by setting their contracted work to zero and for these employees variables and constraints are not initialised. Furthermore, the demand is removed for all employees other than the midwives, to reflect that only midwives should be allocated work shifts. For the instance of 24 midwives, all employees other than the midwives are still removed and half of the existing midwives are also removed from the problem. We choose to remove midwives that give half as many midwives eligible to work night-shifts as before. We also choose to keep midwives that give approximately the same average contracted work for the 24 midwives as it does in the case of 48 midwives. The demand for midwives is divided by 2 and rounded up to the nearest integer, resulting in the same or more work for each midwife in the instance of 24 midwives, compared to the instance of 48 midwives. For the instance of 24 employees, demand is rounded up only during weekdays, since all demands during weekends are even numbers. Since over-coverage is typical during weekdays and staffing during weekends is considered difficult, we believe such a scaling of demand to be appropriate for the instance of 24 employees. We run the instances for the full planning horizon of 27 weeks and get the results

in Table 6.7.

Table 6.7: Run times and gaps for the three instances of different numbers of employees, with the full 27 week planning horizon.

Instances (employees)	Obj. func. value	Run time (s)	Optimality gap (%)	LP gap (%)
24	4051.4	3600	0.025	2.57
48	8182.8	709	Optimal	1.73
69	11897.8	3600	0.159	1.56

The run times and optimality gaps in Table 6.7 do not correspond to our expectations. It seems reasonable to assume that as the number of employees decrease, the instance becomes easier to solve. Although Table 6.7 shows that the optimality gap is reduced for instances of 24 and 48 employees compared to 69 employees, we also see that the optimality gap is higher for 24 employees than for 48 employees. That is surprising. The number of constraints and variables created in the three instances in Table 6.8, indicates that an increased number of employees should make the instance harder to solve. We do see a slight decrease in the LP gap when increasing from 24 to 48 employees, which could lead to a lower run time for the instance with 48 employees. However, the instance with 69 employees has an even smaller LP gap than both other instances, and it still has the highest optimality gap of the three.

Table 6.8: Number of constraints and variables created for the different instances of employees, for 27 weeks.

Instances (employees)	Number of constraints	Number of variables
24	58830	15755
48	114747	30236
69	143043	42186

There are some potential explanations for the surprising results. One potential explanation is that we simply were lucky when running the instance for 48 employees, and that the model's branching strategy accidentally was especially successful for this single instance. On the other hand, it can be that some characteristic for the instance with 48 employees makes the scheduling problem especially easy to solve. For example, it is possible that the decrease in employees from 69 to 48 made it significantly easier for the model to solve the scheduling problem due to the reduction in the problem's size, but that the relative increase in demand per employee when reducing to 24 employees lead to the problem being much harder to solve. Such explanations are possible. However, we can combine reducing the planning horizons with reducing the number of employees and see if similar results occur for these instances.

6.2.3 Combining planning horizon and employees

In this subsection we combine the techniques of reducing the planning horizon and the number of employees. This results in the 9 instances presented in Tables 6.9 and 6.10. The instances are denoted with the numbers of weeks followed by 'W' and the number of employees followed by 'E'.

We note that the values in Table 6.10, the number of constraints and variables seem reasonable.

Inspecting the data in Table 6.9 we see that the instances with a 6 week planning horizon have an increase in the run time as the number of employees increase, and instances seem to be harder to solve for more employees. The same is true for the instances with a 12 week planning horizon, where run times and optimality gaps increase with the number of employees. If there is some special characteristic for the instances with 48 employees that makes it very easily solvable, it is not visible for the instances with 6 and 12 week planning horizons. This implies that the remarkably fast run for the instance of 27 weeks and 48 employees is likely a random event. More thorough analyses would be necessary to be certain of this, and it would

Table 6.9: Run times and gaps for all combinations of the instances of different time horizons and employees.

Instances	Obj. func. value	Run time (s)	Optimality gap (%)	LP gap (%)
6W24E	903.4	27	Optimal	2.96
6W48E	1779.4	239	Optimal	2.73
6W69E	2652.8	3600	0.085	1.81
12W24E	1824.4	29	Optimal	3.11
12W48E	3738.4	3600	0.078	2.17
12W69E	5365.0	3600	0.091	1.32
27W24E	4051.4	3600	0.025	2.57
27W48E	8182.8	709	Optimal	1.73
27W69E	11896.8	3600	0.159	1.56

be outside the scope of this master’s thesis to devote more attention to it.

However, if we ignore the instance of 27 weeks and 48 employees in Table 6.9, the scalability of the model for different staff levels is also very promising. How hard the instances are to solve seem to be affected somewhat more by the level of employees than by the planning horizon. E.g. it is notable how the relatively small instance ‘6W69E’ including 29644 constraints and 9158 variables has a higher optimality gap than the instance ‘27W24E’ that includes 58830 constraints and 15755 variables. The same conclusion can be drawn from comparing instances ‘12W24E’ and ‘6W48E’, where ‘12W24E’ is the largest instance, but also reaches optimality in less time. However, the increase in run times and optimality gaps when the number of employees are increased is still fairly small, implying that the model is very scalable for different different staff levels. This makes it reasonable to assume that scheduling models similar to the MWW scheduling model can be successfully designed for larger wards as well. Lastly, it is very interesting to see that for all planning horizons, the LP gaps seem to

Table 6.10: Number of constraints and variables created for all combinations of the instances of different time horizons and employees.

Instances	Number of constraints	Number of variables
6W24E	12393	3484
6W48E	23002	6441
6W69E	29644	9158
12W24E	40737	10846
12W48E	48568	13046
12W69E	61818	18521
27W24E	58830	15755
27W48E	114747	30236
27W69E	143043	42186

decrease as the number of employees increase. This could be a symptom of the model being developed for the full instance, in that case implying that we have been successful in tailoring the model to the ward's many unique characteristics.

6.3 Staffing analyses

In this section, staffing analyses are performed using the MWW scheduling model. This entails using the MWW scheduling model as an analytical tool, where the effects of decisions such as policy changes can be tested. In this sense the model is a useful tool in the tactical and strategic decision level, and not only usable for solving the scheduling problem.

The policies we consider are chosen after discussions with the ward manager and after receiving input from the board of the Regional Centre of Health Care Development (RSHU). All instances in this section is run for a maximum time of 24 hours, as some interesting instances prove challenging

to solve, but may still provide relevant information.

6.3.1 Open extra bed unit during weekends

The ward manager believes the weekends stand out as the clear bottleneck at MWW. Having enough employees to cover the work during the weekend is the main reason for the current staffing level at the ward. During the weekends, the ward closes one bed unit due to less available staff. This reduces the capacity of MWW for 7 patients, and thus facilitates a lower staff level than during the week-days. However, in practice the number of patients is the same during the weekends as the rest of the week, and most weekends the ward ends up opening the bed unit, without having the appropriate number of staff scheduled. At the same time, the ward experiences a considerable amount of over-coverage during the rest of the week. It is therefore very interesting to perform an analysis where we run the scheduling model with the bed unit open during weekends and implement some policies that can increase the number of weekend-shifts certain employees work. The increase in demand occurring when we open the extra bed unit is presented in Table 6.11 and the resulting legal ranges of demand coverage is presented in Table 6.12. The policies we implement are discussed with the ward manager at MWW and the board of RSHU, and are only implemented for employees the ward manager believes could volunteer to trade working extra weekends for receiving extra off-days. There are 51 such volunteers. The number of extra off-days the employees would need to voluntarily work the extra weekend is not known. We thus test different instances, with different number of days off for each extra weekend.

Instances and model extension

There are three characteristics that differ between the instances we try to open the extra bed unit for. They are listed below and the instances are presented in Table 6.13.

- The maximum number of extra weekends each volunteering employee

Table 6.11: The increase in demand for different days when we open the extra bed unit. Note that night shifts on Mondays begin Sunday 22:15.

Demand	Mon	Fri	Sat	Sun
Midwives day	0	0	1	1
Midwives evening	0	1	1	1
Midwives night	1	0	1	1
Assistant nurses day	0	0	1	1
Assistant nurses evening	0	1	1	1
Assistant nurses night	0	0	0	0
Ward assistant	0	0	0	0

Table 6.12: Legal ranges of demand when we open an extra bed.

Demand	Mon	Tue-Thu	Fri	Sat-Sun
Midwives day	6-8	6-8	7-9	5-7
Midwives evening	5-7	5-7	6-8	5-7
Midwives night	5-6	5-6	6-7	5-6
Assistant nurses day	3	3	4	3-4
Assistant nurses evening	3	3	3-4	3-4
Assistant nurses night	1	1	1	1
Ward assistant	1	1	1	1

can be awarded throughout the planning horizon. This should be as low as possible, to make the policy change realistic. We denote the maximum number of extra weekends each volunteering employee can be awarded throughout the planning horizon in the instances with 'M0' for 0 extra weekends, 'M1' for 1 extra weekend and 'M2' for 2 extra weekends.

- The number of earned off-days the volunteering employees get for working an extra weekend. E.g. if an employee that works an extra weekend earns 3 off-days, the employee will trade 3 off-days for working 2 weekend-shifts, thus acquiring a net gain of 1 off-day. The

number of earned off-days the volunteering employees get for working an extra weekend is denoted 'O3' for 3 extra off-days, 'O4' for 4 extra off-days etc.

- How the off-days traded for extra weekend work should be allocated. For some instances the earned off-days must be allocated within the subsequent 14 days of working the extra weekend as well as being counted as 8 hours worked per extra off-day. For subsequent allocation of off-days, the earned off-days are added to an approximation of the number of off-days the employee would normally have. E.g. an employee normally has 4 off-days during 14 days and earns 3 off-days by working an extra weekend. The employee would thus be allocated minimum 7 off-days in the subsequent 14 days after working the extra weekend. The other way of allocating earned off-days is to allocate them anywhere in the planning horizon. Each earned off-day is counted as 8 hours worked. The two ways of allocating earned off-days are denoted 'S' if they are allocated subsequent to the extra weekends and 'F' if they are allocated freely in the planning horizon.

We make minor changes to the MWW scheduling model to test the above mentioned policy changes. We define new variables, new parameters, adjust certain constraints and add new constraints. We define the variable δ_{nt} for all Sundays in the planning horizon. $\delta_{nt} = 1$ if employee n works an extra weekend containing Sunday t , 0 else. \bar{M}^{EW} is the maximum number of extra weekends any employee can work, M^{EO} is the number of extra off-days awarded for working one extra weekend, H^{EO} is the number of work hours an employee gains from each extra off-day, M_n^{NO} is the normal number of off-days employee n has during two weeks, and P_n^{EW} is 1 if employee n is eligible for working extra weekends, 0 else. The remaining variables, sets and parameters used in constraints (6.1) - (6.6) are defined in Section 5.2.

Constraints (6.1)- (6.2) are added to the mathematical formulation in the case of both methods of allocating extra off-days, i.e. both for instances denoted 'S' and for instances denoted 'F'. Constraints (6.3) are added to

the model only for the instances where off-days are allocated within the subsequent two weeks after an employee works an extra weekend, which are denoted with 'S'.

$$2\delta_{nt} \leq \sum_{s \in \mathcal{S}^W} (x_{nst} + x_{ns(t-1)}), \quad n \in \mathcal{N}, t \in \mathcal{T}^{SUN} \mid P_n^{EW} = 1 \quad (6.1)$$

$$\sum_{t \in \mathcal{T}^{SUN}} \delta_{nt} \leq \overline{M}^{EW}, \quad n \in \mathcal{N} \mid P_n^{EW} = 1 \quad (6.2)$$

$$\sum_{s \in \mathcal{S}^O} \sum_{\tau=t-13}^t x_{nst} - M^{EO} \delta_{nt} \geq M_n^{NO}, \quad n \in \mathcal{N}, t \in \mathcal{T}^{SUN} \mid P_n^{EW} = 1 \quad (6.3)$$

Constraints (6.1) ensure that if an extra weekend containing Sunday t is allocated to employee n , that employee works both Saturday and Sunday. Constraints (6.2) make sure no employees work more extra weekends than allowed. Constraints (6.3) ensure that employees are allocated extra off-days within two weeks after working an extra weekend. The extra off-days must be allocated so that an employee has more off-days during two weeks, than the employee normally has.

Constraints (5.7), (5.15) and (5.16) from Section 5.4 are changed for both cases of allocating extra off-days to employees working extra weekends, and the adapted formulations are found in constraints (6.4), (6.5) and (6.6), respectively.

$$\underline{H}_n^{CW} \leq \sum_{s \in \mathcal{S}^W} \sum_{t \in \mathcal{T}} H_{nst} x_{nst} + V_n^H + \sum_{t \in \mathcal{T}^{SUN}} M^{EO} H^{EO} \delta_{nt} \leq \overline{H}_n^{CW}, \quad n \in \mathcal{N} \quad (6.4)$$

$$\underline{M}_n^{WE} \leq \sum_{s \in \mathcal{S}^W} \sum_{t \in \mathcal{T}^{SUN}} x_{nst} - \sum_{t \in \mathcal{T}^{SUN}} \delta_{nt} \leq \overline{M}_n^{WE}, \quad n \in \mathcal{N} \quad (6.5)$$

$$\sum_{s \in \mathcal{S}^W} \sum_{\tau=0}^{\overline{M}_n^{NW}-1} x_{ns(t-7\tau)} - \sum_{\tau=0}^{\overline{M}_n^{NW}-1} \delta_{n(t-7\tau)} \leq 1, \quad n \in \mathcal{N}, t \in \mathcal{T}^{GEN} \quad (6.6)$$

Constraints (6.4) ensure that employees work a correct number of hours during the planning horizon. If an employee works an extra weekend, that employee will reach the upper limit of work hours through fewer work shifts, than employees not working extra weekends. Employees working extra weekends must therefore be allocated to additional off-days, since no more work hours can be allocated. Constraints (6.5) ensure that if an employee works extra weekends, the minimum number of weekends the employee must work is increased, and the maximum number of weekends the employee can work is increased. Constraints (6.6) ensure that employees can work weekends more frequently than normally allowed, if working an extra weekend. In constraints (6.6) the set \mathcal{T}^{GEN} only includes all Sundays where employee n has not registered requests contradicting this scheduling rule.

Results and analysis

In Table 6.13, 'Infeasible' means that the optimisation software found that no feasible IP solution exists for the instance, while 'feasible' means that an IP solution was found. The notation 'NIS' means that no integer solution was found for the instance, nor was infeasibility proven within the maximum time limit of 24 hours. Thus, we cannot say whether the instance has a feasible solution or not.

We implement the changes described in constraints (6.1)-(6.6) for the instances presented in Table 6.13. Running instance 'M0O3F' is equivalent to not implementing any new policies. However, the increased demand from opening the extra bed unit makes the instance infeasible, as shown in Table 6.13. Furthermore, in instance 'M1O3F' we ran the model with a maximum of 1 extra weekend allocated to each of the employees that can volunteer

Table 6.13: Key data for different instances that address MWW’s problem of opening the extra bed unit during the weekends. %Utilised work hours is defined as scheduled work, hours counted for vacations and extra off-days divided by the total contracted work.

Instance	Feasibility	%Requests	%Utilised work hours	#Extra weekends
M0O3F	Infeasible	NA	NA	NA
M1O3F	Infeasible	NA	NA	NA
M2O3F	Feasible	87.6	97.3	84
M2O4F	Feasible	87.1	97.3	83
M2O5F	Feasible	86.6	97.6	80
M2O6F	Feasible	84.3	97.9	80
M2O7F	Infeasible	NA	NA	NA
M2O3S	Feasible	87.6	97.3	84
M2O4S	Feasible	87.0	97.3	81
M2O5S	Feasible	85.7	97.6	81
M2O6S	NIS	NA	NA	NA

to work extra weekends. In this instance, the employees would only be awarded 3 extra days off, i.e. they trade working 2 extra days during the weekend for 3 extra days off during the planning horizon. This instance proved infeasible, making more generous trades e.g. 4 off-days per extra weekend infeasible as well.

When increasing the maximum number of extra weekends to 2 throughout the planning period of 27 weeks, the model is able to cover the demand, providing feasible solutions. Interestingly, we see that the model is able to provide feasible schedules where each extra weekend is traded for 3 to 5 subsequently allocated off-days or 3 to 6 freely allocated off-days. This is a very promising result, and can solve the ward’s problem of being understaffed during the weekends. It seems likely that many employees would welcome the chance to trade an extra weekend’s work for the most generous trades, thus making the policy changes likely to be implementable in

practise.

From the numbers of extra weekends in Table 6.13 we see that the number of extra weekends allocated are quite similar for the different instances, but somewhat decreasing as each extra weekend is traded for more off-days. This is logical, as the staff becomes a more and more scarce resource when the number of off-days traded for each extra weekend increases. An increase in off-days per extra weekends thus makes it harder for the model to find feasible solutions unless it allocates fewer extra weekends.

Furthermore, we see in Table 6.13 that the % Utilised work hours generally increases as the number of off-days traded for one extra weekend increases. This is due to the fact that some employees are not allocated their full contracted work by the MWW scheduling model (note that the minimum proportion of contracted work was set to 0.90 in Section 6.1.) When these employees volunteer to work extra weekends, their extra available hours, that normally would be allocated lastly by the ward manager, now become off-days that are counted as 8 hours of work each. We also note that the difference between the values for instances 'M2O3F' and 'M2O4F' and the difference between 'M2O3S' and 'M2O4S' is less than 0.1 % Utilised work hours. This is likely due to a decrease in over-coverage for some shifts, rather than the utilisation of previously unused staff.

We also note that the % requests respected is decreasing as the extra weekends are traded for more off-days. We believe this has two main reasons. The first reason is that when the % Utilised work hours approaches 1, there is less and less flexibility for the model to allocate employees to their requests. This implies that the proposed policy change gives MWW the opportunity of trading a reduction in employee preferences for a more effective utilisation of the staff. However, this idea of a trade-off between employee preferences and an effective utilisation of staff is not supported by the second main reason. In fact, most employees know well how many shifts and off-shifts they should work according to their contracts. The requests registered for this planning horizon are not adapted to the policy

change we present in this subsection, meaning that the allocation of extra weekends and extra off-days contradicts the requests of most employees. If this is the case, employees who volunteer to work extra weekend shifts, i.e. have a preference for trading extra weekend shifts for extra off-days, will get a reduced % requests respected. This does not support the notion of a trade-off between employee preferences and a more effective utilisation of staff, as the extra off-days are preferable for the employee. This makes the change in respected requests a flawed measure for the change in employee preference for this particular policy change.

The most notable weakness of our analysis of this policy change is that we cannot analyse the impact this policy change would have in the online operational level. Allocating more off-days than what is normally done entails reducing the over-staffing, which increases the robustness of the schedules. This implies that the extra weekend policy could make it hard for MWW to cover absent staff if too many off-days are traded for each extra weekend.

6.3.2 Reduced staff level

As mentioned before, there seems to be some over-coverage at the ward in the offline operational decision level, especially during weekdays. It would be interesting to see if it is feasible to create schedules for MWW where some employees are excluded from working and the original demand is maintained. This is relevant e.g. if some employees retire, and the management does not hire anyone to compensate for the loss of available labour.

Instances

To analyse this, we select some employees, set their contracted work to zero and see if the normal levels of demand can be covered. In this analysis we make note of which employees that are likely to be close to retirement and which are not. If an employee is entitled to senior off-days, this sug-

gests that the employee is close to retirement and the employee is thus characterised a senior employee. Out of the 69 employees working at the MWW, 10 are entitled to senior off-days, and these employees are thus senior employees. Amongst the senior employees, we focus on those that are midwives (1 senior midwife) and assistant nurses (6 senior assistant nurses).

Employees have different characteristics, e.g. contracted work, skill category, eligibility for night shifts, etc. Preliminary testing showed that these characteristics affect the number of employees that can be removed before the MWW scheduling problem becomes infeasible. Amongst all assistant nurses, 5 out of 15 (33.3 %) can work night shifts while 38 out of 48 (79.2%) midwives can work night shifts. Amongst the senior employees, 0 assistant nurses can work night shifts, but the one senior midwife can.

We run instances where different combinations of the employees are removed. The results of the runs we consider relevant are shown in Table 6.14. All employees with contracted work equal to zero are removed from the set of employees to create schedules for, thus no new constraints are needed to run the instances with reduced staff levels. When we exclude assistant nurses that are not seniors, we exclusively remove assistant nurses that can work night shifts, to examine how that affects the problem. The instances are denoted with the number of midwives removed followed by 'M' and the number of assistant nurses removed followed by 'A'. For senior assistant nurses and midwives we denote the number of removed employees followed by 'As' and 'Ms', respectively. We run one instance where no employees are excluded, but midwives are prevented from covering shifts for assistant nurses, denoted 'No helping midwife'. This instance is run without any alteration to the MWW scheduling model, except that we do not create the variables y_{tu} , used in constraints (9.3), (9.4) and (5.32)-(5.34) in Section 5.4.

Results and analysis

Table 6.14: Table presenting key data from the most interesting instances of reduced staff levels.

Instance	Feasibility	%Requests	#Assistant nurse shifts covered by midwives
1Ms1M	Feasible	89.9	148
2M	Feasible	89.7	189
1Ms2M	Infeasible	NA	NA
3M	Infeasible	NA	NA
2As	Feasible	90.0	284
3As	Feasible	89.7	343
3A	Feasible	89.3	265
4As	Infeasible	NA	NA
4A	Infeasible	NA	NA
1Ms2As	Feasible	89.6	254
1Ms3As	Infeasible	NA	NA
No helping midwives	Infeasible	NA	0

Looking at Table 6.14, it becomes apparent that it is possible to create feasible schedules with a reduced staff level. This is true for a reduction in employees for different combinations of skill-categories.

It is noteworthy how the instance where midwives may not cover assistant nurses' shifts is infeasible, implying assistant nurses are understaffed without the help of midwives. On the other hand, it is feasible to remove three assistant nurses as long as the midwives cover their shifts. This illustrates the flexibility the midwives offer at MWW. While the salaries of midwives are higher than that of assistant nurses, they compensate for this by being able to cover a lack of both other midwives and assistant nurses. This dif-

ference in flexibility is also highlighted when comparing instances '1Ms2M' and '3M' to '3As' and '3A'. The fact that fewer midwives than assistant nurses can be excluded from an instance is likely because there aren't as many other employees that can cover the demand for their skill category when they are removed.

It is interesting that we can remove an equal number of senior assistant nurses and other assistant nurses, because senior assistant nurses do not work night shifts. One could suspect that the employees would not be able to cover all the night shifts when removing three normal assistant nurses, but that is not the case. However, there is a notable difference in the number of assistant nurse shifts covered by midwives when comparing the instances '3A' and '3As', as many of the assistant nurses' night shifts had to be covered by midwives.

We find it interesting that the % requests respected remain quite stable for all instances of excluded employees. This is somewhat in contrast to our expectations, but may have several explanations. One explanation could be that most employees have a good sense of which shifts are needed to cover, and on an aggregated level make requests which coincide with the demand at MWW. Another explanation could be that midwives cover shifts for assistant nurses in such a way that both get their requests granted. E.g. if there are too many midwives requesting to work Thursday evening, but too few assistant nurses requesting it, midwives can cover the assistant nurses' shifts. Thus some extra midwives may get their requests for an evening shift respected, while the assistant nurses can instead be allocated shifts corresponding to their requests that Thursday. This demonstrates how the midwives' flexibility does not just affect the feasibility of the schedule, but also the request fulfilment. Furthermore, it is interesting that removing two senior assistant nurses gives a higher rate of respected requests than running the full instance for 24 hours, which will be presented in Section 6.4. This can be due to the fact that none of the flexible midwives are removed. Also, this can be due to the fact that these specific senior assistant nurses request shifts that are not easy to combine with other employees' shifts.

E.g. they do not work night shifts. When the employees are removed, many popular day and evening shifts becomes free, thus making it possible to respect a larger proportion of requests.

The results from Table 6.14 indicate that demand can be covered with a reduced staff level. Seeing as the instances consist of real-life data, this implies such a downscaling of the staff level is realistic. If an assistant nurse quits or retires, the demand can still be covered by midwives, cutting costs significantly in the offline operational decision level. The downscaling of staff seems to work well when considering the % requests respected due to the flexibility a midwife offers compared to an assistant nurse, although it may not be a popular policy if the ward becomes heavily reliant on midwives to cover assistant nurses' shifts.

If a midwife quits or retires, it is also feasible not to hire a new employee. This would be the most cost-effective alternative in the offline operational decision level, but would likely result in somewhat lower % requests respected as a flexible resource is removed. An alternative to not hiring would be to hire an assistant nurse instead of the midwife that quit. Thus, the demand for assistant nurses could be covered without the help of midwives, if the new assistant nurse could work night shifts. This would represent a middle ground between not hiring and hiring a new midwife with respect to both costs and flexibility offered to the ward when scheduling. Hiring a new midwife would obviously be the most flexible and expensive alternative in the offline operational decision level.

However, our analysis of reducing the staff level does not account for staff absence in the online operational level. Reducing the staff to the minimum feasible level in the offline operational scheduling problem could result in under-coverage in the online operational planning level, where staff absence is a costly and time-consuming component of the scheduling problem.

6.3.3 Open hours

Some members of the board of RSHU were interested in schedules where some of the employees only had 80% of their contracted work allocated as work shifts, i.e. 20% open hours. The idea is that these open hours would serve as a buffer for tackling staffing shortages due to sudden long-term sickness without using employees triggering overtime pay. Governmental regulations require that the employees are told when to work minimum 14 days beforehand. Thus, if an employee becomes sick and is expected to be unavailable for a longer period than 14 days, the ward manager may allocate the sick employee's shifts to the employees with 20% open hours.

Instances

The instances are named such that if 3 employees are allocated 20% open hours it is denoted '3E', if 6 employees are allocated 20% open hours it is denoted '6E' etc. We choose to exclusively model midwives as the employees with open hours, due to their advantageous flexibility to cover shifts for both midwives and assistant nurses. Furthermore, the chosen midwives are all contracted to work the maximum 35.5 hours per week, i.e working full time, meaning scheduling with 20% open hours will imply the chosen midwives are allocated approximately one shift less per week on average. Other than this, the midwives are chosen randomly, implying some of them may have parts of their schedules fixed by preallocated shifts and that their eligibility to work different shifts vary. Other than adding the open hours, the instances are the same as for the full instance of 27 weeks and 69 employees.

Results and analysis

From the instances presented in Table 6.15 we can clearly see that it is possible to create schedules with different ways of allocating open hours. However, it is hard to analyse which of the different instances in Table 6.15 that are most useful to cope with the real-life problems that occur at MWW. A way to estimate how robust the three instances would be if

sudden long-term absence occurs, is to investigate on which days some of the chosen midwives have off-days. If one or more of the chosen midwives are allocated off-shifts on a particular day, this could mean they can cover for the absent employee. We are therefore interested in the number of days some of the chosen employees have an off-day during the planning horizon.

Table 6.15: Data for different instances open hours is implemented for.

Instances	%Requests	Days someone has off-shifts	Days no-one has off-shifts	%Days someone has off-shifts
3E	89.8	165	24	87.3
6E	89.8	185	4	97.9
9E	89.8	189	0	100

In Table 6.15, we see that by just choosing three midwives to be allocated open hours, one or more of them have off-shifts on as much as 87.3% of the days. By choosing six, one or more of the chosen midwives have an off-shift on as much as 97.7% of the days. Choosing nine employees involves at least one midwife having an off-shift on all days in the planning horizon. From these results, one would think allocating open hours to six employees would significantly improve the schedule's ability to tackle long-term absence for one employee. However, it is too simple to assume that midwives with open hours can always be allocated any shift during that day. The shift has to fit the midwife's schedule and respect the scheduling rules. An example of this problem is provided in Figure 6.1. Notice that on the 25th of January only 'Midwife 9' has an off-day. This off-day succeeds a day shift, and because of scheduling rules 'Midwife 9' cannot be allocated a night shift on the 25th of January. As such, if a long-term absent employee was scheduled to work a night shift on 25th of January, that shift cannot be covered by any of the nine employees with open hours.

The fact that one of the midwives with open hours having an off-shift on each day does not ensure that they can cover any shift that day is a prob-

	20. jan. 47	21. jan. 48	22. jan. 49	23. jan. 50	24. jan. 51	25. jan. 52	26. jan. 53	27. jan. 54	28. jan. 55	29. jan. 56	30. jan. 57	31. jan. 58	1. feb. 59	2. feb. 60
Midwife 1	D	F2	F1	F2	D	A	F1	A	A	D	F2	F2	A	A
Midwife 2	D	F2	F1	F2	N	N	A	D	F2	F1	A	F2	F2	F2
Midwife 3	F2	F2	F1	F2	N	N	N	N	F2	F1	V	V	F2	F1
Midwife 4	A	D	A	F2	DPR	A	F2	D	F2	F1	F2	DPR	F2	D
Midwife 5	D	F2	F1	A	D	D	F2	F1	N	N	N	F2	F2	A
Midwife 6	F2	A	D	D	F2	A	D	F2	F2	F1	N	F2	D	F2
Midwife 7	N	F2	F1	A	F2	DU	F1	DU	A	D	F2	A	F2	D
Midwife 8	F2	F2	F1	F2	F2	A	F2	A	F2	F1	A	F1	A	F2
Midwife 9	F2	F2	F1	A	D	F2	F2	F1	N	N	N	A	F2	A

Figure 6.1: Excerpt showing two weeks of the schedule for the midwives chosen to be allocated open hours in the in the '9E' instance. Shifts 'F1' and 'F2' are off-shifts and are highlighted in the figure.

lem. It implies that having several midwives with open hours that have off-days on each day is better than just having one. If many midwives with open hours have an off-day at a day that must be covered, it is more likely that at least one of them can cover the shift for the absent employee. It is thus interesting to see how many employees that have off-days on how many days for the different instances. This is illustrated in Figure 6.2. We see that the more employees with open hours, the higher is the chance that many of them have off-shifts each day. This should also imply a higher chance for at least one of them being able to cover for an absent employee.

However, even for the instance '9E', we cannot be certain that any sudden long-term absence can be covered for, as is the case shown in Figure 6.1. Furthermore, the absent employee would leave weekend-shifts open, forcing someone to work extra weekends. Weekend-shifts are not included in the open hours-strategy, meaning the strategy does not address the part of the schedule the ward has most problems covering. The open hours will also imply extra work. As more midwives that are given open hours, the ward manager will have to allocate extra shifts continuously. Furthermore, the midwives with open hours could become frustrated with regularly being allocated new shifts, affecting morale at the ward.




Figure 6.2: During each day in the planning horizon, midwives with open hours are allocated to a number of off-shifts. The sum of these off-shifts vary on each day, and the columns show the number of times different sums occur in the planning horizon.

To be certain of covering all shifts left uncovered in the case of sudden long-term absence, the scheduling model would have to allocate the open hours more strategically, making sure that some midwife was always available to work any shift at any day. This could be done if a new shift was created. We dub the shift 'availability shift'. The availability shift can be switched for any work shift by the ward manager two weeks before they occur. Each midwife with open hours should be allocated one of these shifts each week on average, and the shifts should be scheduled so that no matter what shift it was swapped for, the scheduling rules are respected. Furthermore, the midwives with open hours would not have to worry that their off-shifts could be swapped for work-shifts two weeks beforehand.

It would probably be possible to implement such availability shifts, but it would take a large amount of resources, as many midwives are needed to cover all days with these shifts. If seven midwives are allocated with the availability shifts and these are allocated perfectly, then there should be enough availability shifts to cover the need. However, the midwives with open hours would need to have some availability shifts during the weekends as well, meaning they would not work as many normal shifts during the weekend as before. This is problematic, as the weekend is already a bottleneck with regards to covering demand. To get a better idea of how this would work in practise, it could be implemented in the MWW scheduling model. However, as this master's thesis is primarily concerned with the offline operational decision level, it is outside the scope of the thesis to do so. Furthermore, sudden long-term absence is, according to the ward manager, not a large problem at MWW. At ward's that do have such problems, implementing availability shifts could be more interesting, especially if the activity is lower during weekends than the rest of the week, making it easier to allocate employees more freely.

We thus conclude that the idea of allocating open hours is interesting, but that the open hours should be allocated in a strategic way. It is not enough to create schedules where there exists excess unused labour, the

employees must also be scheduled in such a way that they can be allocated new shifts when unforeseen changes occur, without breaking scheduling rules. Furthermore, as MWW is not particularly bothered with sudden long-term absence, open hours are more likely to be useful in other wards.

6.4 Comparing schedules

In this section we present some key facts about the full instance produced by our scheduling model and compare it to the manually made schedule that has been created using the current planning process.

We run the full instance of 27 weeks and 69 employees. This results in the key data resented in Table 6.16. Furthermore, the run is presented graphically in Figure 6.3. The objective function value is plotted for the different run times. The red graph represents the best solution that is found, while the yellow graph represents the best bound. Green squares demonstrate that a new integer solution has been found. The model only uses 263 seconds to find the first integer solution, and even better solutions are obtained almost immediately after. This demonstrates that the MWW scheduling model is useful for creating good schedules very fast. Furthermore, the model is able to obtain some even better solutions as it continuous to run, resulting in the key data found in Table 6.16.

Table 6.16: Key data for the full instance.

Obj. func. value	Run time (s)	Optimality gap (%)	LP gap (%)	%Requests respected
11908.6	86400	0.036	1.47	89.8

Furthermore, we present a small excerpt from the output produced when running the full instance in Figure 6.4. For every employee, there is one row describing their proposed schedule for certain days and a second row presenting the requests they have registered. We have highlighted all re-




Figure 6.3: Graphic presentation of running the full instance of 69 employees for a planning horizon of 27 weeks.

spected requests in blue. Note that the shifts 'F1' and 'F2' both denote off-days, thus making a request for an 'F2'-shift respected if the employee is awarded an 'F1'-shift and vice versa. We clearly see that in the excerpt, most proposed shifts correspond to the employees' requests. However, we can also see that two employees who requested to work during the weekend the 4th and 5th of March were allocated off-days instead. This is because there was a large number of employees that wanted to work that specific weekend, forcing the model to allocate some employees to work other weekends instead.

While we have developed the MWW scheduling model and performed analyses with it, the manual 9-step planning process described in Section 2.5 has taken its usual course and has recently reached step 8. We thus present some key information about the manual schedule in Table 6.17 and compare it to similar information about the schedule created by the MWW scheduling model.

One of the strongest features of the manual schedule created through the current 9-step planning process is that it ensures employees' influence on all

6.4 Comparing schedules

		27.feb	28.feb	01.mar	02.mar	03.mar	04.mar	05.mar	06.mar	07.mar	08.mar	09.mar	10.mar	11.mar	12.mar
		85	86	87	88	89	90	91	92	93	94	95	96	97	98
Midwife 1	P:	V	V	F2	F1	V	V	V	F2	F2	N	N	N	F2	F1
Midwife 1	R:	V	V	F2	F2	V	V	V	F2	F2	N	N	N	F2	F1
Midwife 2	P:	VK	A	F2	F2	F1	N	N	V	V	F2	F2	VK	F2	F1
Midwife 2	R:	VK	A	D	F2	F2	N	N	V	V	F2	F2	VK	F2	F2
Midwife 3	P:	A	F2	A	D	D	F2	F1	D	F2	K	A	D	F2	F1
Midwife 3	R:	A	F2	A	D	D	F2	F2	D	F2	K	A	D	F2	F2
Midwife 4	P:	D	D	D	A	F2	F2	F1	F1	F2	D	A	F2	A	D
Midwife 4	R:	D	D	D	A	F2	F2	F2	F2	F2	D	A	F2	N	N
Midwife 5	P:	N	F2	F2	N	N	F2	F1	F2	N	N	F2	F2	F2	F1
Midwife 5	R:	N	F2	F2	N	N	F2	F2	F2	N	N	F2	F2	F2	F2
Midwife 6	P:	F2	A	D	A	F2	F2	F1	A	F2	F1	F2	A	D	A
Midwife 6	R:	F2	A	D	A	F2	F2	F2	N	N	F2	F2	D	D	A
Midwife 7	P:	D	F2	F2	F2	F2	F2	F1	D	F2	K	D	D	F2	F1
Midwife 7	R:	D	F2	F2	F2	F2	A	D	D	F2	K	D	D	F2	F2
Midwife 8	P:	F2	N	A	F2	F1	N	N	N	F2	F2	F2	F2	F2	F1
Midwife 8	R:	F2	N	N	F2	F2	N	N	N	F2	F2	F2	F2	F2	F2
Midwife 9	P:	N	F2	A	A	F2	F2	F1	F2	F2	F2	N	N	F2	F1
Midwife 9	R:	N	F2	A	A	F2	F2	F2	F2	D	F2	N	N	F2	F2
Midwife 10	P:	F2	D	D	D	F2	F2	F1	D	D	F2	D	F2	F2	F1
Midwife 10	R:	F2	D	D	D	F2	F2	F2	D	D	F2	D	F2	F2	F1
Midwife 11	P:	F2	D	D	F2	F2	F2	F1	F2	A	D	D	F2	F2	F1
Midwife 11	R:	F2	D	D	F2	F1	A	D	F2	A	D	D	F2	F2	F2
Midwife 12	P:	F1	A	F2	A	F2	A	D	F2	F2	A	D	F2	F2	F1
Midwife 12	R:	F2	A	F2	A	F2	A	D	F2	D	A	D	F2	F2	F1
Midwife 13	P:	N	F2	F2	F2	F1	N	N	N	F2	F2	N	N	F2	F1
Midwife 13	R:	N	F2	F2	F2	F2	N	N	N	F2	F2	N	N	F2	F2
Midwife 14	P:	F2	F2	F2	F2	F2	F2	F1	F2	F2	F2	N	N	F2	F1
Midwife 14	R:	N	F2	N	N	F2	F2								
Midwife 15	P:	NVK	F2	N	N	F2	F2	F1	F2	NVK	NVK	F2	F2	F2	F1
Midwife 15	R:	NVK	F2	N	N	F2	F2	F2	F2	NVK	NVK	F2	F2	F2	F2
Midwife 16	P:	F2	N	N	A	F2	F2	F1	A	F2	D	F2	F1	A	D
Midwife 16	R:	F2	N	N	A	F2	F2	F2	A	D	D	F2	F2	A	D

Figure 6.4: A small excerpt from an anonymous version of the output. We chose 14 random employees during 2 random weeks, to help visualise the output of the MWW scheduling model. The rows containing "P:" shows the planned schedule while rows containing "R:" shows the employee's requests.

Table 6.17: Information for comparing the manual schedule to the one created by the MWW scheduling model.

Key information	Manual schedule	Scheduling model
Employee influence ensured	Yes	Yes
%Requests respected	85.3	89.8
Max over-coverage respected	No	Yes
All scheduling rules respected	No	Yes
Unbiased shift-allocation	No	Yes
Unfair bartering	Yes	No
Time to create schedule	6 weeks	5min - 24 h

the schedules. This has been of great importance to us when developing the MWW scheduling model as well. Furthermore, the ward manager states that the employees' influence is guaranteed due to the combination of the model prioritising employees' preferences and the possibility for the ward manager and scheduling group to make changes to the produced schedules. This focus on the employees' preferences has also lead to us creating schedules that have a significantly higher proportion of requests respected than the manual schedule. The manual schedule's proportion is indeed also good, highlighting both that employees at MWW seem to make reasonable requests that are aligned with their contracted work and that the management at MWW allow a lot of time and effort to go into creating the manual schedules. However, the proportion of respected requests in the manual schedule is perhaps somewhat higher due to the fact that more over-coverage was allowed. The manual schedule has allocated a greater over-coverage for some shifts on some days than we have allowed the MWW scheduling model to do. This can potentially make it easier to respect many requests, as some shifts on some days are very popular among most employees, but implies that the over-coverage is less evenly distributed and the robustness of the schedule in the online operational level can be reduced.

It should also be mentioned that the MWW scheduling model ensures that no scheduling rules are broken through hard constraints, while this is more

challenging using manual methods. Some scheduling rules tend to be broken in the current 9-step planning process. An example from this manual schedule is that an employee by accident accepted working 5 consecutive night-shifts when bartering. Another strength the MWW scheduling model possesses is its lack of bias when allocating shifts and is thus fair, or 'equally unfair', to all employees. Self-scheduling that includes bartering is a political process. Removing the bartering process is perhaps the single most efficient measure to make the current planning process at MWW more fair. Lastly, the MWW scheduling model creates good schedules very fast. Although some minor adjustments are likely to be done when new instances are run, it is still remarkably faster than the current system and it ties up much fewer employees because the bartering process is removed, thus increasing the resource utilisation of human capital.

We have had several meetings with the ward manager at MWW when developing the MWW model, as mentioned in Section 4.1. The final schedule we have created for the full instance is described by the ward manager as clearly preferable to the manual scheduling that is performed today. The ward manager has stated that the schedules created by the MWW scheduling model is a very good and useful tool to create schedules. Furthermore, the ward manager requests a similar tool for other wards.

Chapter 7

Concluding remarks

In this master's thesis we have analysed the real life planning process at Maternity Ward West (MWW) and modelled their scheduling problem through close co-operation with the ward manager. The model is designed to create useful schedules, securing employee involvement through a large focus on employees' preferences and fairness, while respecting all relevant scheduling rules.

The scheduling model we have designed finds good integer solutions within few minutes and is close to reaching optimality within 24 hours, with an optimality gap of 0.036% for the full real-life instance. The schedules the model produces is able to fulfil 89.8% of employees' requests, and excess labour is successfully distributed between shifts and days by setting maximum levels of over-staffing, thus securing a schedule that is more robust to changes in the online operational decision level. The schedules produced by the model respects all work regulations, ward policies, preferred practises and agreements with trade unions and allocates shifts in an unbiased way. The scheduling model makes the current bartering process at MWW unnecessary, thus facilitating the removal of the unfair bartering.

We have shown that the model is very scalable for different planning hori-

Concluding remarks

zons and different staff levels, implying that the model can be adjusted to work well for similar wards as well. Furthermore, we have shown that the MWW scheduling model can be used as a management tool for tactical and strategic decisions, by implementing changes in policies and staff levels for different instances and producing feasible schedules for these instances. We have proposed and shown the feasibility of policies which realistically allows the MWW to open a much needed extra bed unit during weekends using the MWW scheduling model as an analysis tool. We have also performed similar analyses that suggests it would be realistic to reduce the current staffing level while servicing the current demand. Furthermore, we have performed analyses discussing the possibility of strategic reduction of certain employees' work to let them cover sudden long-term absence at the ward.

We have succeeded in our goal of creating a decision support tool that solves the scheduling problem by producing schedules for MWW of such quality that they are preferable to manually made schedules. The best testament to this is that the ward manager at MWW states that she wants to use this tool and that she requests a similar tool for other wards.

Chapter 8

Future Work

The work presented in this master's thesis has uncovered several interesting areas for future work. As mentioned in Chapter 1, our academic contribution is mostly related to presenting a success story of utilising existing theory in developing a useful scheduling model for a highly detailed real-life scheduling problem. Thus, the areas for future work that have emerged are mostly specific for the situation at MWW and St. Olavs rather than in the academic field. However, some of the topics may very well be typical in the field of nurse scheduling.

After working with the scheduling problem at MWW, it is compelling to explore how to model other hierarchical decision levels to a larger extent. It would be very useful for MWW, and other wards like it, to include the online operational decision level in the model development, thus making the model capable of producing more robust schedules. In the same way, there is a large potential for improvement if tactical and strategic decisions are included in the scope of the scheduling model. Examples of such changes are to include the decision of staffing levels or to make potential changes in policy optional for scheduling models. It would also be very interesting to perform increasingly aggregated planning, such as including several similar wards in one planning problem. E.g. planning for Maternity Ward West

and Maternity Ward East as one large unit, with some common resources, thus increasing the resource utilisation of human capital further.

Including several new aspects in the model could increase the run times of models, making other optimisation approaches relevant to explore. As the MWW scheduling model must be adaptable to changes at the ward, heuristic methods seem promising. Related literature discusses several heuristic and hybrid heuristic methods for the nurse scheduling problem. Applying a combination of heuristic methods with an exact approach would definitely be interesting.

In order to use the schedules produced by the scheduling model, there should be a thorough implementation process, where the model is integrated with the existing IT-system at St. Olavs Hospital. Furthermore, there are some changes to the planning process at MWW that would be interesting to analyse. One example is to create a different request system, where each employee gets a limited number of requests, possibly with the alternative of assigning different weights to different requests. Another example is to introduce 12-hour shifts during the weekends, thus reducing the number of employees that have to work each weekend.

Lastly it would be very interesting to include trade unions in the development of a scheduling model. Currently, their role in the planning process resembles that of a monitor, evaluating finished schedules. It would be much more efficient to include their ideas and opinions of what characterises a high-quality schedule directly into the scheduling model.

Chapter 9

Compressed model

In this chapter the entire model from Chapter 5 is presented without the explanations provided in Chapter 5.

9.1 Definitions

9.1.1 Indices

n	-	employee
s	-	shift
t	-	day
u	-	demand type
k	-	week
p	-	shift-pattern

9.1.2 Sets

\mathcal{N}	-	Set of employees.
\mathcal{N}_u^U	-	Set of employees in the skill category reciprocating demand type u , $\bigcup_{u \in \mathcal{U}} \mathcal{N}_u^U = \mathcal{N}$.

- \mathcal{N}^{DES} - Set of employees who desire shift-patterns rewarded in the objective function.
 \mathcal{N}^{GEN} - A generic set of employees explained further whenever used.
 \mathcal{S} - Set of shifts.
 \mathcal{S}^W - Set of shifts which are work-shifts, $\mathcal{S}^W \subset \mathcal{S}$.
 \mathcal{S}^O - Set of shifts representing all off-days, $\mathcal{S}^O \subset \mathcal{S}$. $\mathcal{S}^O \cup \mathcal{S}^W = \mathcal{S}$.
 \mathcal{S}_u^U - Sets of shifts covering demand type u , $\bigcup_{u \in \mathcal{U}} \mathcal{S}_u^U \cup \mathcal{S}^O = \mathcal{S}$.
 \mathcal{S}^{GEN} - A generic set of shifts explained further whenever used.
 \mathcal{K} - Set of weeks in the planning horizon.
 \mathcal{T} - Set of days in the current planning horizon.
 \mathcal{T}^{SUN} - Set of Sundays in the planning horizon.
 \mathcal{T}^{MON} - Set of Mondays in the planning horizon.
 \mathcal{T}^C - Set of days during Christmas.
 \mathcal{T}^P - Set of days in the current and previous planning horizon. \mathcal{T}^P includes positive and negative values of t where values of $t < 1$ are days in the previous planning horizon and $t = 0$ is the last day in the previous planning horizon. $\mathcal{T} \subset \mathcal{T}^P$.
 \mathcal{T}_k - Set of the days in week k , $\bigcup_{k \in \mathcal{K}} \mathcal{T}_k = \mathcal{T}$.
 \mathcal{T}^{GEN} - Generic set of days explained further whenever used.
 \mathcal{U} - Set of demand types.
 \mathcal{U}^{GEN} - Generic set of demand types explained further whenever used.
 \mathcal{P}_n^{ILL} - Set of shift-patterns which are illegal to allocate to employee n .
 \mathcal{P}_n^{MAN} - Set of shift-patterns which are mandatory to allocate to employee n , under certain circumstances.
 \mathcal{P}_n^{AFT} - Set of shift-patterns that force other shift-patterns to succeed it and the succeeding shift-patterns, for employee n . This is explained thoroughly in example B in Section 5.4.
 \mathcal{P}_n^{DES} - Set of desirable shift-patterns which should be rewarded in the objective function when allocated to employee n .

9.1.3 Parameters

Weighting parameters

- W^R - Weight parameter rewarding requests.
 W^D - Weight parameter rewarding desirable shift-patterns.

Limit parameters

- \underline{D}_{tu} - Minimum number of employees needed to cover the demand for demand type u on day t .
 \overline{D}_{tu}^{OC} - Maximum number of employees allowed to over-cover demand type u on day t .
 \overline{M}_n^{CW} - Maximum number of consecutive work shifts for employee n .
 \overline{M}_{ns}^{CS} - Maximum number of consecutive shifts s for employee n .
 \overline{M}_{ns}^{TS} - Maximum number of shifts s employee n can work in the planning horizon.
 \underline{M}_{ns}^{TS} - Minimum number of shifts s employee n must work during the planning horizon.
 \overline{M}_n^{WE} - Maximum number of weekends employee n can work during the planning horizon.
 \underline{M}_n^{WE} - Minimum number of weekends employee n must work during the planning horizon.
 \overline{M}_n^{NW} - Employee n works one in every \overline{M}_n^{NW} weekends.
 \overline{M}_s^{FT} - Maximum deviation in number of unpopular shifts s allocated to each employee within a certain skill category.
 \overline{H}_n^{7D} - Maximum number of hours employee n can work during any 7-day interval.
 \overline{H}_n^{CW} - Maximum number of hours employee n can work during the planning horizon.
 \underline{H}_n^{CW} - Minimum number of hours employee n must work during the planning horizon.

Indicating parameters

- P_{nst}^{PA} - 1 if employee n should have shift s preallocated on day t , 0 else.
 P_{nst}^R - 1 if employee n requests shift s on day t , 0 else.
 P_n^{DE} - 1 if employee n should always be allocated its requests, 0 else.
 $P_{s_1s_2}^A$ - 1 if there is sufficient time between shifts s_1 and s_2 on days $t - 1$ and t respectively for an employee to work them both, 0 else.
 $P_{s_1s_2}^{F1}$ - 1 if there is sufficient time between shifts s_1 and s_2 on days $t - 2$ and t respectively for an employee to be allocated to an 'F1' off-day on day $t - 1$, 0 else.
 $P_{n_1n_2}^{TOG}$ - 1 if employees n_1 and n_2 can work together during the same shift.
 $P_{u_1u_2}^U$ - 1 if $\mathcal{S}_{u_1}^U = \mathcal{S}_{u_2}^U$, i.e., 1 if the two demand types, u_1 and u_2 , are serviced by the same set of shifts, 0 else.
 P_{ns}^S - 1 if employee n can work shift s , 0 else.
 $P_{np_1p_2}^{AFT}$ - 1 if shift-pattern p_2 must be allocated to employee n immediately after shift-pattern p_1 is allocated, 0 else.

Dynamic parameters

- S_{stp}^{ILL} - Illegal shift-patterns. $S_{stp}^{ILL} = 1$ if shift-pattern p includes shift s on day t of the shift-pattern, and 0 else.
 S_{stp}^{AFT} - Shift-patterns that must be allocated after certain shift-pattern is allocated. $S_{stp}^{AFT} = 1$ if shift-pattern p includes shift s on day t of the shift-pattern, and 0 else.
 S_{stp}^{OBJ} - Shift-patterns that are rewarded in the objective function. $S_{stp}^{OBJ} = 1$ if shift-pattern p includes shift s on day t of the shift-pattern, and 0 else.
 S_{ntp}^{MAN} - Shift-pattern that must be allocated to an employee if the employee is allocated any of the shifts included in the shift-pattern. $S_{stp}^{MAN} = 1$ if shift-pattern p includes shift s on day t of the shift-pattern, and 0 else.

General parameters

- H_{nst} - Calculated duration of shift s , in hours, for employees n on day t .
- V_n^H - Due to vacations, the number of hours employee n should work is reduced by V_n^H .
- V_{ns}^S - Due to vacations, the number of shifts s employee n should work is reduced by V_{ns}^S .
- $\bar{\tau}_p$ - Number of days spanned by shift-pattern p .

9.1.4 Variables

- x_{nst} - 1 if employee n works shift s on day t , 0 else.
- q_{ntp} - 1 if employee n works a desirable shift-pattern p ending on day t , 0 else.
- y_{tu} - Integer variable taking value when a demand type u is covered on day t by an employee not in \mathcal{N}_u^U . This is relevant for midwives and assistant nurses.

9.2 Objective function

$$\max Z = W^R \sum_{n \in \mathcal{N}} \sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{T}} P_{nst}^R x_{nst} + W^D \sum_{n \in \mathcal{N}^{DES}} \sum_{t \in \mathcal{T}} \sum_{p \in \mathcal{P}_n^{DES}} q_{ntp} \quad (9.1)$$

9.3 Constraints

9.3.1 Covering demand

$$\sum_{s \in \mathcal{S}} x_{nst} = 1, \quad n \in \mathcal{N}, t \in \mathcal{T} \quad (9.2)$$

$$y_{tu_1} + y_{tu_2} = 0, \quad t \in \mathcal{T}, u_1, u_2 \in \mathcal{U} \mid P_{u_1 u_2}^U = 1 \quad (9.3)$$

$$\underline{D}_{tu} \leq \sum_{n \in \mathcal{N}_u^U} \sum_{s \in \mathcal{S}_u^U} x_{nst} + y_{tu} \leq \underline{D}_{tu} + \overline{D}_{tu}^{OC}, \quad t \in \mathcal{T}, u \in \mathcal{U} \quad (9.4)$$

$$\sum_{n \in \mathcal{N}_u^U \cap \mathcal{N}^{GEN}} \sum_{s \in \mathcal{S}_u^U} x_{nst} \geq 1, \quad t \in \mathcal{T}, u \in \mathcal{U}^{GEN} \quad (9.5)$$

$$x_{n_1st} + x_{n_2st} \leq 1, \quad n_1, n_2 \in \mathcal{N}, s \in \mathcal{S}, t \in \mathcal{T} | P_{n_1n_2}^{TOG} = 0 \quad (9.6)$$

9.3.2 Work hours

$$\underline{H}_n^{CW} \leq \sum_{s \in \mathcal{S}^W} \sum_{t \in \mathcal{T}} H_{nst} x_{nst} + V_n^H \leq \overline{H}_n^{CW}, \quad n \in \mathcal{N} \quad (9.7)$$

$$\sum_{s \in \mathcal{S}^W} \sum_{\tau=t-6}^t H_{nst} x_{ns\tau} \leq \overline{H}_n^{TD}, \quad n \in \mathcal{N}, t \in \mathcal{T} \quad (9.8)$$

9.3.3 Required rest

$$x_{ns_1(t-1)} + \sum_{s_2 \in \mathcal{S} | P_{s_1s_2}^A = 0} x_{ns_2t} \leq 1, \quad n \in \mathcal{N}, s_1 \in \mathcal{S}, t \in \mathcal{T} \quad (9.9)$$

$$x_{ns_1(t-2)} + x_{n'F1'(t-1)} + \sum_{s_2 \in \mathcal{S} | P_{s_1s_2}^{F1} = 0} x_{ns_2t} \leq 2, \quad n \in \mathcal{N}, s_1 \in \mathcal{S}, t \in \mathcal{T} \quad (9.10)$$

$$\sum_{t \in \mathcal{T}_k} x_{n'F1't} = 1, \quad n \in \mathcal{N}, k \in \mathcal{K} \quad (9.11)$$

$$\sum_{s \in \mathcal{S}^W} \sum_{\tau=t-\overline{M}_n^{CW}}^t x_{ns\tau} \leq \overline{M}_n^{CW}, \quad n \in \mathcal{N}, t \in \mathcal{T} \quad (9.12)$$

$$\sum_{\tau=t-\overline{M}_{ns}^{CS}}^t x_{ns\tau} \leq \overline{M}_{ns}^{CS}, \quad n \in \mathcal{N}, s \in \mathcal{S}^W, t \in \mathcal{T} \quad (9.13)$$

9.3.4 Weekends

$$\sum_{s \in \mathcal{S}^W} (x_{ns(t-1)} - x_{nst}) = 0, \quad n \in \mathcal{N}, t \in \mathcal{T}^{SUN} \setminus \mathcal{T}^C \quad (9.14)$$

$$\underline{M}_n^{WE} \leq \sum_{s \in \mathcal{S}^W} \sum_{t \in \mathcal{T}^{SUN} \setminus \mathcal{T}^C} x_{nst} \leq \overline{M}_n^{WE}, \quad n \in \mathcal{N} \quad (9.15)$$

$$\sum_{s \in \mathcal{S}^W} \sum_{\tau=0}^{\overline{M}_n^{NW} - 1} x_{ns(t-7\tau)} \leq 1, \quad n \in \mathcal{N}, t \in \mathcal{T}^{GEN} \quad (9.16)$$

9.3.5 Shift-patterns

$$\sum_{s \in \mathcal{S}} \sum_{\tau=t-\bar{\tau}_p+1}^t S_{s(\tau-t+\bar{\tau}_p)p}^{DES} x_{ns\tau} - \bar{\tau}_p q_{ntp} \geq 0, \quad n \in \mathcal{N}^{DES}, t \in \mathcal{T}, p \in \mathcal{P}_n^{DES} \quad (9.17)$$

$$\sum_{s \in \mathcal{S}} \sum_{\tau=t-\bar{\tau}_p+1}^t S_{s(\tau-t+\bar{\tau}_p)p}^{ILL} x_{ns\tau} \leq \bar{\tau}_p - 1, \quad n \in \mathcal{N}, t \in \mathcal{T}^{GEN}, p \in \mathcal{P}_n^{ILL} \quad (9.18)$$

$$\sum_{s \in \mathcal{S}} \sum_{\tau=t-\bar{\tau}_p+1}^t S_{s(\tau-t+\bar{\tau}_p)p}^{MAN} x_{ns\tau} - \bar{\tau}_p \sum_{s \in \mathcal{S}} S_{s(1)p}^{MAN} x_{ns(t-\bar{\tau}_p+1)} = 0, \quad (9.19)$$

$$n \in \mathcal{N}, t \in \mathcal{T}^{GEN}, p \in \mathcal{P}_n^{MAN}$$

$$\sum_{s \in \mathcal{S}} \sum_{\tau_1=t-\bar{\tau}_{p_1}-\bar{\tau}_{p_2}+1}^{t-\bar{\tau}_{p_2}} S_{s(\tau_1-t+\bar{\tau}_{p_1}+\bar{\tau}_{p_2})p_1}^{AFT} x_{ns\tau_1} - \frac{1}{\bar{\tau}_{p_2}} \sum_{\tau_2=t-\bar{\tau}_{p_2}+1}^t S_{s(\tau_2-t+\bar{\tau}_{p_2})p_2}^{AFT} x_{ns\tau_2} \leq \bar{\tau}_{p_1} - 1, \quad (9.20)$$

$$n \in \mathcal{N}, t \in \mathcal{T}^{GEN}, p_1, p_2 \in \mathcal{P}_n^{AFT} \mid P_{np_1 p_2}^{AFT} = 1$$

9.3.6 Related to fairness

$$\sum_{t \in \mathcal{T}} x_{nst} + V_{ns}^S \geq \underline{M}_{ns}^{TS}, \quad n \in \mathcal{N}, s \in \mathcal{S} \quad (9.21)$$

$$\sum_{t \in \mathcal{T}} x_{nst} + V_{ns}^S \leq \overline{M}_{ns}^{TS}, \quad n \in \mathcal{N}, s \in \mathcal{S} \quad (9.22)$$

$$\sum_{s \in \mathcal{S}^{GEN}} \sum_{t \in \mathcal{T}} (x_{n_1st} - x_{n_2st}) \leq \overline{M}_s^{FT}, \quad n_1, n_2 \in \mathcal{N}^{GEN} \quad (9.23)$$

9.3.7 Variable declarations and fixations

$$x_{nst} \in \{0, 1\}, \quad n \in \mathcal{N}, s \in \mathcal{S}, t \in \mathcal{T}^P \quad (9.24)$$

$$q_{ntp} \in \{0, 1\}, \quad n \in \mathcal{N}, t \in \mathcal{T}, p \in \mathcal{P}^{DES} \quad (9.25)$$

$$y_{tu} \in \{-\underline{D}_{tu}, \dots, 0\}, \text{integer} \quad t \in \mathcal{T}, u \in \mathcal{U}^{GEN} \quad (9.26)$$

$$y_{tu} \in \{0, \dots, \underline{D}_{tu}\}, \text{integer} \quad t \in \mathcal{T}, u \in \mathcal{U}^{GEN} \quad (9.27)$$

$$y_{tu} = 0, \quad t \in \mathcal{T}, u \in \mathcal{U}^{GEN} \quad (9.28)$$

$$x_{nst} = 1, \quad n \in \mathcal{N}, s \in \mathcal{S}, t \in \mathcal{T}^P \mid P_{nst}^{PA} = 1 \quad (9.29)$$

$$\sum_{s \in \mathcal{S}} P_{nst}^R x_{nst} = 1, \quad n \in \mathcal{N}, t \in \mathcal{T} \mid P_n^{DE} = 1 \quad (9.30)$$

$$x_{nst} = 0, \quad n \in \mathcal{N}, s \in \mathcal{S}, t \in \mathcal{T} \mid P_{ns}^S = 0 \quad (9.31)$$

Bibliography

- M. Akbari, M. Zandieh, and B. Dorri. Scheduling part-time and mixed-skilled workers to maximize employee satisfaction. *The International Journal of Advanced Manufacturing Technology*, 64(5):1017–1027, 2012.
- G. H. W. Alliance. Health workforce 2030, 2012. URL http://www.who.int/hrh/documents/strategy_brochure9-20-14.pdf?ua=1.
- M. Azaiez and S. A. Sharif. A 0-1 goal programming model for nurse scheduling. *Computers & Operations Research*, 32(3):491 – 507, 2005.
- J. F. Bard and H. W. Purnomo. Preference scheduling for nurses using column generation. *European Journal of Operational Research*, 164(2): 510 – 534, 2005.
- F. R. Beckmann and K. K. Klyve. Telling a success story of nurse scheduling -model development and analysis, 2016. Specialisation project.
- J. V. d. Bergh, J. Beliën, P. D. Bruecker, E. Demeulemeester, and L. D. Boeck. Personnel scheduling: A literature review. *European Journal of Operational Research*, 226(3):367 – 385, 2013.
- M. J. Bester, I. Nieuwoudt, and J. H. Vuuren. Finding good nurse duty schedules: A case study. *J. of Scheduling*, 10(6):387–405, Dec. 2007.
- E. K. Burke, P. De Causmaecker, G. V. Berghe, and H. Van Landeghem. The state of the art of nurse rostering. *Journal of Scheduling*, 7:441 – 499, 2004.

- M. L. Christiansen, N. C. Petersen, and T. M. Range. Physician rostering: From academia to implementation. Working paper, University of Southern Denmark, 2015.
- Dagsavisen. Ny trend i norsk arbeidsliv, 2016. URL <http://www.dagsavisen.no/innenriks/ny-trend-i-norsk-arbeidsliv-1.775657>.
- A. Ernst, H. Jiang, M. Krishnamoorthy, and D. Sier. Staff scheduling and rostering: A review of applications, methods and models. *European Journal of Operational Research*, 153:3–27, 2004.
- E. W. Hans, M. Van Houdenhoven, and P. J. H. Hulshof. A framework for health care planning and control. Memorandum 1938, Department of Applied Mathematics, University of Twente, Enschede, February 2011.
- Helse-Midt-Norge. St. Olavs Hospital styrets beretning og årsregnskap 2009, 2009. URL <https://ekstranett.helse-midt.no/1001/Foretaksmtter/St.\%201avs\%20Hospital\%20\%C3\%85rsregnskap\%20og\%20styrets\%20\%C3\%A5rsberetning\%202009.pdf>.
- Helse-Midt-Norge. St. Olavs Hospital styrets beretning og årsregnskap 2015, 2015. URL <https://ekstranett.helse-midt.no/1001/Foretaksmtter/Styreberetning\%20og\%20\%C3\%A5rsregnskap\%202015.pdf>.
- D. L. Kellogg and S. Walczak. Nurse scheduling: From academia to implementation or not? *Interfaces*, 37:355 – 369, 2007.
- B. Maenhout and M. Vanhoucke. An integrated nurse staffing and scheduling analysis for longer-term nursing staff allocation problems. *Omega*, 41(2):485 – 499, 2013. Management science and environmental issues.
- NRK. Flere leger ønsker å forlate sykehusene, 2016. URL https://www.nrk.no/nordland/_-flere-leger-onsker-a-forlate-sykehusene-1.13179691.

- T. Range. Interview conducted February 11, 2016 at the Norwegian University of Science and Technology, 2016.
- J. Ruzzakiah, I. Wan Rosmanira, L. C. Yuen, and A. Oughalime. A cyclic nurse scheduling using goal programming. *A Cyclic Nurse Scheduling Using Goal Programming*, 43A:151 – 164, 2011.
- E. Rönning and T. Larsson. Automating the self-scheduling process of nurses in swedish healthcare: a pilot study. *Health care management science*, 13:35 – 53, 2010.
- E. Rönning, T. Larsson, and A. Bertilsson. Automatic scheduling of nurses: What does it take in practice? *Springer Optimization and Its Applications*, 74:151 – 178, 2012.
- S. Simoens, M. Villeneuve, and J. Hurst. Tackling nurse shortages in oecd countries, 2005. URL <http://www.oecd.org/els/health-systems/34571365.pdf>.
- P. M. Vaz and M. Moz. Solving a bi-objective nurse rostering problem by using a utopic pareto genetic heuristic. *Journal of Heuristics*, 14(4): 359–374, 2007.
- J. Vissers, J. Bertrand, and G. D. Vries. A framework for production control in health care organizations. *Production Planning & Control*, 12 (6):591–604, 2001.
- D. M. Warner. Scheduling nursing personnel according to nursing preference: A mathematical programming approach. *Operations Research*, 24 (5):801, 1976.
- W. Zijm. Towards intelligent manufacturing planning and control systems. *OR-Spektrum*, 22(3):313–345, 2000.