

Block Scheduling at Magnetic Resonance Imaging Labs

Anders N. Gullhav^{*1}, Marielle Christiansen¹, Bjørn Nygreen¹, Mats M. Aarlott¹, Jon Erik Medhus¹, Johan Skomsvoll², and Per-Olav Østbyhaug²

¹Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology (NTNU), Trondheim, Norway

²Center for health care improvement in Mid-Norway, St. Olavs hospital, Trondheim, Norway

Abstract

This paper considers a tactical block scheduling problem at a major Norwegian hospital. Here, specific patient groups are reserved time blocks for scanning at a heterogeneous set of Magnetic Resonance Imaging (MRI) labs. The time blocks consist of several time slots, and one or more patients from the same group are scanned in a block. A total weekly number of time slots for each specific patient group is given through demand forecast and negotiations, and several restrictions apply to the allocation of time blocks. Only part of the week is allocated to blocks for the specific patient groups. The rest is classified as open time. Thus, the MRI block scheduling problem consists of finding a cyclic weekly plan where one or more time blocks are to be allocated to each specific patient group, by deciding the day, start time and length, to minimise unfavourable patient group allocations, as well as allocations of open time. For the problem, we propose an integer programming model with an objective function that combines penalties for allocating time blocks to patient groups at unfavourable time slots and labs, and rewards for advantageous positioning of open time slots. The aim of the optimisation model is to facilitate the coordination of the MRI resources between the hospital departments, that are responsible for the specific patient groups, to achieve a fair distribution of time slots to the specific patient groups and open time blocks. The computational study is based on the real problem as well as artificially generated instances. Real-sized instances for our case hospital can be solved in short time. We illustrate how the model can be used to produce Pareto optimal solutions, and how these solutions can provide the decision makers with managerial insight.

Keywords— block scheduling; MRI scheduling; tactical planning; integer programming

1 Introduction

The demand for health care services is increasing worldwide, and the demand for MRI (magnetic resonance imaging) scans is no exception. According to the Organisation for Economic Collaboration and Development (OECD) [19, 20], the number of MRI scans performed in western countries has increased dramatically in the last decade (see Figure 1). The number of MRI scans has grown at a higher rate than the number of MRI units, indicating a better utilisation of the MRI units by time. However, the pressure on good utilisation of such facilities is increasing because the MRI units are often major bottlenecks for the patient flow at hospitals.

Several external factors influence operations management decision making at public hospitals [27]. These external factors include budgetary targets, demographic changes, medical technological advances, and public's increased access to information. Public hospitals need to

^{*}corresponding author (anders.gullhav@iot.ntnu.no)




Figure 1: MRI scans and MRI units per 1,000,000 inhabitants. The numbers are based on data from OECD [19, 20]

adjust to the changes in these factors, and governmental regulations and requirements often drive this change. Governments also impose direct requirements on public hospitals and other public health care providers to increase their efficiency, such as defining guidelines for maximum waiting times for examination and treatment [27]. Diagnostic imaging departments often have a large impact on these patient throughput times because an MRI scan is more often one of several health care services in a patient pathway. Utilising the existing capacity more efficiently is the key to improving patient flow because imaging diagnostic facilities typically represent large investments with tight budgetary limits.

The development in the demand for health care services in general, and MRI diagnostics in particular, is also experienced by St. Olavs hospital, one of Norway’s major hospitals. The MRI unit is part of the Clinic of Radiology and Nuclear Medicine, and functions as a hub for many patients at the hospital. MRI has replaced many of the diagnostic tasks previously performed by clinical staff, which means that more patients flow through this unit than before. In addition, the Norwegian government has tightened requirements for patients’ waiting times, the time between a scan request and the scan (also referred to as access time in the literature). These developments clearly increase pressure on the MRI unit as a scarce resource at the hospital. Thus, the need for utilising the existing resources more efficiently is critical. The utilisation of resources for a health care provider depends on sophisticated processes for planning and scheduling to avoid under-utilisation, staff overtime and long patient waiting times. Hence, improving planning and scheduling processes at the Clinic of Radiology and Nuclear Medicine is believed to increase the efficiency and performance at large parts of the hospital.

To reduce the vast complexity of the entire set of health care services required by all patients and to improve service management, clinical pathways (also called integrated care pathways [6]) have been developed for groups of patients sharing similar symptoms or diagnoses. Essentially, a clinical pathway prescribes which health care services should be performed, in which sequence, and by whom. By introducing clinical pathways, St. Olavs hospital aims to improve the quality of its care, utilise their resources better, avoid mistakes and reduce uncertainty. Some of the larger clinical pathways that include MRI scanning as one of the health care services are classified as patient groups. In addition, other patients are grouped in the schedules due to some specific

requirement, such as cardiology (heart) patients where a radiologist needs to be present during the scan.

Hans et al. [13] propose a holistic framework for planning and control for a health care provider, which consists of four managerial areas, combined with a hierarchical decomposition of decision-making levels. Figure 2 illustrates the combinations of managerial areas and decision-making levels, and presents examples of planning and control functions for each combination. This paper considers the managerial area *resource capacity planning*. Strategic decisions involve dimensioning of resource capacities, while online operational decisions are related to the execution of the health care delivery process. Between them are the tactical and offline operational decisions, which involve the planning and organisation of the execution of health care delivery processes. They are based on more aggregated information and forecasts than the online operational level, and the planning horizon is longer. Specifically, good decision support regarding the block scheduling (also referred to as block planning) and appointment scheduling processes for the MRI unit should provide a basis for better utilisation of the existing MRI resources. In short, block schedules are high-level cyclic plans where time blocks are reserved for specific patient groups at the different MRI labs. These plans must allow for factors such as the technical differences between the MRI labs, equipment requirements, and staff capacities. Moreover, appointment scheduling relates to the patient-to-appointment assignment based on the defined block schedule [16].

	Medical planning	Resource capacity planning	Materials planning	Financial planning	↑ hierarchical decomposition ↓
Strategic	Research, development of medical protocols	Case mix planning, capacity dimensioning, workforce planning	Supply chain and warehouse design	Investment plans, contracting with insurance companies	
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 2: Framework for health care planning and control [13]

Today the block and appointment scheduling for the MRI unit is done manually at St. Olavs hospital. The current and future expected number of MRI scans, as well as strict new governmental regulations for waiting times, makes the assignment of patients to different MRI labs at particular times a highly complex task. Therefore, St. Olavs hospital has realised the need for decision support for these planning issues.

In this paper, we consider the block scheduling problem for MRI labs. The hospital has a given number of MRI labs, and the block scheduling problem is to develop a cyclic weekly schedule where specific patient groups are reserved time blocks at the different MRI labs. Such a weekly schedule is normally repeated for half a year and up to a year. The specific patient groups that have reserved time blocks are typically patient groups belonging to a certain clinical pathway, which have strict deadlines for diagnostics and treatment. For instance, in Norway some of the different clinical pathways for cancer care have deadlines for diagnostics in the range of five to seven days. A time block consists of a number of time slots, and the length of each time block is part of the decision problem. Before the block schedule is constructed, the weekly number of time slots for a specific patient group is fixed and decided through negotiations and

forecasted demand. The weekly demand for examinations is uncertain and varies from week to week, so the specific patient groups with strict deadlines are typically allocated more time slots than the expected demand to hedge against uncertainty. This strategy will lead to time blocks not being fully utilised all the time. Hence, unused time slots are released to other patients at the latest on the day before, so that the released time slots can be utilised. Each specific patient group might be allocated several time blocks of different lengths at different labs in the weekly schedule. Depending on the length of the block, one or more patients from a specific patient group can be assigned to a particular block.

For each specific patient group, one department is responsible for the patients in the group. This department schedules the appointments for their patients. This means that the hospital decentralises the appointment scheduling and detailed planning of each patient's health care services for the specific patient groups. The decentralised scheduling facilitates the operational planning at the different departments, and is particularly beneficial for patients belonging to a clinical pathway where the various health care services need to be coordinated in time. As experienced in the study by Van Sambeek et al. [25], a high percentage of blocks in a weekly schedule increases the waiting time for MRI scans. Therefore, only around 30% of the weekly block schedule consists of blocks belonging to the specific patient groups at St. Olavs hospital at present. The rest of the time slots are denoted open (or unallocated) time slots. The centralised appointment planners have the responsibility for all patients who do not belong to the specific patient groups and can allocate these patients to open time slots. Because of this decentralised appointment scheduling policy, the manual appointment scheduling task for the planners has been more manageable.

In short, the block scheduling problem considered consists of allocating one or more blocks to each specific patient group and deciding the day, start time and length of each block to minimise unfavourable patient group allocations as well as open block allocations, while the total number of time slots for each specific patient group is given. The optimisation model is supposed to facilitate the coordination of the MRI resources between the hospital departments that are responsible for the specific patient groups. The aim is to produce a fair distribution of time slots to the specific patient groups and open time blocks and resolve conflicts in cases where several departments want to be allocated the same time slots at the same MRI labs for their patient groups.

The purpose of this paper is to present an integer programming model for an MRI block scheduling problem with a partially decentralised appointment decision policy. The model is based on a planning problem at St. Olavs hospital, but the model is general and should therefore be useful at other hospitals as well. We introduce multiple objective functions with penalties for unfavourable block lengths and spread of blocks for the specific patient groups and rewards for advantageous positioning of open time slots. The multiple objective functions are combined to form a single objective. We are not aware of any existing mathematical model for MRI block scheduling, so our contribution to the literature is a first integer programming model for this type of problem where the length and start time of each block are variables. In addition, several blocks on the same day and MRI lab can be reserved for the same patient group. Furthermore, the model is intended to be used as an active tool in the development of a new block schedule at St. Olavs hospital, and can also become an important element in a future portfolio of decision support tools at the MRI unit at St. Olavs hospital. The computational study provides some useful managerial insight to a sensible setup of a new block schedule and shows that the model can solve instances of considerably larger size in reasonable time.

The remainder of this paper is organised as follows. Section 2 gives a related literature review. The block scheduling problem is described in Section 3, while Section 4 presents the mathematical formulation of the problem. Extensive computational results are reported and discussed in Section 5, followed by some concluding remarks in Section 6.

2 MRI planning and related literature

As described in the introduction, MRI planning belongs to the managerial area *resource capacity planning* according to Hans et al. [13], and a variety of problems discussed are relevant for MRI planning at different planning levels, ranging from strategic to operational. The Operations Research (OR) literature on MRI planning is particularly scarce with just a few studies reported. Fortunately, there are some similarities between operating room planning and MRI planning. Therefore, we will give a brief literature review and relate the literature within operating room planning to MRI planning. See the excellent surveys on operating room planning and scheduling by Cardoen et al. [8] and Guerriero and Guido [12], as well as the more recent ones focusing on achievements, challenges and pitfalls by Samudra et al. [21] and trade-offs between electives and emergencies by Van Riet and Demeulemeester [24]. In this section, we will mention some problems and literature on each of the strategic, tactical, and operational levels. We limit ourselves to the OR literature with a focus on optimisation-based solution approaches. However, our focus is the tactical planning level as the block scheduling problem belongs to this level.

At the strategic level, the problems range from case mix planning and capacity dimensioning to workforce planning. Case mix planning concerns choosing the best composition and volume of patients in a hospital, see the recent survey by Hof et al. [14]. Relevant issues within MRI case mix planning would be to determine which patient groups to treat and block schedule as well as the number of time slots for each patient group. A related issue concerns decisions regarding which patient groups to handle at the hospital and which ones to outsource to the private market due to limited MRI capacity. Capacity dimensioning is another main strategic problem where both the number of MRI labs and their characteristics must be decided.

Regarding tactical planning, block (or master) scheduling is one of the main planning problems. The MRI block scheduling problem has several similarities to the master schedules used for surgery planning. Several of the contributions which developed master surgery schedules base their approach on (mixed) integer programming models, e.g. [2, 3, 5, 7, 15, 18, 22, 23].

Traditionally, the lengths of the blocks are fixed. However, more recent papers include a variable length of the blocks. However, the model of Beliën et al. [3] does not give the start time of each block. This is taken into account in Mannino et al. [18], where they generate a priori all daily patterns, which are sets of continuous time slots within one day. This is a useful approach for their problem with relatively few time slots each day. They operate with time slots of 2 hours, while a typical time slot for MRI labs is 15-30 minutes. Herein, we therefore present an innovative integer programming model with binary variables giving information about the start time and block length for a block belonging to a (patient group, day, MRI lab) combination.

The uncertainty in surgery duration is a complicating factor that is considered in several mathematical programming models. See for instance [2, 3, 7, 15, 18, 23]. However, the duration of an MRI scan for a given patient has limited uncertainty compared to surgery duration, so this issue is disregarded in our study.

As described in Mannino et al. [18], numerous objectives have been used: balancing patient queue lengths among different patient groups, minimising the changes to existing master schedules, levelling the resulting bed occupancy, maximising operating room utilisation, minimising operating room or staff costs, and maximising profit. None of these objectives are reasonable to use in our case, for various reasons. The number of time slots allocated to each patient group is fixed in advance, so the modelling of resource utilisation or the patient queue lengths is not meaningful. Moreover, the costs of operating the labs can be assumed to be sunk as long as time blocks are allocated during normal working hours. Instead of using the objectives that are typically used in the master surgery scheduling literature, we use a non-monetary objective function based on preferences, which penalises and rewards the positioning and lengths of the blocks. As in some of the papers mentioned above [3, 7, 23], we operate with multiple objective functions.

Guerriero and Guido [12] introduce the term modified block scheduling, and this concept has

some similarities with the planning policy at St. Olavs hospital. In modified block scheduling, only some portion of the schedule is allocated to certain patient groups, and this increases the flexibility of the schedule. For the rest of the time, capacity is scheduled at a later stage when demand is more certain. The open time slots in the MRI block scheduling problem are managed by the centralised planners, who release these time slots gradually for different urgency classes and perform the appointment scheduling. Van Sambeek et al. [25] performed simulation experiments to examine the effects of reducing the number of blocks in MRI planning. The study showed that block reduction lead to a significant decrease in access time.

An integrated strategic and tactical master surgery scheduling approach is described by Fügener [11] where an integer programming model provide both the optimal allocation of how many (strategic) and what (tactical) operation room blocks to assign to each medical specialty. The study includes considerations on the impact on the downstream resources such as intensive care units and general patient wards. Further applications of the model include analysis of the value of flexible resources and the simulation of specific resource expansions. Fügener et al. [10] also consider a master surgery scheduling problem with regards to multiple downstream units. Both exact and heuristics algorithms are proposed to solve the problem.

Another tactical problem concerns decisions related to the number and position of scheduled appointments combined with the appointment scheduling. This is relevant both for a hospital department or primary care having both non-scheduled walk-in patients and scheduled patients with an appointment. The challenge is to match capacity with patient requests and provide as few appointment slots as necessary. Wiesche et al. [28] present a mixed integer programming model (MIP) where the minimum number of appointments scheduled for a weekly profile is determined. The demands related to the number of urgent patients, inter-arrivals and service times are uncertain. Therefore, the optimal appointment schedule from the MIP is evaluated by a comprehensive stochastic simulation model.

A classical appointment scheduling problem with given slots available for appointments is a planning problem at the operational level. Several appointment scheduling studies are reported in the literature. Vermeulen et al. [26] describe an operational adaptive and flexible approach for appointment scheduling for a CT (Computed Tomography) unit consisting of several CT scanners. Simulation experiments showed that their approach of adaptive capacity allocation improved the performance of scheduling patient groups with different characteristics, and the CT capacity was used efficiently. Bhattacharjee and Ray [4] present a case study for appointment scheduling of patients to an MRI lab of the Radiology department at a hospital situated in eastern India. The patient flows in the system are modelled using discrete-event simulation where several appointment scheduling policies are analysed and evaluated.

3 The MRI block scheduling problem

To prepare for the mathematical formulation of the MRI block scheduling problem, we will here give a detailed description of the problem. First, we describe important characteristics regarding the patients requiring MRI scans, the MRI labs, the time aspects and the blocks as well as restrictions and relationships between these. Finally, we give an example to ease the understanding of the problem.

There exists a set of specific *patient groups* requiring MRI scans, such as prostate cancer patients, thorax (breast, lung, heart) patients, arthrography patients, upper gastro patients, etc. Here, a patient group consists of patients sharing certain medical characteristics or belong to the same clinical pathway. Additionally, patients can belong to the same group if they require coordination with specialised resources and personnel (general anaesthesia, radiologists present, a contrast medium injected into the joint area to help highlight structures of the joint for arthrography patients etc.) or that incur high setup costs. In addition, some MRI patients do not belong to any of the specific patient groups.

The hospital has a given number of *MRI labs*, or in short labs, which perform MRI scans. The MRI labs are heterogeneous, in terms of technical properties, the competence of the staff associated with the lab and the geographical locations. Furthermore, the labs might have different opening hours, and can have pre-allocated time slots to research or other activities that make them unavailable for scheduling at given times during a week. This heterogeneity influences the capabilities and preferences for examining the different patient groups at the different labs. For instance, a specific patient group requiring a neuro (head) MRI scan might preferably be allocated to an MRI lab with a high magnetic field strength (3 Tesla). Another technical limitation might be that patient groups requiring a specific coil for setting up the magnetic field could only be allocated to labs equipped with this type of coil. Moreover, some patient groups require specialised staff, such as medical specialists and nurse anaesthetists, who are strategically located at or close to specific labs. For some patient groups, it is practical to limit the number of labs to which the patient group can be allocated, even if the specific patient group can be allocated to a greater number of labs.

The block schedule is a weekly plan that is repeated for a longer time horizon. Each week is divided into a set of *days*, and each day is divided into a homogeneous set of *time slots*. Before the block schedule is constructed, each patient group has been assigned a given number of time slots for scans each week, and this number is typically set higher than the expected weekly demand for the patient group to hedge against uncertain and variable demand. We denote a set of consecutive time slots allocated to the same patient group a *time block*. The length of a time block for a specific patient group is chosen so that several different sets of scans for this group will fit in the block. This means that a time block can serve one or more patients from one patient group. There is a minimum and maximum length of a time block for each patient group. For some patient groups, there is a preference towards allocating time blocks that are long enough to contain multiple scans to ensure a smooth operation of the MRI labs with same personnel and equipment. Although long time blocks might be efficient, some patient groups should be allocated time slots on multiple days during a week due to requirements of the clinical pathway and other resources required. Therefore, there is a minimum number of days in which each specific patient group should be allocated a time block.

In addition to the preferences for scanning specific patient groups in specific labs, the departments responsible for a patient group might have preferences on what times of the week or day it is desirable to be allocated time slots. This could be due to dependencies and coordination with other departments, personnel or other hospital resources. For instance, some patients need an MRI scan before brain surgery, and children often need to be placed under general anaesthesia before the scan starts and should therefore be assigned blocks in the morning.

The appointment scheduling is not directly considered in this research work. However, the block schedule should also allow for subsequent planning of appointments for patients who do not belong to the specific patient groups. To ease the work of the operational planners, one should ensure that there is a minimum level of *open slots* (or unallocated slots) during a day and possibly during the hours in the middle of the day. It is also important that the open time slots are consecutive to a certain degree, so that there is enough time to allocate appointments in the open time slots.

Finally, the block scheduling problem consists of designing a weekly allocation of favourable time blocks to specific patient groups as well as allocating beneficial open time slots with regards to the time of the day, the spread over days, the length of the blocks, and MRI labs. The main decisions for each patient group as well as for the open time are to decide the number of blocks, the start time and the length of each block.

Figure 3 illustrates a possible block schedule for a particular MRI lab. The planning horizon is five days and each day is divided into 16 time slots. Five different patient groups are scheduled for this MRI lab. Patient group 2 has only one block, while patient group 3 has four blocks of equal length. At the end of the day, there is one hour of open time slots that might be used

for patients with high urgency. We see that patient groups 1, 4, and 5 have blocks of different lengths. Patient group 4 has two blocks assigned on the same day. Otherwise, the blocks for a specific patient group are spread over several days (except patient group 2 with just one block). There is also a preference for some open time slots in the middle of the day, and we can find such slots on Tuesday and Thursday. The open time slots are also spread over all days in the week. The department responsible for patient group 3 has typically negotiated the same time blocks four days a week due to other health care services for this patient group before and after the scan. We see that some departments want short time blocks spread over several days, while other departments prefer longer time blocks, such as for patient groups 1 and 5, due to longer scans and the specialised staff and equipment required. Nearly half of the time slots are allocated to the specific patient groups, and this number is higher than the present average rate at the hospital.

	Monday	Tuesday	Wednesday	Thursday	Friday
08:00	Open	Open	Open	Patient group 2	Open
08:30					
09:00					
09:30	Patient group 1	Patient group 4		Open	
10:00					Patient group 5
10:30			Patient group 1		
11:00		Open			
11:30					
12:00				Patient group 5	Open
12:30					
13:00	Open	Patient group 4	Open		
13:30					
14:00	Patient group 3	Patient group 3	Patient group 3		Patient group 3
14:30				Open	
15:00	Open	Open	Open		Open
15:30					

Figure 3: Example of a weekly block schedule with five patient groups on an MRI lab

4 Mathematical formulation of the MRI block scheduling problem

This section presents the mathematical formulation of the block scheduling problem described in the previous section. We start by listing the sets, parameters and variables of the model, before we proceed with explaining the objective function and constraints.

4.1 Sets

The different sets used in the model are given below.

\mathcal{P}	set of patient groups, indexed by p
\mathcal{P}^S	set of patient groups that cannot be allocated time blocks at multiple labs simultaneously
\mathcal{B}_p	set of time block length alternatives for patient group p , indexed by b
\mathcal{B}^O	set of time block length alternatives for allocation of open slots
\mathcal{M}	set of MRI labs, indexed by m
\mathcal{M}_p^P	set of MRI labs available for patient group p
\mathcal{D}	set of days in the planning horizon (one week), indexed by d
\mathcal{D}_p^P	set of days in the planning horizon to which patient group p can be allocated a time block
\mathcal{D}_m^M	set of days in the planning horizon for which MRI lab m is available
\mathcal{D}_{pm}^C	set of days in the planning horizon for which MRI lab m is available for patient group p
\mathcal{T}_{pd}^P	set of available time slots for patient group p at day d
\mathcal{T}_{md}^M	set of available time slots at MRI lab m at day d
\mathcal{T}_{pbmd}^B	set of available time slots at MRI lab m at day d where it is possible to start time block alternative b for patient group p
\mathcal{T}_{bmd}^O	set of available time slots at MRI lab m at day d where it is possible to start open time block alternative b

To ease the exposition of the model, we denote the tuple of an MRI lab, day and time slot (m, d, t) , $m \in \mathcal{M}, d \in \mathcal{D}_m^M, t \in \mathcal{T}_{md}^M$, a *lab-time slot* in the following. In Section 3, we discussed the requirement for having enough consecutive open time slots. This requirement is handled indirectly in the specification of the set \mathcal{B}^O defining the possible lengths of an open time block. For instance, if a total length of consecutive open time slots of less than one hour is forbidden, all lengths of open time blocks defined by \mathcal{B}^O should be at least one hour.

4.2 Parameters

The different parameters used in the model are summarised below.

C_{pm}^{PM}	penalty for allocating patient group p to lab m
C_{pdt}^{PT}	penalty for allocating patient group p to time slot t on day d
R_{mdt}	reward for allocating an open slot to lab-time slot (m, d, t)
α_{pb}	parameter to scale the penalty for time block alternative b for patient group p
α_b^O	parameter to scale the reward for open time block alternative b
T_b	number of time slots in time block alternative b
T_{bdt}^H	number of time slots during the hours in the middle of the day that is allocated if a time block alternative b starts at time slot t on day d
N_p	number of time slots (demand) requested for patient group p
L_p^D	minimum number of days to which patient group p should be allocated a time block
U_p^M	maximum number of distinct labs that can be used for allocation of time blocks for patient group p
F_d	minimum number of open time slots during day d
F_d^H	minimum number of open time slots during the hours in the middle of day d

The length of each time block alternative b is specified by the parameter T_b , and certain lengths are penalised by the scaling parameters α_{pb} for patient group p and α_b^O for open time

blocks.

4.3 Variables

The variables of the model are listed below.

x_{pbmdt}	1, if patient group p is allocated time block alternative b that starts at lab-time slot (m, d, t) ; 0, otherwise
x_{pd}^D	1, if patient group p is allocated at least one time block at any lab on day d ; 0, otherwise
x_{pm}^M	1, if patient group p is allocated at least one time block at lab m on any day; 0, otherwise
y_{bmdt}	1, if lab-time slot (m, d, t) is the first slot for open time block alternative b ; 0, otherwise

Thus, we have two main types of variables: the x_{pbmdt} variables indicating the allocation of time blocks to the specific patient groups, and the y_{bmdt} variables indicating the allocation of open blocks.


4.4 Objective function

The objective function (1) represents the cost of allocating time blocks to the specific patient groups on the MRI labs, and the reward of maintaining open time blocks at favourable time slots. These costs are not monetary values. Instead, they are based on considerations that reflect the preferences of the departments that are responsible for the patient groups and the preferences of the management and planners of the MRI unit.

$$\begin{aligned}
\min z = & \sum_{p \in \mathcal{P}} \sum_{b \in \mathcal{B}_p} \sum_{m \in \mathcal{M}_p^P} \sum_{d \in \mathcal{D}_{pm}^C} \sum_{t \in \mathcal{T}_{pbmd}^B} \alpha_{pb} \left(T_b C_{pm}^{PM} + \sum_{\tau=t}^{t+T_b-1} C_{pd\tau}^{PT} \right) x_{pbmdt} \\
& + \sum_{b \in \mathcal{B}^O} \sum_{m \in \mathcal{M}} \sum_{d \in \mathcal{D}_m^M} \sum_{t \in \mathcal{T}_{bmd}^O} \alpha_b^O \left(\sum_{\tau=t}^{t+T_b-1} R_{md\tau} \right) y_{bmdt}
\end{aligned} \tag{1}$$

The objective function can be regarded as a combination of multiple objective functions. Basically, the term including the x_{pbmdt} variables penalises time allocation to patient groups in undesirable lab-time slots, which is a sum of two terms. The coefficients C_{pm}^{PM} and $C_{pd\tau}^{PT}$ are penalties for allocating patient group p to lab m , and penalties for allocating patient group p to time slot t of day d , respectively. These penalties are given values so that a smaller value is preferable. Moreover, the term including the y_{bmdt} variables rewards the allocation of open time blocks to the desired lab-time slots. We want all the variables in the objective to have positive coefficients. For that reason, we also use smaller positive reward values $R_{md\tau}$ for lab-time slots we prefer to be open. Moreover, the penalty and reward coefficients are assumed to be given values on the same scale, e.g. between 1 and 10. Both the penalty and reward terms contain a factor, α_{pb} and α_b^O , indexed by the time block alternative b , which purpose is to enable the decision maker to scale the penalty or reward for certain time block lengths. For instance, for specific patient groups that cause a significant setup cost before the first scan in a block, it is beneficial to allocate long blocks that allow several similar scans in a row, and thus one would like to give long time block lengths smaller values on the α_{pb} parameter. The α_{pb} parameter is set according to the piecewise linear function in Figure 4a, based on the time block length T_b and parameter \underline{T} . Similarly, longer open time blocks might be preferred over very short open time blocks. However, it might not be beneficial to have very long open time blocks. Thus, based on the time block length T_b and two parameters \underline{T} and \bar{T} , the α_b^O parameter is according to

the piecewise linear function in Figure 4b. However, it is not reasonable to penalise completely open days, so the scaling parameter corresponding to a full day has the same value as the value between the break points. The minimum value of the scaling parameters should be greater than 0, and here we have set this value to 1.



(a) The scaling factor α_{pb} is set based on a piecewise linear function with a breakpoint \underline{T} . (b) The scaling factor α_b^O is set based on a piecewise linear function with breakpoints \underline{T} , \bar{T} , and a full day.

Figure 4: Piecewise linear function for setting the scaling factors based on the time block length T_b

4.5 Constraints

The block scheduling model contains different kinds of constraints which are structured in three groups: allocation and scheduling constraints, demand constraints, and requirements for open time slots. These groups are discussed accordingly, before we conclude the model by stating the variable definitions. In the constraints, we assume that the variables are only defined for the index combinations as specified in the variable definitions (12) - (15), and thus, the variables are not necessarily defined for all time slot indices summed over.

4.5.1 Allocation and scheduling

To obtain a valid schedule without conflicts one must ensure that a lab-time slot (m, d, t) is allocated to at most one patient group. This is specified in equations (2) which state that a time slot is either used by a patient group or set to be open, and which also prevent overlap of two or more time blocks. Moreover, constraints (3) avoid allocation of time slots to different labs simultaneously for patient group $p \in \mathcal{P}^S \subseteq \mathcal{P}$. Some patient groups have a strict limitation on the number of labs they can be allocated to, even if several labs can possibly be used for an MRI scan. The binary variable x_{pm}^M is forced by constraints (4) to take value 1 if patient group p is allocated a time block on lab m . Specifically, if x_{pm}^M equals 0, patient group p cannot be allocated time blocks on lab m at any time. Conversely, if x_{pm}^M equals 1, the constraints for the combination of p and m are redundant. Constraints (5) ensure that the number of labs allocated to patient group p is within the upper limit. Due to the penalties for long open time blocks in the objective function, one must prevent two open time blocks from being adjacent. This is ensured by constraints (6). An analogous set of constraints for the x_{pbmdt} variables is not required, as we do not directly penalise long time blocks for patient groups.

$$\sum_{p \in \mathcal{P}} \sum_{b \in \mathcal{B}_p} \sum_{\tau=t-T_b+1}^t x_{pbmd\tau} + \sum_{b \in \mathcal{B}^O} \sum_{\tau=t-T_b+1}^t y_{bmd\tau} = 1, \quad m \in \mathcal{M}, d \in \mathcal{D}_m^M, t \in \mathcal{T}_{md}^M \quad (2)$$

$$\sum_{b \in \mathcal{B}_p} \sum_{m \in \mathcal{M}_p^P} \sum_{\tau=t-T_b+1}^t x_{pbmd\tau} \leq 1, \quad p \in \mathcal{P}^S, d \in \mathcal{D}_p^P, t \in \mathcal{T}_{pd}^P \quad (3)$$

$$x_{pbmdt} - x_{pm}^M \leq 0, \quad p \in \mathcal{P}, b \in \mathcal{B}_p, m \in \mathcal{M}_p^P, d \in \mathcal{D}_{pm}^C, t \in \mathcal{T}_{pbmd}^B \quad (4)$$

$$\sum_{m \in \mathcal{M}_p^P} x_{pm}^M \leq U_p^M, \quad p \in \mathcal{P} \quad (5)$$

$$y_{bmdt} + \sum_{\beta \in \mathcal{B}^O} y_{\beta md(t+T_b)} \leq 1, \quad b \in \mathcal{B}^O, m \in \mathcal{M}, d \in \mathcal{D}_m^M, t \in \mathcal{T}_{bmd}^O \quad (6)$$

4.5.2 Time block demand

Constraints (7) ensure that each specific patient group is allocated a given number of time slots. Furthermore, some patient groups have a lower bound on the number of allocated days during a week, and we use x_{pd}^D to keep track of the days where patient group p is allocated time blocks. Constraints (8) force the allocation of a positive number of time slots to patient group p if x_{pd}^D is 1. Otherwise, the constraints are redundant. The lower bounds on the number of days are ensured by constraints (9).

$$\sum_{b \in \mathcal{B}_p} \sum_{m \in \mathcal{M}_p^P} \sum_{d \in \mathcal{D}_{pm}^C} \sum_{t \in \mathcal{T}_{pbmd}^B} T_b x_{pbmdt} = N_p, \quad p \in \mathcal{P} \quad (7)$$

$$\sum_{b \in \mathcal{B}_p} \sum_{m \in \mathcal{M}_p^P} \sum_{t \in \mathcal{T}_{pbmd}^B} x_{pbmdt} - x_{pd}^D \geq 0, \quad p \in \mathcal{P}, d \in \mathcal{D}_p^P \quad (8)$$

$$\sum_{d \in \mathcal{D}_p^P} x_{pd}^D \geq L_p^D, \quad p \in \mathcal{P} \quad (9)$$

4.5.3 Allocation of open time slots

A block schedule should also ensure a sufficient number of open time slots during the day, especially during the hours in the middle of the day. Constraints (10) guarantee a minimum number of open time slots during each day, while constraints (11) ensure a minimum number of open slots during the hours in the middle of each day.

$$\sum_{b \in \mathcal{B}^O} \sum_{m \in \mathcal{M}} \sum_{t \in \mathcal{T}_{bmd}^O} T_b y_{bmdt} \geq F_d, \quad d \in \mathcal{D} \quad (10)$$

$$\sum_{b \in \mathcal{B}^O} \sum_{m \in \mathcal{M}} \sum_{t \in \mathcal{T}_{bmd}^O} T_{bdt}^H y_{bmdt} \geq F_d^H, \quad d \in \mathcal{D} \quad (11)$$

4.5.4 Variable definitions

All variables are defined as binary in (12) - (15).

$$x_{pbmdt} \in \{0, 1\}, \quad p \in \mathcal{P}, b \in \mathcal{B}_p, m \in \mathcal{M}_p^P, d \in \mathcal{D}_{pm}^C, t \in \mathcal{T}_{pbmd}^B \quad (12)$$

$$x_{pd}^D \in \{0, 1\}, \quad p \in \mathcal{P}, d \in \mathcal{D}_p^P \quad (13)$$

$$x_{pm}^M \in \{0, 1\}, \quad p \in \mathcal{P}, m \in \mathcal{M}_p^P \quad (14)$$

$$y_{bmdt} \in \{0, 1\}, \quad b \in \mathcal{B}^O, m \in \mathcal{M}, d \in \mathcal{D}_m^M, t \in \mathcal{T}_{bmd}^O \quad (15)$$

4.6 Use of the optimisation model in the planning process

As described in the introduction, the block scheduling process is performed manually at St. Olavs hospital today. In brief, the planners at the MRI unit discuss and negotiate the schedule with the departments that are responsible for the different patient groups, and they try to find a good compromise. However, in such processes, some departments typically end up with better schedules due to their ability to negotiate, and the resulting schedule can be sub-optimal.

The block scheduling model can improve the planning process. Instead of creating a schedule based on informal discussions between the MRI unit and the departments, the model requires the departments to put numbers on their preferences, i.e., the penalty parameters C_{pm}^{PM} and C_{pdt}^{PT} . Similarly, the planners at the MRI unit should put numbers on the rewards for open slots, R_{mdt} . This parameter collection process is certainly an important step in the planning process, and should be formally managed.

Thus, the goal of the optimisation model is to coordinate the shared MRI resources and avoid conflicts among the various departments responsible for the specific patient groups. When constructing a block schedule, it might be desirable to discuss solutions with different emphasis on the objective function terms, that is, investigate the trade-off between allocating favourable lab-time slots to the patient groups versus maintaining open time blocks at appropriate lab-time slots. This can be done by introducing weight parameters in the objective function, and running the model with different weights. Alternatively, one could take a formal approach to construct an approximate Pareto front as we do in Section 5.5. In light of this, one can view the model as having at least two objectives: one that optimises the schedule of the specific patient groups, and one that optimises the schedule of the open time blocks.

5 Computational study

The purpose of the computational study is to illustrate the scalability and functionality of the optimisation model of the MRI block scheduling problem, and discuss how the model can be used to provide managerial insight for the decision makers at the hospital. First, we describe the test instance used in this section, and discuss implementation details.

5.1 Test instances

We have performed experiments using both real data provided by St. Olavs hospital, and artificial data which have a structure that reflects the real-world instances. To be able to test the model on multiple instances of different sizes, we have designed an instance generator with randomness controlled by a seed to construct the artificial data. This generator is constructed to generate test instances that have the structure and characteristics of the real data. Due to confidentiality, all the results presented in this computational study are based on the generated test instances.

Currently, the real instance at St. Olavs hospital consists of 16 specific patient groups in the block schedule, and 6 MRI labs with available time slots for the specific patient groups between 08:00 and 16:00. In total, the block schedule covers about 30% of the available time slots on the 6 MRI labs, and the rest are open time slots. Furthermore, some of the labs are open in the evenings on week-days, and in emergencies it is possible to open labs at night and during weekends. However, no specific patient group is currently allocated time blocks in the evenings or weekends.

In the instance generation, we constructed cases with 16, 32 and 48 patient groups, and 6 and 12 MRI labs, which are available for time blocks during the specified opening hours from 08:00 to 16:00. In these cases, the opening hours are divided into time slots of either 15 minutes or 30 minutes. There can be several types of scans with different durations for each patient group. For each patient group, we randomly draw a typical duration of a scan to be an integer

number of time slots. When using 30-minute slots, the typical duration is drawn to be either 1 or 2 time slots with equal probability, and when using 15-minute slots, the typical duration is drawn to be an integer between 2 and 5 time slots. Moreover, the cases are constructed with different amounts of *total patient group demand*, where approximately 30%, 50% or 70% of the available time slots during the opening hours are to be allocated to specific patient groups. The total patient group demand is distributed among the patient groups, where the demand from each patient group, N_p , is randomly drawn from a uniform distribution with mean equal to the total demand divided by the number of patient groups, but that must be a multiple of the typical duration of a scan. Thus, the sum of the demand from each patient group is close but not necessarily equal to the total patient group demand. Moreover, one fourth of the patient groups can only be allocated time blocks on one lab ($U_p^M = 1$), while the other patient groups do not have this requirement. The minimum number of days a patient group should be allocated time blocks, L_p^D , is randomly drawn from a uniform distribution between 1 and 3. The time block alternatives for the patient groups \mathcal{B}_p are specified with all lengths that are multiples of the typical scan duration up to the minimum of a full day and N_p . The scaling factor α_{pb} is set based on a piecewise linear function (cf. Figure 4a) where the break point \underline{T} is equal to $0.75N_p/L_p^D$. The open time block alternatives \mathcal{B}^O are specified from a minimum length of time slots corresponding to one hour and up to the number of time slots in a day. The scaling factor α_b^O is defined with a valley between break points $\underline{T} = 3$ hours and $\bar{T} = 5$ hours.

The minimum amounts of open time slots during a day F_d (and during the middle of the day F_d^H) are identical for all days, and set equal to the number of time slots on all labs in a day (during the middle of the day) minus the number of time slots corresponding to the total patient group demand divided by 2. Thus, if the number of time slots on all machines during a day is 96 (e.g., 16 slots/day and 6 MRI labs) and the total patient group demand is 30%, F_d is set to $\lceil 96 \cdot (100\% - 30\%) / 2 \rceil = 34$.

The structure of the penalties and rewards C_{pm}^{PM} , C_{pdt}^{PT} , and R_{mdt} was obtained from discussions with the managers of the MRI unit at St. Olavs hospital. In the instance generation, we randomly generate values for the penalties and rewards, in the range [1,10], that have this realistic structure. For instance, for a patient group, the feasible labs are graded on the scale from 1 to 10, which results in coefficients C_{pm}^{PM} . Some labs might also be equally good for a given patient group. Moreover, it is typical that a department responsible for a patient group has some preferred days during the week that fit well for the time blocks. In addition, a department may prefer time blocks between specific hours, e.g., between 10:00 and 12:00, to conform with other activities at the department. This is reflected in the C_{pdt}^{PT} coefficients.

In the following, we distinguish between the notions *case* and *instance*. The number of patient groups, MRI labs and time slots per hour, in addition to the total patient group demand specify a case. Based on the seed, the instance generator produces different data, which are referred to as instances. For each case, we have generated five test instances. The test instances are labelled based on the case and seed, according to the structure mM_pP_qQ.tT_sS, where M denotes the number of labs; P denotes the number of patient groups; Q denotes the total patient group demand; T denotes the number of time slots per hour; and S denotes the seed used in the generation.

5.2 Model implementation

We have implemented the optimisation model of the MRI block scheduling problem in FICO Xpress Mosel 4.0. All experiments are run on a CentOS 6.8 machine with a dual core 3.0 gigahertz Intel E5472 Xeon processor and 16 gigabytes of memory, using the MIP solver of FICO Xpress Optimization Suite 8.0. The MIP solver has utilised up to 8 threads in the tree search. In all experiments, we have given the MIP solver a maximum run time of 3 hours.

To speed up the solution process, we directed the MIP solver to prioritise to branch on the x_{pm}^M variables since these binary variables are important in the solution structure. If a single

x_{pm}^M variable is branched to zero, many x_{pbmdt} variables are forced to zero through constraints (4). Moreover, only a few of the x_{pm}^M variables can be one, cf. constraints (5). Thus, we also directed the MIP solver to select the branch and bound node of the up branch, before the down branch.

Due to the piecewise linear structure of the α_{pb} and α_b^O parameters that has a flat part where the parameters have equal values, the solver might find multiple solutions with equal objective function value and quite similar structure. To speed up the solution process, we suggest to perturb the α_{pb} and α_b^O parameters so that these similar solutions have slightly different objective function values. The perturbation scheme is chosen so that the optimal solution is unchanged. Initial experiments showed that both the branching prioritisation and the perturbation had positive effect on the solution times, and thus, the results presented in this section are obtained using both techniques.

5.3 Model scalability

Here, we present detailed results from the experiments with test instances discussed in Section 5.1. While the instances that resemble the real problem at St. Olavs hospital are the m6_p16_q30_t2 case, it is valuable to test the model on larger instances to assess its scalability. For the instances with time slots of 30 minutes, Table 1 shows the number of variables, constraints, and visited nodes in the branch and bound tree, in addition to the solution time, objective function value and relative optimality gap¹. The results reveal that almost all instances are solved to optimality within three hours, and that most of the smaller instances are solved within 90 seconds. All instances of the size of the real problem are solved within one minute. Two of the largest instances are not solved to optimality, but the optimality gaps are less than one percent, and in practice, these small gaps are of no significance. We also see that the solver requires a certain amount of time, and needs to visit a large number of branch and bound nodes to solve the larger cases. Thus, the complexity of the problem seems to grow exponentially with the problem size, given by the number of variables and constraints. Regarding the optimal objective function value, it remains around the same region for the instances with an equal number of MRI labs (and thus, time slots). For the instances with six labs, the objective function values are roughly between 1300 and 2300, while the objective function value for the instances with 12 labs is roughly in the area from 3200 to 3600.

While 30-minute time slots might be appropriate in tactical decision problems, such as this block scheduling problem, there certainly exist cases where 15-minute slots are preferable and more reasonable. In case the typical duration of a scan for a patient group is 45 minutes or 75 minutes, which is not unusual in an MRI context, one needs a time resolution of 15 minutes to model all possible time block lengths. Table 2 presents the same type of results as Table 1 for the test instances with time slots of 15 minutes. In comparison to the number of variables and constraints for the instances with time slots of 30 minutes, we can see that the numbers of variables and constraints are roughly doubled. Generally, we also see an increase in the solution time, and three instances are not solved to optimality. Still, most of the solution times of the smaller instances are below 10 minutes, and the optimality gaps of the unsolved instances are less than one percent. However, we would like to note that solution times of some hours are rarely an issue for tactical decision problems with a planning horizon of several weeks, such as the problem studied here. Therefore, the computational complexity should not be a barrier if one constructs block schedules with 15-minute time slots. Similarly to the way the objective function values increased with the number of MRI labs in Table 1, we also see an increase in the objective function value when the number of time slots per hour is increased to four. However, this does not imply that a given instance from Table 1, e.g., m6_p16_q50_t2_s1, results in a better solution than the corresponding instance in Table 2, m6_p16_q50_t4_s1. The objective function

¹The relative optimality gap is defined as (Objective function value of best solution - best bound)/best bound.

Table 1: Results from experiments with test cases with time slots of 30 minutes.

Instance	# Variables	# Constraints	# Nodes	Solution time (s)	Objective function value	Gap (%)
m6_p16_q30_t2_s1	29,335	15,084	1,782	76	1,712.6	
m6_p16_q30_t2_s2	23,005	12,931	261	49	1,882.0	
m6_p16_q30_t2_s3	31,019	17,924	46	39	1,691.1	
m6_p16_q30_t2_s4	21,059	12,766	114	21	1,789.3	
m6_p16_q30_t2_s5	22,275	11,705	1	8	1,720.8	
m6_p16_q50_t2_s1	34,239	16,834	31	49	1,671.1	
m6_p16_q50_t2_s2	26,564	14,361	1	35	1,878.9	
m6_p16_q50_t2_s3	35,555	20,418	1	43	1,644.2	
m6_p16_q50_t2_s4	24,406	14,141	49	31	1,864.3	
m6_p16_q50_t2_s5	26,433	12,821	3	23	1,698.6	
m6_p16_q70_t2_s1	35,671	17,143	87	69	1,889.6	
m6_p16_q70_t2_s2	27,903	14,809	11	50	2,288.1	
m6_p16_q70_t2_s3	36,934	21,542	109	70	2,087.6	
m6_p16_q70_t2_s4	25,626	14,326	265	70	2,225.9	
m6_p16_q70_t2_s5	27,880	13,202	1,461	67	1,994.9	
m6_p32_q30_t2_s1	30,268	19,560	90	41	1,647.4	
m6_p32_q30_t2_s2	31,402	17,408	229	31	1,678.0	
m6_p32_q30_t2_s3	31,670	19,804	899	55	1,558.8	
m6_p32_q30_t2_s4	31,252	16,967	1,121	44	1,647.7	
m6_p32_q30_t2_s5	30,107	17,013	71	34	1,632.3	
m6_p32_q50_t2_s1	40,336	24,251	287	57	1,520.7	
m6_p32_q50_t2_s2	40,848	21,079	15	66	1,607.6	
m6_p32_q50_t2_s3	43,022	25,968	661	90	1,360.3	
m6_p32_q50_t2_s4	39,946	20,816	1	31	1,542.8	
m6_p32_q50_t2_s5	39,892	21,085	438	62	1,501.4	
m6_p32_q70_t2_s1	47,723	27,979	1	55	1,427.6	
m6_p32_q70_t2_s2	48,942	24,546	705	97	1,610.4	
m6_p32_q70_t2_s3	51,100	30,115	2,039	148	1,278.0	
m6_p32_q70_t2_s4	47,908	23,879	1,041	108	1,523.3	
m6_p32_q70_t2_s5	47,369	24,187	50	76	1,424.1	
m12_p48_q30_t2_s1	93,127	46,644	3,839	481	3,528.4	
m12_p48_q30_t2_s2	84,539	42,536	24,221	1,243	3,543.6	
m12_p48_q30_t2_s3	91,060	50,575	26,702	1,834	3,589.8	
m12_p48_q30_t2_s4	83,737	43,419	3	111	3,513.6	
m12_p48_q30_t2_s5	82,294	39,794	40,790	2,649	3,545.7	
m12_p48_q50_t2_s1	146,221	69,494	1	267	3,217.2	
m12_p48_q50_t2_s2	130,287	61,450	108,716	10,033	3,231.3	
m12_p48_q50_t2_s3	143,900	77,891	21,997	2,609	3,164.8	
m12_p48_q50_t2_s4	131,177	63,571	189	340	3,283.8	
m12_p48_q50_t2_s5	126,535	57,017	2,661	469	3,182.1	
m12_p48_q70_t2_s1	174,918	80,575	5,983	2,690	3,362.8	
m12_p48_q70_t2_s2	152,940	70,237	2,869	1,018	3,315.4	
m12_p48_q70_t2_s3	170,701	89,618	18,219	10,800	3,153.9	0.35
m12_p48_q70_t2_s4	157,256	74,454	7,649	5,490	3,515.6	
m12_p48_q70_t2_s5	148,585	65,807	13,174	10,800	3,452.3	0.47

value is not directly comparable across instances with different numbers of time slots.

5.4 Evaluating the solution quality

The goal of the optimisation model is to coordinate the MRI resource between the hospital departments that are responsible for the specific patient groups and the planners at the MRI unit, according to their preferences. However, the objective function does not directly measure how satisfied each department responsible for a specific patient group is with their allocated time blocks, and how satisfied the planners at the MRI unit are with their allocated open time blocks. Below, we introduce a way to measure the solution quality from the viewpoint of each department and the planners at the MRI unit.

With z_p^P denoting the total penalties of the time blocks of patient group p , and z_p^P and

Table 2: Results from experiments with test cases with time slots of 15 minutes.

Instance	# Variables	# Constraints	# Nodes	Solution time (s)	Objective function value	Gap (%)
m6_p16_q30_t4_s1	50,569	31,420	276	121	3,531.9	
m6_p16_q30_t4_s2	43,485	27,032	12,203	684	3,826.6	
m6_p16_q30_t4_s3	54,245	36,638	21	103	3,463.1	
m6_p16_q30_t4_s4	42,851	27,816	13,354	561	3,637.3	
m6_p16_q30_t4_s5	41,256	26,214	1,797	159	3,548.3	
m6_p16_q50_t4_s1	57,248	34,132	12	140	3,618.6	
m6_p16_q50_t4_s2	48,756	28,734	1,573	233	4,025.2	
m6_p16_q50_t4_s3	60,906	40,824	9	173	3,534.6	
m6_p16_q50_t4_s4	48,495	29,937	6,551	491	3,920.7	
m6_p16_q50_t4_s5	47,484	28,047	275	102	3,549.9	
m6_p16_q70_t4_s1	58,790	34,518	3	116	4,331.3	
m6_p16_q70_t4_s2	50,811	29,451	63	149	5,615.8	
m6_p16_q70_t4_s3	62,422	42,143	29	208	5,103.8	
m6_p16_q70_t4_s4	50,300	30,071	14	138	4,849.2	
m6_p16_q70_t4_s5	49,187	28,492	23	159	4,347.1	
m6_p32_q30_t4_s1	55,202	37,404	7,745	564	3,298.1	
m6_p32_q30_t4_s2	54,926	35,460	239	113	3,421.2	
m6_p32_q30_t4_s3	54,862	39,786	331	120	3,132.7	
m6_p32_q30_t4_s4	56,870	36,861	241	107	3,315.8	
m6_p32_q30_t4_s5	52,759	34,235	5,569	344	3,341.8	
m6_p32_q50_t4_s1	71,071	44,709	350	221	3,098.6	
m6_p32_q50_t4_s2	70,368	41,828	35	189	3,330.1	
m6_p32_q50_t4_s3	71,435	49,279	109	179	2,875.1	
m6_p32_q50_t4_s4	70,854	43,358	1	93	3,126.9	
m6_p32_q50_t4_s5	67,412	40,697	15	121	3,038.0	
m6_p32_q70_t4_s1	82,179	50,217	65	217	3,015.4	
m6_p32_q70_t4_s2	81,638	47,278	2,243	779	3,452.2	
m6_p32_q70_t4_s3	84,212	56,176	15	233	2,747.2	
m6_p32_q70_t4_s4	83,289	48,610	799	448	3,126.2	
m6_p32_q70_t4_s5	78,999	45,377	29	196	2,944.4	
m12_p48_q30_t4_s1	159,351	87,564	1,818	974	7,069.1	
m12_p48_q30_t4_s2	143,599	80,678	9	583	7,144.3	
m12_p48_q30_t4_s3	156,927	93,978	4,368	1,850	7,192.8	
m12_p48_q30_t4_s4	154,082	88,570	1	386	7,042.2	
m12_p48_q30_t4_s5	138,933	75,695	49,471	10,800	7,161.8	0.21
m12_p48_q50_t4_s1	242,651	121,462	4	906	6,633.6	
m12_p48_q50_t4_s2	216,039	109,286	396	1,527	6,675.3	
m12_p48_q50_t4_s3	237,180	133,775	172	1,722	6,508.8	
m12_p48_q50_t4_s4	230,172	121,962	155	928	6,760.4	
m12_p48_q50_t4_s5	204,909	102,549	135	1,016	6,622.6	
m12_p48_q70_t4_s1	286,092	137,341	2,265	3,663	7,314.8	
m12_p48_q70_t4_s2	248,693	121,892	4,793	10,800	7,175.1	0.53
m12_p48_q70_t4_s3	278,521	151,102	1,143	4,018	6,611.7	
m12_p48_q70_t4_s4	270,408	140,402	436	2,109	7,472.2	
m12_p48_q70_t4_s5	236,137	115,177	7,648	10,800	7,466.8	0.36

\bar{z}_p^P being the best possible (minimum) and worst possible (maximum) penalties, respectively, for patient group p , a *patient group score*, s_p^P , can be computed as in (16). Similarly, we can compute an *open block score*, s^O , based on the total rewards for allocating open time blocks, z^y , which corresponds to the second main term of (1), and \underline{z}^y and \bar{z}^y that refer to the best possible and worst possible reward for allocating open time blocks.

$$s_p^P = \frac{z_p^P - \bar{z}_p^P}{\bar{z}_p^P - \underline{z}_p^P} \cdot 9 + 1, \quad \forall p \in \mathcal{P} \quad \text{and} \quad s^O = \frac{z^y - \underline{z}^y}{\bar{z}^y - \underline{z}^y} \cdot 9 + 1 \quad (16)$$

Each of the patient group scores takes a value between 1 and 10, where a score of 1 means that the given patient group has been allocated time blocks in the best possible way for this group, and a score of 10 means that the given patient group has been scheduled in the worst

possible way. The open block score also takes a value between 1 and 10, and can be interpreted in a similar fashion with regards to the allocation of open time blocks. In the evaluation of a solution one could thus look at the scores of the different patient groups and open blocks, and use that as an indication of the satisfaction, where lower scores indicate higher satisfaction.

In Section 5.5, we will use this score when we are to evaluate different solutions for a test instance that is quite close to the real-world problem at St. Olavs hospital.

5.5 Managerial insight by Pareto optimal solutions

The block scheduling model is supposed to be an active decision support tool available to the managers and planners at the MRI unit at St. Olavs hospital when developing new block plans. In this section, we illustrate how the model can be used in the decision-making process by searching for solutions with different qualities.

As mentioned in Section 4.6, one could have introduced weight parameters in the objective function to obtain solutions with different emphasis on the terms in the objective function. Thus, one could also consider the problem as a multi-objective optimisation problem.

In the process of solving a multi-objective optimisation problem in practice, one can distinguish between two tasks: find a set of Pareto optimal solutions; and choose the most appropriate solution from this set. The first task is done by the optimisation software, while the second is carried out by the decision maker. A way to obtain a set of Pareto optimal solutions is to solve the model with different weights in the objective function. If the weights sum to one, and at least one weight is positive, one will obtain a Pareto optimal solution. Another technique used to obtain a set of Pareto optimal solutions, which is applied here, is to solve an optimisation problem iteratively based on the ϵ -constraint method [17]. The ϵ -constraint method for solving multi-objective optimisation problems is based on minimising one of the objective functions while the other(s) are implemented as constraints upper bounded by some value ϵ . To find a set of Pareto optimal solutions, one solves this single objective problem using different values for ϵ . Specifically, we implement this method by minimising only the first term of the original objective function, that is given in equation (17), subject to the new ϵ -constraint (18) and the original constraints (2)-(15).

$$\min z^x = \sum_{p \in \mathcal{P}} \sum_{b \in \mathcal{B}_p} \sum_{m \in \mathcal{M}_p^P} \sum_{d \in \mathcal{D}_{pm}^C} \sum_{t \in \mathcal{T}_{pbmd}^B} \alpha_{pb} \left(T_b C_{pm}^{PM} + \sum_{\tau=t}^{t+T_b-1} C_{pd\tau}^{PT} \right) x_{pbmdt} \quad (17)$$

$$\sum_{b \in \mathcal{B}^O} \sum_{m \in \mathcal{M}} \sum_{d \in \mathcal{D}_m^M} \sum_{t \in \mathcal{T}_{bmd}^O} \alpha_b^O \left(\sum_{\tau=t}^{t+T_b-1} R_{md\tau} \right) y_{bmdt} \leq \epsilon \quad (18)$$

When finding several Pareto optimal solutions, one can illustrate these solutions by an approximate Pareto front in the space of the objective functions, i.e., the main terms of the original objective function (1). Here, we solve the m6_p16_q30_t2_s1 instance with ϵ taking values in the set of all integers between 1040 and 1665 that are divisible by 5. With ϵ less than 1040, the problem is infeasible, and with ϵ larger than 1665, the ϵ -constraint is non-binding. Since we cannot guarantee that we have found all Pareto optimal solutions when solving the problem with the chosen values of ϵ , we refer to the Pareto front as approximate. Figure 5 shows the approximate Pareto front of the m6_p16_q30_t2_s1 instance, where z^y refers to the left-hand-side of constraint (18), and z^x is the value of the objective function (17).

The optimal objective function value from the experiments presented in Section 5.3 is 1712.6 where the two main terms in the objective function are $z^x = 480.6$ and $z^y = 1232.0$. To illustrate the differences in the solutions for different values of ϵ , Figure 6 shows the block schedule of MRI lab 3 when ϵ equals 1080 and 1500, respectively. This lab is allocated slightly different subsets of five and six of the 16 patient groups in the two different solutions. We have chosen two rather




Figure 5: Approximate Pareto front for the m6_p16_q30_t2_s1 instance.

extreme ϵ values, so that the different characteristics of the solutions are pronounced. The block schedule in Figure 6a ($z^x = 736.7, z^y = 1080.0$) is produced with high emphasis on the rewards for allocation of open blocks, while the block schedule in Figure 6b ($z^x = 351.0, z^y = 1500.0$) is produced with high emphasis on the preferences for allocation of time blocks for patient groups.

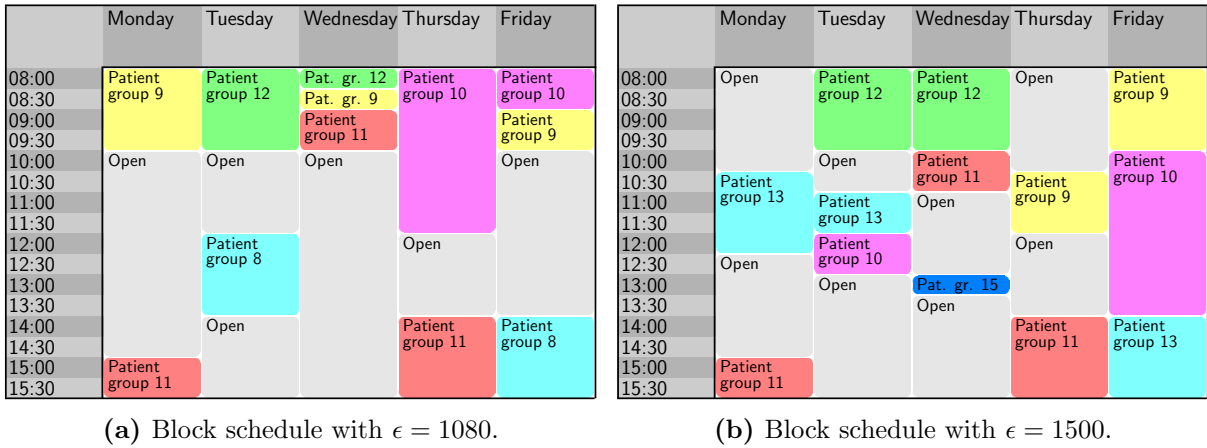


Figure 6: Block schedule of MRI lab 3 for different values on ϵ for the m6_p16_q30_t2_s1 instance.

The differences between the solutions might affect the operation of the MRI labs. In the operational planning, the planners at the MRI unit will schedule patients who do not belong to one of the specified patient groups to the open blocks. Since the scans to be scheduled have different durations, the planners are faced with a combinatorial optimisation problem where one of the objectives is to maximise the utilisation of the open blocks. It is easier to pack long blocks more efficiently. Thus, for the planners, it would be beneficial to have few long open blocks, compared to many short ones. Based on this, the planners would prefer the block schedule in Figure 6a ($\epsilon = 1080$). For the departments that are responsible for the different patient groups, a schedule with short time blocks might lead to a less efficient operation through higher setup costs

due to more coil changes. In the block schedule in Figure 6a, the patient groups 9 and 12 are assigned short time blocks on Wednesdays. Moreover, in this solution, the patient groups tend to be assigned fewer of their preferred time slots, and there are patient groups that are assigned to their less preferred labs. These issues might also lead to worse operational performance at the departments, which implies that they might prefer the block schedule in Figure 6b ($\epsilon = 1500$).

Table 3 provides a more quantitative insight into the differences between the two solutions with $\epsilon = 1080$ and $\epsilon = 1500$, in addition to the solution obtained by minimising the original objective function (1). The rightmost columns give the patient group scores, cf. (16), of the 16 patient groups. In addition, the table presents the average patient group score and the open block score of each solution. One can observe that the open block score and the average patient group score vary between the solutions, which clearly indicates that the solutions are perceived differently by the different planners at the MRI unit and the departments responsible for the specific patient groups. In comparison to the two solutions with $\epsilon = 1080$ and $\epsilon = 1500$, the solution produced with the original objective has a smaller difference between the open block score and average patient group score. When one puts high emphasis on the allocation of open time blocks ($\epsilon = 1080$), one obtains a better open block score than in the solution with the original objective, in sacrifice of higher patient group scores for most of the specific patient groups. Conversely, when one puts little emphasis on the allocation of open time blocks ($\epsilon = 1500$), 8 of the patient groups obtain a score of 1.00, which means that these patient groups are given the best possible schedule, at the expense of allocating worse lab-time slots for the open time blocks. This means that it was possible to assign half of the patient groups to their best lab-time slots while avoiding conflicts. Moreover, the patient group scores are quite evenly distributed in the solution obtained with $\epsilon = 1500$, which is not the case in the solution obtained with $\epsilon = 1080$. In the latter, several patient groups have scores around and below 1.5, at the same time as some patient groups have scores above 4.0. Due to this uneven distribution, some departments might perceive this solution as less fair.

Table 3: Computed patient group scores and open score for three solutions of the m6_p16_q30_t2_s1 instance, produced by either minimising the *original* objective function (1) or minimising the altered objective function (17) in the ϵ -constraint method with $\epsilon = 1080$ or $\epsilon = 1500$.

Solution	Open block score	Avg. PG score	Patient group score															
			1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Orig. obj.	2.11	1.61	1.96	1.52	1.37	2.80	1.46	2.52	2.19	1.35	1.38	1.13	1.17	1.24	1.30	1.38	1.63	1.37
$\epsilon = 1080$	1.28	2.51	4.75	1.52	1.51	2.20	2.29	2.98	4.63	4.81	1.41	1.73	1.00	1.89	4.11	1.89	1.65	1.85
$\epsilon = 1500$	3.56	1.06	1.08	1.00	1.00	1.16	1.18	1.03	1.00	1.00	1.00	1.04	1.00	1.24	1.00	1.00	1.27	1.00

As illustrated, the block scheduling model might provide the decision maker with valuable insight and trade-offs in the process of developing block schedules. However, while the block scheduling model is capable of producing solutions with different characteristics, we want to emphasise that it is up to the decision maker to select a particular solution for implementation.

6 Concluding remarks

The purpose of this paper is to present an optimisation model for an MRI block scheduling problem, where a central motivation for the block allocation is grounded in the idea of decentralising the appointment decision policy. We performed a comprehensive computational study with different numbers of MRI labs and patient groups, and imposed both 15-minute and 30-minute time slots. Additionally, we conducted tests with different levels of demand for the patient groups. Using a general-purpose MIP solver, the problem could be solved to optimality within a reasonable amount of time, even for most of the larger test instances. The case that

resembles today's planning problem at St. Olavs hospital is solved within one minute. While 30-minute time slots might be reasonable for tactical decision problems, there are no computational issues involved in introducing time slots of 15 minutes, and thereby allowing time blocks of finer granularity for patient groups where this is beneficial.

The block scheduling model is intended to be an active decision support tool available to the managers and planners at the MRI unit at St. Olavs hospital when developing new block schedules. The model can be regarded as a multi-objective optimisation model where multiple objectives formed by the penalty and reward coefficients are combined to obtain a single objective function. Since the objective function has great effect on the resulting block schedule, careful specification of the coefficients is required. Moreover, one can add weights to the different terms of the objective function to obtain solutions with characteristics that are preferred by the decision maker. However, this is difficult in practice, and we suggest using the model to find a set of Pareto optimal solutions by iterative use of the ϵ -constraint method, and thereafter, letting the decision maker select the most suitable solution from this set. We show the usefulness of such an approach, by illustrating and quantifying the differences between the solutions.

The work conducted in this paper relies on that the subset of patient groups to explicitly plan for in the block schedule, and the total amount of time allocated to each such patient group have been decided in advance. We see a potential in developing optimization-based decision support models for these decisions. Because of the uncertain demand, different decisions on the number of time slots allocated to the patient groups might have different operational consequences, such as patients' waiting times or resource utilisation. Therefore, these decisions could be evaluated, for instance in a simulation-optimisation framework, where patient group demand is simulated. See Fu [9] for an introduction to the topic simulation-optimisation, and Amaran et al. [1] for a recent review of algorithms and applications. Furthermore, the unused time slots for a patient group are released at least a day in advance, and used for patients with high urgency. However, if there are too few patients to allocate to the released time slots, the utilisation of the labs is reduced. The use of time blocks to urgent cases might also affect the requirements for an optimal allocation of time blocks. Another direction for future research is to develop decision support tools for the centralised appointment scheduling, and possibly connect this tool with the block scheduling model.

Acknowledgements

We would like to thank the two anonymous reviewers for their valuable comments and suggestions that helped us speed up the optimisation model and improve the writing of this paper.

References

- [1] Satyajith Amaran, Nikolaos V. Sahinidis, Bikram Sharda, and Scott J. Bury. Simulation optimization: a review of algorithms and applications. *Annals of Operations Research*, 240(1):351–380, 2016.
- [2] Jeroen Beliën and Erik Demeulemeester. Building cyclic master surgery schedules with leveled resulting bed occupancy. *European Journal of Operational Research*, 176(2):1185 – 1204, 2007.
- [3] Jeroen Beliën, Erik Demeulemeester, and Brecht Cardoen. A decision support system for cyclic master surgery scheduling with multiple objectives. *Journal of scheduling*, 12(2): 147–161, 2009.
- [4] Papiya Bhattacharjee and Pradip K. Ray. Simulation modelling and analysis of appointment

- system performance for multiple classes of patients in a hospital: A case study. *Operations Research for Health Care*, 8:71 – 84, 2016.
- [5] John T. Blake, Franklin Dexter, and Joan Donald. Operating room managers’ use of integer programming for assigning block time to surgical groups: a case study. *Anesthesia & Analgesia*, 94(1):143–148, 2002.
- [6] Harry Campbell, Rona Hotchkiss, Nicola Bradshaw, and Mary Porteous. Integrated care pathways. *BMJ: British Medical Journal*, 316(7125):133, 1998.
- [7] Paola Capanera, Filippo Visintin, and Carlo Banditori. A goal-programming approach to the master surgical scheduling problem. In *Health Care Systems Engineering for Scientists and Practitioners*, pages 155–166. Springer, 2016.
- [8] Brecht Cardoen, Erik Demeulemeester, and Jeroen Beliën. Operating room planning and scheduling: A literature review. *European Journal of Operational Research*, 201(3):921–932, 2010.
- [9] Michael C. Fu. Optimization for simulation: Theory vs. practice. *INFORMS Journal on Computing*, 14(3):192–215, 2002.
- [10] Andreas Fügener, Erwin W Hans, Rainer Kolisch, Nikky Kortbeek, and Peter T Vanberkel. Master surgery scheduling with consideration of multiple downstream units. *European journal of operational research*, 239(1):227–236, 2014.
- [11] Andreas Fügener. An integrated strategic and tactical master surgery scheduling approach with stochastic resource demand. *Journal of Business Logistics*, 36(4):374–387, 2015.
- [12] Francesca Guerriero and Rosita Guido. Operational research in the management of the operating theatre: a survey. *Health care management science*, 14(1):89–114, 2011.
- [13] Erwin W. Hans, Mark Van Houdenhoven, and Peter J.H. Hulshof. A framework for health-care planning and control. In *Handbook of healthcare system scheduling*, pages 303–320. Springer, 2012.
- [14] Sebastian Hof, Andreas Fügener, Jan Schoenfelder, and Jens O. Brunner. Case mix planning in hospitals: a review and future agenda. *Health Care Management Science*, 20(2):207–220, 2017.
- [15] Matias Holte and Carlo Mannino. The implementor/adversary algorithm for the cyclic and robust scheduling problem in health-care. *European Journal of Operational Research*, 226(3):551–559, 2013.
- [16] Peter J.H. Hulshof, Nikky Kortbeek, Richard J. Boucherie, Erwin W. Hans, and Piet J.M. Bakker. Taxonomic classification of planning decisions in health care: a structured review of the state of the art in OR/MS. *Health systems*, 1(2):129–175, 2012.
- [17] Antonio L. Jaimes, Saúl Z. Martinez, and Carlos A.C. Coello. An introduction to multiobjective optimization techniques. *Optimization in Polymer Processing*, pages 29–57, 2009.
- [18] Carlo Mannino, Eivind J. Nilssen, and Tomas E. Nordlander. A pattern based, robust approach to cyclic master surgery scheduling. *Journal of Scheduling*, 15(5):553–563, 2012.
- [19] OECD. Magnetic resonance imaging (MRI) exams (indicator), December 2015. doi: 10.1787/1d89353f-en Accessed on 08 December 2015.
- [20] OECD. Magnetic resonance imaging (MRI) units (indicator), December 2015. doi: 10.1787/1a72e7d1-en Accessed on 08 December 2015.

- [21] Michael Samudra, Carla Van Riet, Erik Demeulemeester, Brecht Cardoen, Nancy Vansteenkiste, and Frank E Rademakers. Scheduling operating rooms: achievements, challenges and pitfalls. *Journal of Scheduling*, 19(5):493–525, 2016.
- [22] Pablo Santibáñez, Mehmet Begen, and Derek Atkins. Surgical block scheduling in a system of hospitals: an application to resource and wait list management in a british columbia health authority. *Health care management science*, 10(3):269–282, 2007.
- [23] Jeroen M. van Oostrum, Mark Van Houdenhoven, Johann L. Hurink, Erwin W. Hans, Gerhard Wullink, and Geert Kazemier. A master surgical scheduling approach for cyclic scheduling in operating room departments. *OR Spectrum*, 30(2):355–374, 2008.
- [24] Carla Van Riet and Erik Demeulemeester. Trade-offs in operating room planning for electives and emergencies: a review. *Operations Research for Health Care*, 7:52–69, 2015.
- [25] Jasper R.C. Van Sambeek, Paul E. Joustra, Saskia F. Das, Piet J. Bakker, and Mario Maas. Reducing MRI access times by tackling the appointment-scheduling strategy. *BMJ quality & safety*, 20(12):1075–1080, 2011.
- [26] Ivan B. Vermeulen, Sander M. Bohte, Sylvia G. Elkhuisen, Han Lameris, Piet J.M. Bakker, and Han La Poutré. Adaptive resource allocation for efficient patient scheduling. *Artificial intelligence in medicine*, 46(1):67–80, 2009.
- [27] Jan Vissers and Roger Beech. *Health operations management: patient flow logistics in health care*. Psychology Press, 2005.
- [28] Lara Wiesche, Matthias Schacht, and Brigitte Werners. Strategies for interday appointment scheduling in primary care. *Health care management science*, pages 1–16, 2016. doi: 10.1007/s10729-016-9361-7.