Stochastic Scheduling Techniques for Integrated Optimization of Catheterization Laboratories and Wards

Irving I. van Heuven van Staereling · René Bekker · Cornelis P. Allaart

Abstract In catheterization laboratories (cath labs), doctors are required to perform invasive cardiovascu-lar procedures under a variety of specific constraints. Patients undergoing a treatment in a cath lab, generally also require preparative and aftercare at one of the hospital's wards, which complicate the scheduling pro-cess significantly. Still, in practice, scheduling of procedures for cath labs is mainly done by hand, which partly can be explained by the fact that many models make simplistic assumptions that ignore the actual practical complexity of the problem, such as the inherent randomness.

In this paper, we propose an Integer Linear Programming based technique that integrates optimization for both cath labs and wards, while incorporating randomness within the scheduling process. Since the natural objective function is non-linear, the key insight for applying this method is that the objective function can be linearized under specific assumptions. These models have been tested on a case study of the VU Medical Center, for which the results are shown to be effective, as useful blueprints for the daily schedules are generated according to the preference of the hospital.

Keywords Cath lab · Integer Linear Programming · Multi-objective Optimization · Stochastic Timetabling

1 Introduction

A catheterization laboratory (or cath lab or cardiac catheterization room) is a special type of hospital room in which doctors (mainly cardiologists) perform invasive cardiovascular procedures to diagnose, visualize and treat cardiovascular diseases. For a good utilization of these cath labs, advanced scheduling is crucial [14]. However, scheduling of cath interventions is complicated due to many factors, including variability in the number of patients, urgency of patients, and the duration of such interventions, see e.g. [10, 13, 15, 16, 18]. Next to the cath lab, the availability of a hospital bed is crucial for treating a patient, making scheduling even more involved.

Due to (semi-)urgent patients and the uncertainty in the procedure lengths, planners have complications to estimate the number of patients that can be treated at a specific day. When too few patients are scheduled, the available time of the staff and rooms are not used efficiently, which decreases the profit and increases the length of the waiting list. On the other hand, when too many patients are scheduled, staff is forced to work longer than desired or patients need to be rescheduled on another day with adverse effect for patient experiences.

The choice of patients that need to be scheduled strongly depends on the procedure type, but also depends on the availabilities of surgeons and the degree to which urgent surgeries need to be taken into

I.I. van Heuven van Staereling
Centrum Wiskunde & Informatica, Science Park 123, 1098 XG Amsterdam, The Netherlands
Vrije Universiteit Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands
E-mail: heuven@cwi.nl

R. Bekker
Vrije Universiteit Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands
E-mail: r.bekker@vu.nl

C.P. Allaart
VU University Medical Center, De Boelelaan 1117, 1081 HV Amsterdam, The Netherlands
E-mail: cp.allaart@vumc.nl

account. But more importantly, all patient types require their own preparation and aftercare at wards that have a limited capacity, meaning that the schedules at the cath labs are limited by the available capacity at the wards. Typically, the capacity constraints at the wards are more stringent towards the end of the week (Thursdays and Fridays) than during the weekend and early in the week, see e.g. [2]. This can be explained by the admission pattern, as elective procedures are not scheduled during weekends leading to peaks in bed demand towards the end of the week. Thus, scheduling of patients/procedures at cath labs requires an integral approach involving both cath labs and the beds at the wards.

The aim of this paper is to provide a mathematical model for integrated scheduling of cath lab procedures at the labs and wards taking the inherent variability into account. Our contribution is that we fully incorporate the randomness in procedure and hospitalization time in the scheduling model. We do this by formulating an Integer Linear Program (ILP) where the required excess capacity is approximated by a piecewise linear function. Moreover, the case study gives insight in the practical aspects of cath lab scheduling.

From an operations research perspective, the optimization of the schedules of cath labs may be considered to be a specific type of Operating Room (OR) scheduling. The differences between cath labs and Operating Theatre's (OT's) are primarily of a practical and organizational nature, having an impact on the mathematical formulation as well. The first difference is simply the size; while many large hospitals manage at least 10 to 20 OR's, the number of cath labs is typically only a few (for instance, at the VUmc there are 3). The organization of cath labs and OT's is also considerably different, affecting the constraints of a model. The OT is used by many medical disciplines involving different departments, whereas the cath lab is dedicated to cardiology. As an intermediate step, OT's typically first divide their sessions over the different disciplines. For cath lab scheduling, this intermediate step is not necessary but more flexibility is often possible, for instance, in the duration or starting time of a session. Furthermore, cath labs generally need to perform procedures of which the duration is very hard to estimate; after all, for a relatively large part of the procedures, the required actions only become known during that procedure. Also, there are (semi-)urgent patients and restrictions on which patients can be treated in which cath lab due to availability of equipment per lab. Hence, cath labs contain many complicating scheduling features. Due to the single department involved and the relatively smaller size, we envisage that cath labs are an ideal example of OR scheduling.

Related work Scheduling of operating rooms is a well studied area in the operations research literature, see e.g. [6, 12, 17] for some literature reviews in health care. We refer the reader to [5, 7] for literature reviews specifically related to operating room planning and scheduling. For scheduling in cath labs in particular, we were not able to find any references; cath labs are also not mentioned at all in the above mentioned reviews.

Papers considering the planning of OR's taking the bed occupancy of the wards into account as well are [1, 3, 4, 8, 11, 19]. These papers are focused on surgery scheduling (or scheduling of OT time blocks or elective admissions) with the objective that the resulting bed occupancy is balanced, whereas the scheduled procedures fit in the available OR time. For the occupation of the OR and the wards, only the mean is considered, such that the scheduling problem can essentially be formulated as an ILP model. In the current paper, we include the expected overtime and the expected excess at the wards in our ILP. The randomness in bed demand as a result of the surgical schedule is fully analyzed in [20]; however, this paper does not contain an optimization algorithm and improvement is accomplished by trial and error.

We note that even when the cath labs are disregarded (and only the excess bed demand at the wards are optimized) the problem is known to be NP-hard, meaning that creative techniques are required to optimize the problem. Yet, to make the mathematical models relevant for practice, it is essential to incorporate realistic assumptions. The current paper aims to make an important step in that direction.

Overview This paper is structured as follows. In Section 2, we illustrate the motivation and research objectives of this paper. Subsequently, we formally define the problem in Section 3 by introducing the required notation in Section 3.1. For the reader's convenience, we first present a model without variability, to illustrate the basic ideas of the model in a simple setting. Then, randomness of the procedures times and bed occupancy (at the cath labs and wards, respectively) is incorporated in the model. In this section, we also include optional, realistic extensions that illustrate the flexibility of the presented model. We proceed by a case study in Section 4, where results are presented that illustrate the practical use of the models of this paper. Finally, we complete the paper by a few concluding remarks in Section 5.

2 Background

The research presented in this paper is motivated by the scheduling difficulties from several years ago at the Department of Cardiology of the VU University medical center (VUmc). Currently, VUmc has three cath labs at their disposal where different types of cardiological procedures can be performed (heart catheterizations, implant placements, etc.). Prior to such procedures, patients require a specific preparation (which generally takes hours to days) at a suitable ward, where usually also the required aftercare is provided for the patients. Due to specialized equipments and personnel, not every patient can be hospitalized at every ward or be treated at every cath lab. The subsets of suitable wards and cath labs per patient depend not only on the patient's procedure, but also on the urgency of the patient, since specific wards are primarily intended for urgent or elective patients. The available wards and cath labs can be found in Tables 1 and 2.

Ward	# Beds	Urgencies	Procedures
5B	14	Elective, semi-urgent, urgent	CAG/PCI's, Implants
5C	4	Elective, semi-urgent	EFO/ablations, Implants
CCU	6	Urgent, semi-urgent	CAG/PCI's, Implants
EHH	6	Urgent, semi-urgent	CAG/PCI's, Implants
SCAR	4	Elective, semi-urgent	CAG/PCI's

Table 1 Wards cardiology VU medical center

Room	Starting time	Ending time	Procedures
1	08:15	16:30	EFO/ablations, Implants
2	09:30	16:30	CAG/PCI's, Implants
3	08:15	16:30	CAG/PCI's

Table 2 Cardiac Catheterization Rooms cardiology VU medical center

So far, patient scheduling has been done entirely manually, which has regularly led to last-minute or ad hoc improvisations in order to accommodate incoming patients properly. Even though no severe accidents have resulted from manual scheduling, there has been a strong conjecture at VUmc that improvement could be obtained with patient scheduling, particularly because the scarce number of beds at the five wards are poorly taken into account during the scheduling process. Ideally, the utilization at each ward is as stable as possible to balance the workloads, but arguably more important, to make sure that there is room available for (semi-)urgent patients as often as possible. For this reason, the VU medical center approached the VU University to perform a data analysis on their cardiological patient flow and to construct quantitative methods that can improve their utilization, of which the results are presented in this paper.

2.1 Data analysis

To admit patients at the cath lab, there is also a bed required at a ward, requiring synchronization during the scheduling process. In this section, we give an impression of patient flow at the wards due to patients at the cath lab. It is important to plan patients at wards such that the load is stabilized, thereby increasing the buffer for accepting (semi-urgent or urgent) patients and to balance the workloads of nurses at the wards.

To plan this appropriately, the scheduler should have knowledge regarding the distribution of the hospitalization times of patients. To provide an intuition for this duration, the data analysis of the hospitalization times is given in Figure 1.

This multi-modal pattern is due to two reasons. First of all, it is very unusual that patients are transferred or sent home from the hospital during the evening at night. Patients are generally received during the morning, meaning patients are unlikely to have a hospitalization time between 16 or 20 hours. Also, most of the patients visit the hospital for a CAG/PCI-procedure, for which the total hospitalization time lasts between 4 and 10 hours. Moreover, there are groups of patients that remain slightly over 1 (24h) or 2 (48h) days.

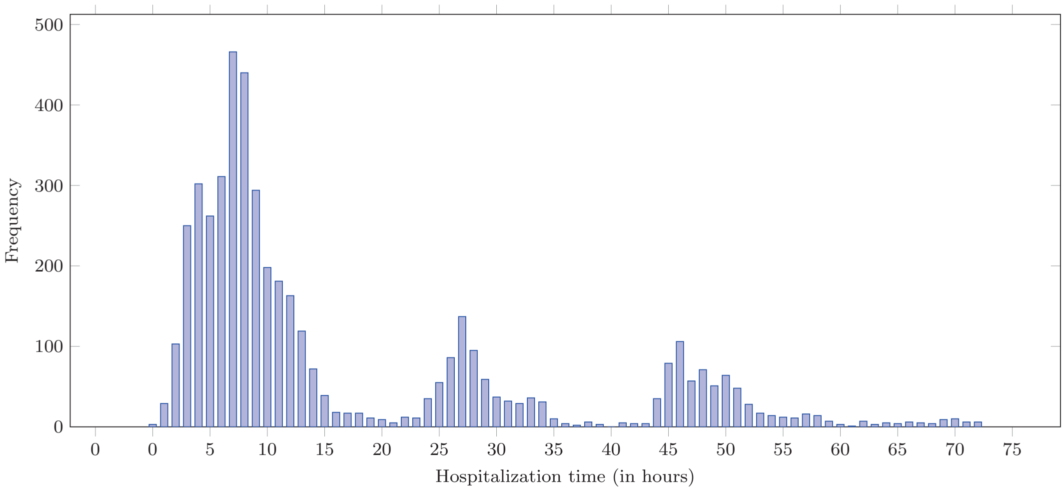

Fig. 1 Distribution of hospitalization times

Daily patterns The mentioned hospitalization times clearly affect the load on the wards throughout the day. To give an impression of the workload, the daily patterns are visualized in Figure 2. Specifically, Figure 2 displays the average number of occupied beds over all five wards.

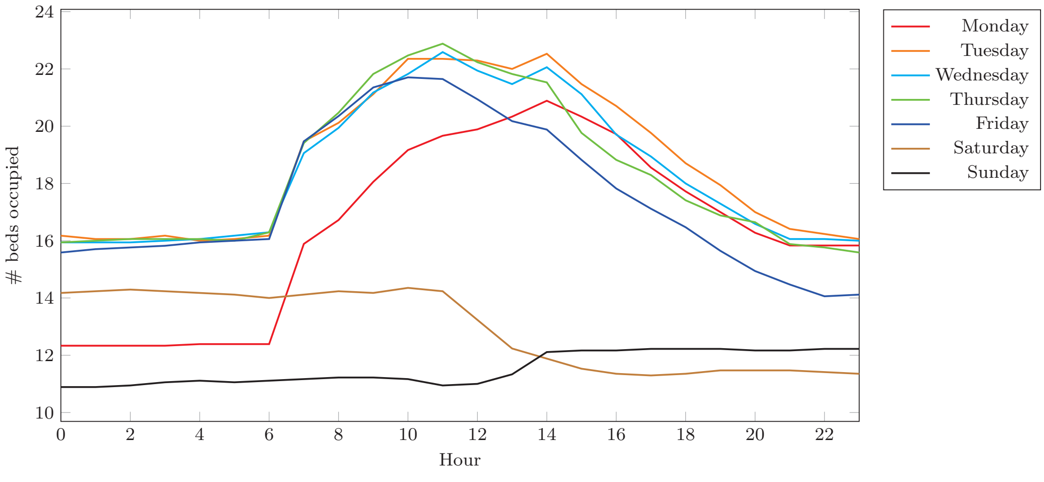

Fig. 2 Average bed occupancy of all wards per day

It is clear that the number of occupied beds increases as soon as the working day starts (mainly due to the Special Care, that is the only ward that is not open during the night and weekend) and decreases as the end of a working day approaches. For similar reasons, fewer patients are within the hospital during the weekend. However, on Sundays, the occupancy increases since some patients need to be hospitalized the day in advance.

Weekly patterns To obtain a vivid comparison between the weekdays, the weekly pattern is given in Figure 3. For this graph, we considered three different points of time: 08:00, 12:00 and 16:00. It should be no surprise that 12:00 is a peak moment within the day (apart from Saturday and Sunday). Even though it is intended that all patients are received before 08:00, it frequently occurs that patients arrive later.

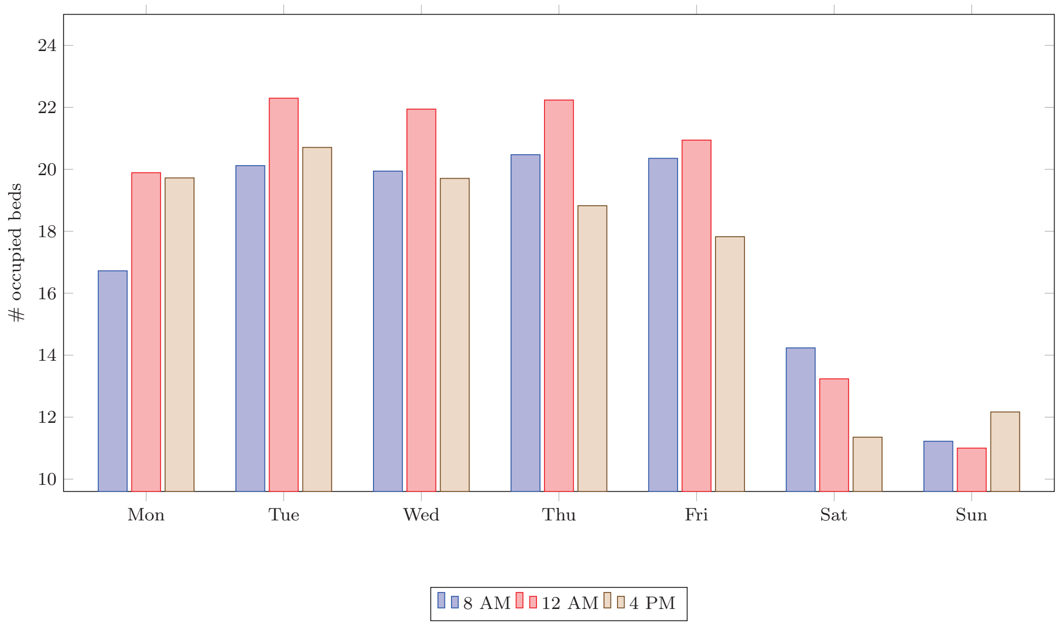

Fig. 3 Average bed occupancy all wards per week

Variability Figures 2 and 3 show the variability in bed occupancy during the day and week. Next to these patterns, there is also a considerable amount of variability (or random fluctuations) around these patterns. To accommodate for this variability excess capacity is required. In particular, when the mean duration of all admitted patients fit in the available capacity, overtime and excess of the bed capacity may still occur. With advanced scheduling, the aim is to minimize the frequency of such occurrences.

2.2 Research objectives

From a high-level point of view, the scheduling process should lead to a blueprint specifying which type of procedures (or patient category) should be carried out during which day (and time slot). Every procedure category has the following characteristics:

- An expected demand (e.g., per week or per year).
- A preparation time at a ward before the procedure.
- An expected procedure time on a cath lab and corresponding standard deviation.
- An expected aftercare time at a ward and corresponding standard deviation.

These procedures need to be scheduled such that the cath lab and ward capacities are used efficiently while keeping overtime and excessive ward occupancy at an acceptable level. This is formulated in terms of an objective function consisting of multiple performance indicators (PI's):

- The expected number of performed procedures (e.g., per week).
- The expected overtime of the personnel at the cath labs.
- The expected overload on the beds at the wards.
- The extent to which (semi-)urgent patients can be treated on time.
- The expected waiting time of a patient for a procedure. For elective patients, this can be expressed in the number of days that need to be waited at home, while this waiting time can be expressed in minutes for urgent patients, since the waiting time for this group naturally should be in the order of minutes.

The objective of this research is to design practical optimization methods that can produce a schedule within a reasonable amount of time, that is able to take the above parameters, constraints and objectives into account. Clearly, these PI's are not equally important.

3 Optimization models

In this section, we introduce our ILP to generate a blueprint for scheduling cath lab procedures. We start by introducing a considerable amount of notation in Subsection 3.1. To improve readability, we formulate an ILP for a model where variability is mainly neglected in Subsection 3.2; that is, we only consider the mean cath lab and ward occupations throughout the week and do not take the random fluctuations around these means into account. We observe that the mean occupancy at the ward requires the full distribution of the length of stay as in Figure 1. In Subsection 3.3 we take the random fluctuations into account and consider the expected overtime. A linear approximation that can be used in the ILP is elaborated upon in Subsection 3.4. Subsection 3.5 provides several possible model extensions.

3.1 Parameters and variables

For the models presented in this paper, we proceed by defining a considerable number of parameters and variables that illustrate the complexity of the process. Even though the number of parameters and variables is relatively large, we will argue that this model represents the practical setting very accurately. We adopt the following notational conventions: we define $[X]$ as the set $\{1, \dots, X\}$ for any natural number $X \in \mathbb{N}$. Also, for a variable x_{ij}, we use the shorthand notation $x = (x_{ij})_{i,j}$.

Define R as the number of rooms for procedures, and W as the number of wards that can provide the required preparation and aftercare for patients. We assume that every patient in the model always undergoes a procedure and is assigned to a ward. The schedule will be created for a fixed time horizon, comprising D days; in practice, it is natural to set this value to 7 to create a weekly schedule. Every patient of the C categories (of procedures) has its own probability distribution with respect to his procedure time, as well as his after care time. Therefore, we introduce the following random variables:

$$\begin{aligned} H_c &= \text{required hospitalization time (in days) for category } c \text{ at a ward} \\ P_c &= \text{required procedure time (in hours) for category } c \text{ at a cath lab} \\ Z_{cd} &= \#\text{ (semi-)urgent patients of category } c \text{ on day } d \end{aligned}$$

We assume that the mean and variance of the above random variables are known. For our case study, these are based on an extensive data analysis, which has been partly provided in Section 2. As mentioned before, the creation of a schedule is simply the assignment of procedures to time slots, and for this reason, the following decision variable for our stochastic program are introduced:

$$x_{cdrw} = \#\text{ patients of category } c \text{ scheduled on day } d \text{ at room } r \text{ and ward } w$$

where $c \in [C]$, $d \in [D]$, $r \in [R]$ and $w \in [W]$. The extensive decision variable, having 4 subscripts, gives the impression that the model will become too large to solve, even using the current state-of-the-art ILP solvers. The case study will show that such a model is practically implementable. But first, we introduce the following variables to derive an initial integer program:

$$\begin{aligned} \tau_{dr} &= \text{available time for procedures on day } d \text{ at room } r \text{ (in hours)} \\ \upsilon_{dr} &= \text{expected overtime on day } d \text{ at room } r \text{ (in hours)} \\ \beta_{dw} &= \text{available number of beds on day } d \text{ at ward } w \\ \lambda_{dw} &= \text{expected overload on day } d \text{ at ward } w \text{ (in beds)} \end{aligned}$$

Depending on the practical situation, τ_{dr} and β_{dw} can either be decision variables or given constants. The former would be the case if the tactical decision would involve the opening times of the cath labs in addition to the blue print. Moreover, υ_{dr} and λ_{dw} are used to obtain the expected overtime and expected ward overload as performance indicators.

Due to the expensive (and therefore limited) equipment at the rooms, patients of specific categories cannot be treated at every room, and even not on every day, due to, e.g., the availabilities of surgeons. Furthermore, not every patient can be hospitalized at every ward, primarily due to the specialized skills of the nurses at the corresponding ward. In other words, for the rooms, wards and also days, there is only a set of categories of patients that can be treated. To account for these so-called "compatibilities", we introduce the following sets, which are assumed to be given:

$$\begin{aligned}\delta_d &= \text{set of patient categories that can be treated on day } d\\ \rho_r &= \text{set of patient categories that can be treated at room } r\\ \omega_w &= \text{set of patient categories that can be hospitalized at ward } w\end{aligned}$$

Note that $\delta_d, \rho_r, \omega_w \subseteq [C]$ for $d \in [D]$, $r \in [R]$ and $w \in [W]$. The following constants are also expected to be part of the input, and is closely related to the performance indicators. In particular, these constants represent the relative weights for each of the five different performance measures. Such a weighted combination of performance measures is a standard mathematical approach for comparing the quality of different schedules for multiple criteria. Hence, we require the following five constants for our model:

$$\begin{aligned}V_{\tau,r} &= \text{loss per opened hour of room } r \text{ (including staffing costs)}\\ V_{\upsilon,r} &= \text{loss per hour overtime at room } r\\ V_{\beta,w} &= \text{loss per reserved bed at ward } w \text{ (including staffing costs)}\\ V_{\lambda,w} &= \text{loss per overloaded bed at ward } w\\ R_c &= \text{profit of treating a patient of category } c\end{aligned}$$

Finally, we need to pose some bounds on the number of elective procedures, opening times of the cath labs and number of beds available. As mentioned, it depends on the practical situation whether these three aspects are given or part of the optimization. Let us briefly comment on these three aspects:

- If the number of elective procedures is a decision variable, this implies that there exist a strategic decision about the case mix. Moreover, due to arrangments with insurance companies, there might be either upper or lower bounds on the number of elective procedures that should be carried out. In any case, there will be an upper bound due to the finite population requiring specific treatments.
- For the number of beds at a ward, there is typically an upper bound due to space limitations or limited staff availability. To avoid small wards, there may exist a lower bound as well. Often the avaible number of beds is given, although closing some beds during the week (or weekend) might be a good tactical decision.
- For the opening time per room there may be a lower bound due to working restrictions. Similarly, there will be an upper bound on the working time. For OR's, the opening times are typically given.

The abovementioned constraints can easily be incorporated by introducing the following constants:

$$\begin{aligned}E_{min,c} &= \text{minimum number of elective patients to plan of category } c\\ E_{max,c} &= \text{maximum number of elective patients to plan of category } c\\ T_{min,r} &= \text{minimum opening time of room } r\\ T_{max,r} &= \text{maximum opening time of room } r\\ B_{min,w} &= \text{minimum number of beds to reserve at ward } w\\ B_{max,w} &= \text{maximum number of beds to reserve at ward } w\end{aligned}$$

We note that by taking the upper bound equal to the lower bound, the corresponding quantity is actually given and no longer a decision variable.

3.2 Basic ILP formulation

In the first model formulation, we neglect the random fluctuations of the cath lab and ward occupancies and only consider their (time-dependent) means. For cath lab r at day d, the mean occupancy can be calculated by $\sum_{c=1}^{C} \sum_{w:c\in\omega_w} \mathbb{E}(P_c) \cdot x_{cdrw}$. For the expected ward occupancy, we need to take admissions of (all) previous days into account as well. For instance, if we consider day d, then a patient admitted at day $d-i$ us still present at day d with probability $\mathbb{P}(H_c \geq i)$. Summing over all previous days, the mean occupancy of ward w at day d is $\sum_{c=1}^{C} \sum_{r:c\in\rho_r} \sum_{i=0}^{\infty} \mathbb{P}(H_c \geq i) \cdot x_{c(d-i)rw}$. We assume that the blueprint is periodic, such that $x_{cdrw} = x_{c(d-kD)rw}$ for any $k \in \mathbb{N}$.

Now, we state the ILP where the objective function and constraints are described below.

$$\min \Pi(x, \tau, \upsilon, \beta, \lambda) \tag{1}$$

$$\text{s.t.} \sum_{c=1}^{C} \sum_{w:c\in\omega_w} \mathbb{E}(P_c) \cdot x_{cdrw} \leq \tau_{dr} + \upsilon_{dr} \qquad d \in [D]; r \in [R] \tag{2}$$

$$\sum_{c=1}^{C} \sum_{r:c\in\rho_r} \sum_{i=0}^{\infty} \mathbb{P}(H_c \geq i) \cdot x_{c(d-i)rw} \leq \beta_{dw} + \lambda_{dw} \qquad d \in [D]; w \in [W] \tag{3}$$

$$\upsilon_{dr} \geq 0 \qquad r \in [R]; d \in [D] \tag{4}$$

$$\lambda_{dw} \geq 0 \qquad d \in [D]; w \in [W] \tag{5}$$

$$B_{min,w} \leq \beta_{dw} \leq B_{max,w} \qquad d \in [D]; w \in [W] \tag{6}$$

$$E_{min,c} \leq \sum_{w:c\in\omega_w} \sum_{d:c\in\delta_d} \sum_{r:c\in\rho_r} x_{cdrw} \leq E_{max,c} \qquad c \in [C] \tag{7}$$

$$T_{min,r} \leq \tau_{dr} \leq T_{max,r} \qquad d \in [D]; r \in [R] \tag{8}$$

$$x_{cdrw} \in \mathbb{N}_0 \tag{9}$$

Objective function The objective function $\Pi(x, \tau, \upsilon, \beta, \lambda)$, depending on the matrices of five-tuple (possible) decision variables consists of five terms:

$$\Pi(x, \tau, \upsilon, \beta, \lambda) = \pi_\tau(\tau) + \pi_\upsilon(\upsilon) + \pi_\beta(\beta) + \pi_\lambda(\lambda) - \pi_P(x)$$

which are defined as follows:

– $\pi_P(x)$ is the expected profit obtained from all elective procedures that are performed during the scheduling horizon. Since R_c is defined as the profit obtained for treating a patient of category c, this simply implies that:

$$\pi_P(x) = \sum_{c=1}^{C} \sum_{d:c\in\delta_d} \sum_{r:c\in\rho_r} \sum_{w:c\in\omega_w} R_c \cdot x_{cdrw}$$

– $\pi_\tau(\tau)$ is the expected loss from staffing/operational costs due to the opening time of the rooms, i.e.:

$$\pi_\tau(\tau) = \sum_{d=1}^{D} \sum_{r=1}^{R} V_{\tau,r} \cdot \tau_{dr}$$

– $\pi_\upsilon(\upsilon)$ is the expected loss from overtime at rooms, i.e.:

$$\pi_\upsilon(\upsilon) = \sum_{d=1}^{D} \sum_{r=1}^{R} V_{\upsilon,r} \cdot \upsilon_{dr}$$

– $\pi_\beta(\beta)$ is the expected loss from staffing/operational costs due to the available beds at the wards, i.e.:

$$\pi_\beta(\beta) = \sum_{d=1}^{D} \sum_{w=1}^{W} V_{\beta,w} \cdot \beta_{dw}$$

– $\pi_\lambda(\lambda)$ is the expected loss from overloaded wards, i.e.:

$$\pi_\lambda(\lambda) = \sum_{d=1}^{D} \sum_{w=1}^{W} V_{\lambda,w} \cdot \lambda_{dw}$$

Observe that the objective function is clearly linear in all its decision variables. The decision on the relative values of the weights R_c and vector V is typically a managerial decision.

Constraints

- Constraint (2) implies that the sum of the scheduled procedures on every room per day, must be smaller than the opening time plus the overtime of that room. Since the overtime of the room, υ_{dr}, is a decision variable that needs to be minimized, this constraint is a soft constraint which always can be fulfilled. In this constraint, (semi-)urgent patients are not yet taken into account, which clearly can affect the optimal opening time of the rooms. A simple way to account for (semi-) urgent patients is to replace x_{cdrw} by $\left(x_{cdrw} + \frac{\mathbb{E}(Z_{cd})}{|\{r:c\in\rho_r\}|}\right)$. Here, $\mathbb{E}(Z_{cd})$ is the expected number of (semi-)urgent patients of category c on day d, and $|\{r : c \in \rho_r\}|$ is the number of rooms at which these patients can be treated. The additonal term represents the number of (semi-)urgent patients that on average need to be treated per day, assuming that these patients are spread evenly among the rooms.
- Constraint (3) implies that the expected occupation on every ward per day, must be smaller than the number of reserved beds (the capacity) plus the excess demand (also referred to as overload). Similar to (2), this constraint is a soft constraint that always can be fulfilled. (Semi-)urgent patient can also be accounted for in the context of wards, by replacing $x_{c(d-i)rw}$ by $\left(x_{c(d-i)rw} + \frac{\mathbb{E}(Z_{cd})}{|w:c\in\omega_w|}\right)$. Also, the model can account for the preparational care for patients, by adding the term:

$$\sum_{i=1}^{\infty} \mathbb{P}(\text{“\# days of preparation care of category } c\text{”} \geq i) \cdot x_{c(d+i)rw}$$

at the left-hand side of the inequality-sign (which is not added for simplicity).
- Constraints (4) and (5) imply that the expected overtime per room and expected overload per ward has to be at least 0, such that underload is not rewarded.
- Constraints (6), (7), and (8) state that the available number of beds, the number of admitted patients over the time horizon, and the opening hours of the rooms, respectively, should be between their upper and lower bound.
- Constraint (9) is the standard requirement that the number of elective procedures per day cannot be negative and must be integer.

3.3 Overtime and excess demand

In this section we consider the overtime and excess demand (or overload) if we take the randomness in procedure times and occupancy into account. Specifically, let μ_{dr} and σ^2_{dr} denote the mean and variance of the total procedure time (sum of all procedures) at day d for room r. These first two moments depend on the decision variables x_{cdrw} and are directly given by

$$\mu_{dr} = \sum_{c=1}^{C}\sum_{w=1}^{W} \mathbb{E}(P_c) \cdot x_{cdrw}, \quad \sigma^2_{dr} = \sum_{c=1}^{C}\sum_{w=1}^{W} Var(P_c) \cdot x_{cdrw} \qquad (10)$$

for $d \in [D]$ and $r \in [R]$. We assume here that the **total** procedure time on a day follows a normal distribution. This assumption is motivated by the Central Limit Theorem stating that the scaled sum of indepent and identically distributed random variables converges to a normal distribution. Although we typically see that the duration of a single procedure is not normally distributed, the sum of a couple of procedures often is.

The expected overtime at any room and day can now be determined using the overshoot over a fixed level of the normal distribution. Let T denote the opening time and drop the indices from the notation of υ_{dr}, μ_{dr}, and σ_{dr} for now for notational convenience. Then,

$$\upsilon = \int_T^{\infty} (t-T) \cdot \frac{1}{\sqrt{2\pi\sigma^2}} \cdot e^{-\frac{1}{2}\left(\frac{t-\mu}{\sigma}\right)^2} dt$$

A similar expression can be obtained for the expected overload at a ward, λ_{dw}. Observe that the occupied beds at day d consists of patients that arrived at some day $d-i$ for $i \geq 0$. A patient arriving at day $d-i$ is still present at day d with probability $\mathbb{P}(H_c \geq i)$. That is, for each patient, there is a Bernoulli random variable indicating whether the patient is present at day d or not. The total occupancy may thus be represented by a sum of Bernoulli distributions (partly with different parameters), naturally leading to a normal approximation for the ward occupancy as well. To obtain the first two moments of the ward occupancy, we note that the variance of a Bernoulli random variable with parameter p equals $p(1-p)$. Thus, the mean (μ_{dw}) and variance (σ^2_{dw}) at day d for ward w are

$$\mu_{dw} = \sum_{c=1}^{C} \sum_{r:c\in\rho_r} \sum_{i=0}^{\infty} \mathbb{P}(H_c \geq i) \cdot x_{c(d-i)rw},$$
$$\sigma^2_{dw} = \sum_{c=1}^{C} \sum_{r:c\in\rho_r} \sum_{i=0}^{\infty} (1 - \mathbb{P}(H_c \geq i)) \cdot \mathbb{P}(H_c \geq i) \cdot x_{c(d-i)rw}$$

A key issue for the ILP is the non-linear expression for the expected overtime υ; the expressions for μ_{dr} and σ^2_{dr} are linear in the decision variables, but these terms also appear in the exponent within the integral. To overcome this issue, we will use a piecewise linear approximation of the expected overtime, as explained in the next section.

3.4 Linear approximation of overtime

In the ILP model, the expected overtime υ is part of the minimization, whereas it contains non-linear expressions in terms of the decision variables. For the moment, let us assume that μ, σ and T are given, also neglecting the subscripts again. Using standard calculus, the expected overtime can be rewritten as

$$\begin{aligned}\upsilon(\mu, \sigma^2, T) &= \int_T^{\infty} (t - T) \cdot \frac{1}{\sqrt{2\pi\sigma^2}} \cdot e^{-\frac{1}{2}\left(\frac{t-\mu}{\sigma}\right)^2} dt \\ &= \int_T^{\infty} t \cdot \frac{1}{\sqrt{2\pi\sigma^2}} \cdot e^{-\frac{1}{2}\left(\frac{t-\mu}{\sigma}\right)^2} dt - T \cdot \int_T^{\infty} \frac{1}{\sqrt{2\pi\sigma^2}} \cdot e^{-\frac{1}{2}\left(\frac{t-\mu}{\sigma}\right)^2} dt.\end{aligned}$$

Using a change of variable for the first term, we have

$$\begin{aligned}\upsilon(\mu, \sigma^2, T) &= \int_{T-\mu}^{\infty} (t + \mu) \cdot \frac{1}{\sqrt{2\pi\sigma^2}} \cdot e^{-\frac{t^2}{2\sigma^2}} dt - T \cdot \left(1 - \Phi\left(\frac{T-\mu}{\sigma}\right)\right) \\ &= \int_{T-\mu}^{\infty} t \cdot \frac{1}{\sqrt{2\pi\sigma^2}} \cdot e^{-\frac{t^2}{2\sigma^2}} dt + \mu \int_{T-\mu}^{\infty} \frac{1}{\sqrt{2\pi\sigma^2}} \cdot e^{-\frac{t^2}{2\sigma^2}} dt - T\left(1 - \Phi\left(\frac{T-\mu}{\sigma}\right)\right) \\ &= \left[-\frac{\sigma^2}{\sqrt{2\pi\sigma^2}} \cdot e^{-\frac{t^2}{2\sigma^2}}\right]_{t=T-\mu}^{\infty} + \mu \cdot \left(1 - \Phi\left(\frac{T-\mu}{\sigma}\right)\right) - T \cdot \left(1 - \Phi\left(\frac{T-\mu}{\sigma}\right)\right) \\ &= \frac{\sigma}{\sqrt{2\pi}} \cdot e^{-\frac{(T-\mu)^2}{2\sigma^2}} + (\mu - T) \cdot \left(1 - \Phi\left(\frac{T-\mu}{\sigma}\right)\right)\end{aligned}$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution. The first expression can also be numerically determined using a Riemann sum, which might be practically preferable.

For the linearization, we make one very important assumption, namely that the ratio between the variance and the expectation of a procedure team of every category (on a room) is approximately the same. In other words, the variance of a procedure time increases almost linear with the expected procedure time, which is a realistic assumption since longer procedure times usually tend to have more variation. More formally, this means that for all rooms $r = 1, \ldots, R$, it must be that $Var(P_{c_1})/\mathbb{E}(P_{c_1}) \approx Var(P_{c_2})/\mathbb{E}(P_{c_2})$ for every $c_1, c_2 \in \rho_r$. Therefore, we drop the assumption that the opening time is a decision variable, i.e., we need to replace τ_{dr} by a given constant T.

Consider a room r on day d and denote μ_{dr} as the expected procedure time and σ^2_{dr} as its corresponding variance. Due to the earlier mentioned reason, we can express σ^2_{dr} as a linear function of μ_{dr}, i.e., $\sigma^2_{dr} = Q_r \cdot \mu_{dr}$, where Q_r is the (weighted) average ratio $Var(P_c)/\mathbb{E}(P_c)$ for every category $c \in \rho_r$. Under this assumption, υ_{dr} is a function that only depends on μ_{dr}, which is linear in the decision variables.

To incorporate $\upsilon_{dr}(\mu_{dr})$ in our ILP, we adopt a piecewise linear approximation using tangent lines of the function $\upsilon_{dr}(\mu_{dr})$ (which is now only a function of μ_{dr}). Let $N_{dr} \in \mathbb{N}$ be the number of tangent lines that touch $\upsilon_{dr}(\mu_{dr})$. Then, every tangent line $n \in [N_{dr}]$ must be of the form:

$$y_n(\mu_{dr}) = \alpha_{drn} \cdot \mu_{dr} + \beta_{drn}$$

where the constants α_{drn} and β_{drn} for day d, room r and the nth tangent line should be chosen such that

$$\upsilon_{dr}(\mu_{dr}) = y_n(\mu_{dr}), \quad \text{and} \quad \upsilon'_{dr}(\mu_{dr}) = \alpha_{drn}$$

The use of these tangent lines is crucial for the model to become linear again. For the ILP, the constants α_{drn} and β_{drn} should be determined first. Then Equation (2) should be replaced by

$$\upsilon_{dr} \geq \alpha_{drn}\mu_{dr} + \beta_{drn} \qquad n \in [N_{dr}]; d \in [D]; r \in [R]$$

where μ_{dr} is given by Equation (10). Clearly, if N_{dr} increases, the accuracy increases, but also the computation time. The same can be done for overloads at wards, where μ_{dr} and σ_{dr} should be replaced by μ_{dw} and σ_{dw} given above.

3.5 Extensions of the model

In this section, we consider some extensions of the ILP that are motivated by practical considerations observed at the VUmc.

Non-linear performance indicators So far, we have put a linear penalty on e.g., the overtime for every room. This suggests that the first hour of overtime is experienced just as inconveniently as the second hour, but this is not always the case in practice. For instance, some overtime may be considered to be inevitable and be part of the job. Excessive overtime may be perceived as poor scheduling and impacts satisfaction of staff members. We illustrate how to overcome this limitation by introducing the following constants.

$$\begin{aligned}
V_{\upsilon 1,r} &= \text{loss of the first hour overtime at room } r \\
V_{\upsilon 2,r} &= \text{loss per hour overtime after the first hour at room } r \\
V_{\omega 1,w} &= \text{loss of the first overloaded bed at ward } w \\
V_{\omega 2,w} &= \text{loss per overloaded bed after the first overloaded bed at ward } w
\end{aligned}$$

We have made a distinction between the loss for the first hour overtime, and the loss per overtime after the first hour. The first hour of overtime is less inconvenient that the following hours, i.e., $V_{\upsilon 1,r} < V_{\upsilon 2,r}$ for all $r \in [R]$. For overload at wards a similar effect occurs, albeit for a slightly different reason. The first overloaded bed may perhaps be solved with some shuffling among the wards, but when there are two or more beds required, this is difficult ro arrange physically, which may result in e.g., patients waiting before they can enter the ward.

Closing rooms We mentioned earlier that it was possible that there is a minimum opening time per room. For example, staff needs to work at least 4 hours per day, as it may be undesirable that surgeons or nurses work for 3 hours or less. However, depending on the number of prodedures over the entire scheduling horizon, it may be beneficial to close a room for the entire day. In terms of our mathematical model, we desire from our decision variable for the opening time, τ_{dr}, to be in element of $\{0\} \cup [T_{min,r}, T_{max,r}]$ instead of only $[T_{min,r}, T_{max,r}]$. Such a sitaution may be implemented by adding the following constraints:

$$\begin{aligned}
\tau_{dr} &\geq T_{min,r} \cdot y_{dr} & d \in [D]; r \in [R] \qquad (11) \\
\tau_{dr} &\leq y_{dr} M & d \in [D]; r \in [R] \qquad (12)
\end{aligned}$$

where $y_{dr} \in \{0, 1\}$ is a decision variable and M is sufficiently large.

4 Case study: VU medical center

As mentioned in Section 2, the research in this paper is motivated by VUmc, which therefore acts as a case study for the mathematical models presented in Section 3. At VUmc. the procedures are subdivided in more categories than the three categories mentioned in Table 1 and 2, because the results would be too inprecise. After all, within one of the three categories, there are multiple variants with a different procedure time and hospitalization time. For example, there are different type of ablation procedures, of which some is well-known that these last one and a half to two times as long as standard ablations.

4.1 Input data

When creating a schedule of patients, one naturally can only work with information that is known prior to the procedure. The inconvenience of the acquired dataset for this research is that only the information after the procedure is known; for example, a standard PCI procedure may be planned, but during the procedure, this may turn out to be a PCO CTO procedure, which last two to three times as long. Hence, we cannot derive by our data was the original diagnosis was. For this reason, we will work only with categories for which we know that the patient belongs to this category. We have chosen to only work with categories for which enough procedures/measurements in the data could be found, i.e., at least 40 procedures per year.

This has led to the following input data regarding the procedures; the procedure times are in **hours**, while the hospitalization times are expressed in **days**.

Category	#	μ_v	σ_v	μ_p	σ_p	μ_n	σ_n
CAG/PCI	18	0	0	1,33	0,5	0,2	0,1
CAG/PCI 5B	1	0	0	1,5	0,6	0,8	0,5
SwanGanz5C	1	1	0,5	1,25	0,6	0,5	0,6
Implant	5	*	*	2	0,75	0,7	0,7
Short Ablation	2	*	*	2,5	0,9	0,75	0,4
Long Ablation	3	*	*	3,5	0,8	0,5	0,6
Semi-urgent PCI 5B	5	2	2,3	1,25	0,5	1,1	1,1
Semi-urgent PCI CCU	5	0,7	1,2	1,2	0,3	0,8	1,8
Semi-urgent PCI EHH	1	0,5	0,4	1,05	0,5	0,2	0,2
Semi-urgent PCI SCAR	5	0,1	0,1	1,2	0,5	0,2	0,1
Semi-urgent implant 5B	1	0,5	0,4	1,5	0,6	2,6	4,4
Urgent PCI CCU	4	0,1	0,2	1,0	0,4	0,6	1,1
Urgent PCI EHH	2	0,1	0,2	1,0	0,5	0,2	0,2

Table 3 Average number of procedures per week

For simplicity, the numbers in Table 3 have been rounded/corrected where required. Whenever a star is mentioned at the preparation time, the hospitalization time depends on the time of the day where the procedure takes place. If a patient for an implant is scheduled in the morning on a cath lab, he or she will be hospitalized the day prior to the procedure at a ward. When the patient is scheduled in the afternoon, there is enough preparation time for the patient to be hospitalized in the morning of the same day.

As a minor remark, prior to every procedure, the room needs to be prepared and the preceding patient needs to depart from the room, which in total takes approximately 15 minutes. For this reason in our implementation, approximately 15 minutes have been added to every procedure.

Choice of the weights The data presented in Table 3 does not suffice to complete the desired input for the model in Section 3

$$V_{\tau,r} = 3$$
$$V_{\upsilon 1,r} = 5$$
$$V_{\upsilon 2,r} = 9$$
$$V_{\beta,w} = 0.25$$
$$V_{\omega 1,w} = 3$$
$$V_{\omega 2,w} = 9$$
$$R_c = 5 \text{ for dotters, } 10 \text{ for implants, } 20 \text{ for ablations}$$

4.2 Results

To illustrate the results, it is not informative to mention the objective value, since this is a meaningless number for the reader. Instead, we illustrate the quality of the solution of our case study by depicting the plan that is derived on an average day. In other words, the model outputs a list of instruction in an average week. This provides fundamental insights to the scheduler that can also be used in irregular weeks.

As a clarification regarding the blueprint of room 1: the small, red thin blocks represent the average duration per day that the room is busy with a (semi-)urgent patient for an implant. Depending on the urgency, it is sensible to treat these at Monday or Wednesday because these days contain more free space. The blueprint for room 2 only consists of two PCI's per day en requires no explanation.

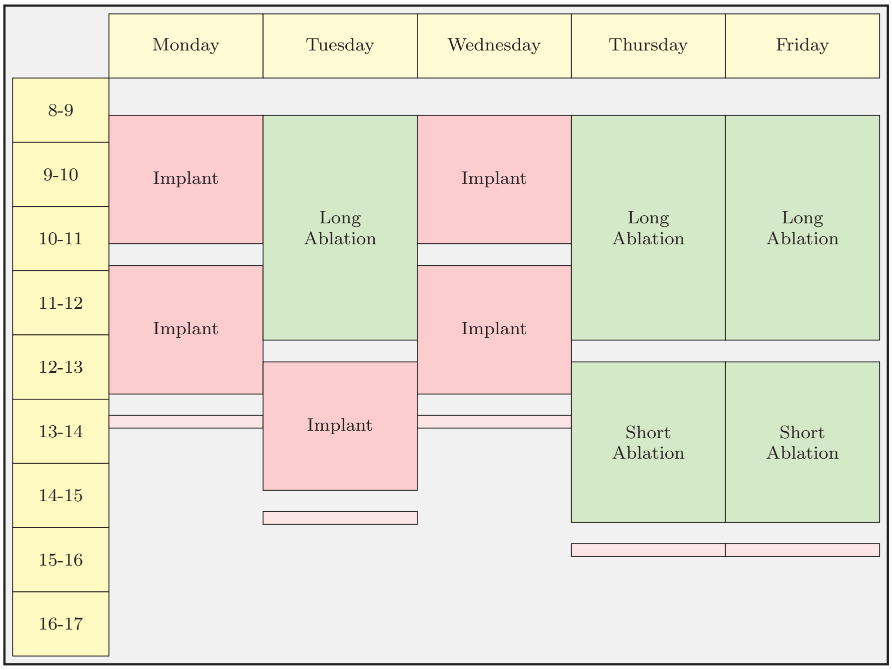

Fig. 4 Blueprint room 1

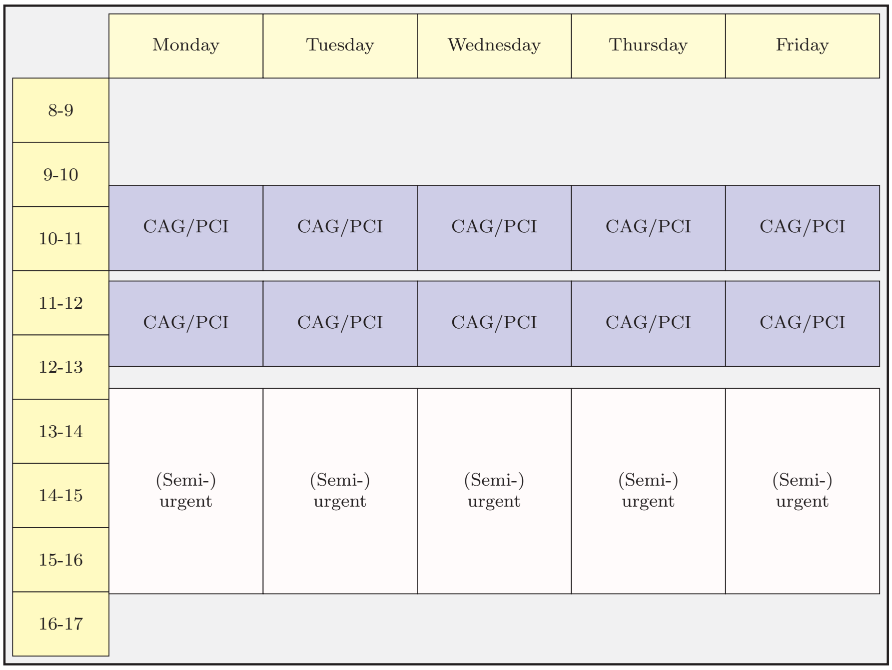

Fig. 5 Blueprint room 2

Regarding the blueprint of room 3, it must be noticed that the Swan Ganz procedure and Dotter5B (a PCI with a much longer aftercare) on respectively Monday and Friday are scheduled. After all, a Swan Ganz

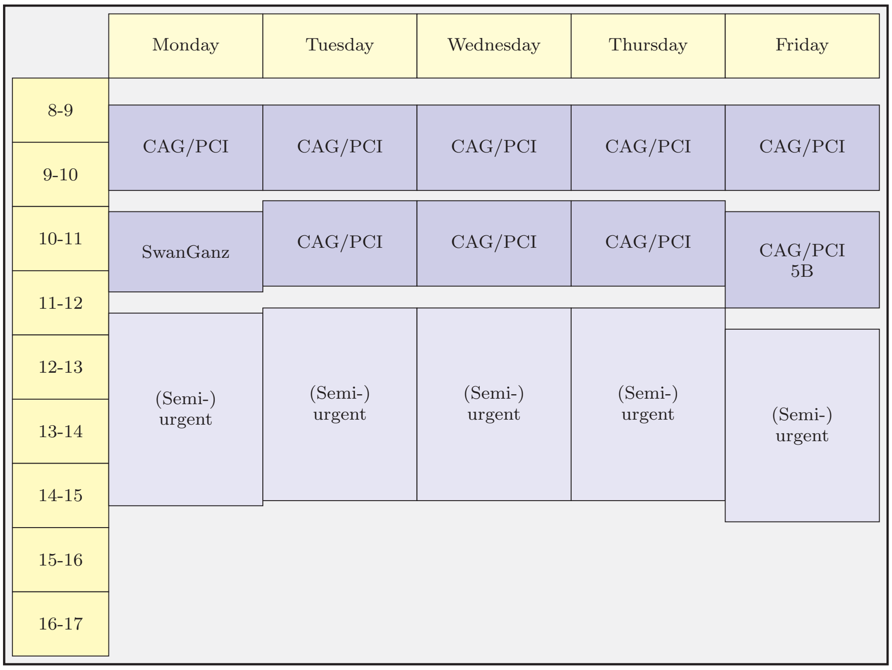

Fig. 6 Blueprint room 3

procedure generally has a much longer procedure time, while a PCI-procedure often results in an extra day aftercare at ward 5B. This stabilizes the bed occupancy on 5B and 5C as beds are more often used in the weekend, creating more space throughout the week.

A rule of thumb is therefore clearly to plan the procedures with the longest preparation on the Monday (possibly on Tuesday, since the Monday is usually not crowded), and that procedures with the longest aftercare are to be scheduled at Friday. Note that the occupation at 5B on Monday according to this blueprint can be equal to 4; 2 implants and a Swan Ganz procedures are treated on Monday, while a patient with a long ablation on Tuesday also needs to be hospitalized that Monday.

Occupancy rate 5C Using these blueprints, we consider the occupancy rate of ward 5C. This ward has shown in practice to have the most fluctuations. By simulating the working process on the ward he simulation model, a comparison can be made between utilization using the current scheduling methods, and the utilization using the blueprints following from the model, as visualized below.

For ward 5C, a notable improvement can be seen as the peak on the Thursday has been reduced, and a much more stable pattern throughout the whole week can be seen. This pattern is mainly due to the fact that procedures with a long preparation time have been scheduled on the Monday (such that the Sunday is used optimally).

4.3 Scenarios

The possibilities of the mathematical model, combined with the simulation model used for this research, are elaborated on in this section by considering scenarios. Using the blueprints that follow from the model, the current situation is once more simulated, after which a variety of other scenarios are simulated. Within this research, a scenario is referred to as a specific change in structure or size regarding the logistic process. A different distribution of the beds among the ward is therefore an example of such a scenario, if one wants to investigate whether this can reduce the overload at a ward.

These simulations were requested by the VU University Medical Center, as specific changes are very realistic to happen such as a significant increase in patients in a specific category, as new hospitals were

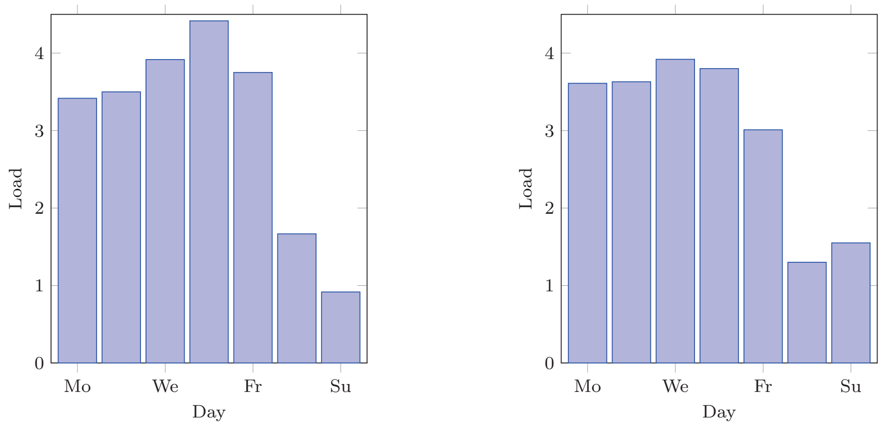

Fig. 7 Average maximum load 5C current (left) vs. optimal (right)

planning to redirect their patients to VUmc. An overview of these scenarios are explained below, after which the corresponding results can be found at the end of this section.

Current situation Although simulations are meant to research changes in the process, it is of vital importance to simulate the current situation. In this way, the simulation model can be validated to see whether it actually simulates the correct process and to see whether the results are in correspondence with the practice (i.e., the data analysis). In other words, by simulation the current situation, one gains insight in how close the model lies to reality.

In case the utilization rate of a specific cath lab lies a few percent below the percentage in the data analysis, one should take at the interpretation of the simulation results (of other scenarios) into account that the simulated utilization rate also should be a few percentages lower.

Growth scenario A very realistic scenario for VUmc is the scenario where the number of patients increases significantly, because VUmc will most likely take over the cardiological group of patients of a different hospital. In expectation, this leads to an increase of 240 CAG/PCI's (1 per day), 40 implants and 80 EFO/ablations per year. To facilitate this growth, the current idea is to increase the number of beds from the Special Care from 4 to 7 beds, while 5C gains one extra bed.

Different starting times cath labs The growth scenario is very likely to force extra capacity on the wards, but in case there is also too little capacity on the cath labs, the hospital considers to change the starting time of room 2 from 09:30 to 08:15.

Four days EFO/ablations Currently, EFO/ablations are performed on only three days (Wednesday, Thursday and Friday) due to the availability of the cardiologists. However, when the number of EFO/ablations increases with 80 per year (2 per week), the cardiologists may consider to extend their availability to four days.

The results for the simulation of specific scenarios are given in Table 4. For completeness, a simulation has been performed of one run, representing 100 years of the exact same setting.
This table contains a large variety of information, but it is important to verify the simulation model first (i.e., to see whether the simulation model is close to the realistic process), which can be seen in the first two columns. The simulation model produces utilizations (of the cath labs) that lie approximately two percent lower than in practice (i.e., the utilization that follows from the data analysis). As a heuristic way to correct the results, one could simply add two percent to the utilization rate to fix the gap between the simulation model and the realistic process. These extra two percents clearly have not been included in any of the last four columns for consistency, but will be referred to in the remainder of this section. An explanation for this difference is that the schedulers of the hospitals most likely can solve unexpected difficulties in a more efficient way. The overtime from the model is very close to the results from the data analysis, but the undertime shows a significant, but not extreme, difference of approximately 5 to 20 minutes.

	Data analysis	Current situation	Growth	Growth + extension room 2	Growth + extension room 2 + 4 days EFO
Utilization room 1	66.4%	64.5%	75.8%	73.7%	74.1%
Utilization room 2	70.3%	69.3%	76.9%	74.5%	75.3%
Utilization room 3	73.9%	70.5%	78.5%	76.7%	77.0%
Undertime room 1	69.1 m	86.4 m	38.2 m	45.0 m	46.5 m
Undertime room 2	43.9 m	58.4 m	39.1 m	50.7 m	46.8 m
Undertime room 3	46.9 m	50.6 m	43.5 m	51.8 m	50.1 m
Overtime room 1	9.9 m	11.2 m	21.2 m	18.2 m	14.8 m
Overtime room 2	18.5 m	19.1 m	32.6 m	17.3 m	16.7 m
Overtime room 3	30.8 m	29.0 m	38.5 m	22.4 m	22.2 m

Table 4 Overview simulation results of different scenarios

With this information, one can interpret the three scenarios of interest much more accurate. Clearly, the utilization rate will increase due to the increase in patients (in all scenarios). In this case study, an increase of the utilization of 10% is very realistic to happen for room 1 using this model, while room 2 and 3 are expected to have an extra 5 to 10% increase. For the scenario with all three possible changes (growth + extension room 2 + 4 days EFO), an average utilization of respectively 79%, 77% and 77% is obtainable (after the simulation correction of 2%). For the department of cardiology, this would a very positive development, as more efficient usage of their rooms is of high importance.

The situation is investigated in which cardiologists that can perform EFO/ablations should be available four days per week (instead of three). Minor improvements are visible (an utilization increase of about 0.5%), but it is questionable whether this improvement is significant. For the cardiologists at room 1, overtime decreases with 6 minutes on average per day, which is considerable, but the cardiologists are free to choose whether this is worth for them.

5 Concluding remarks

In this paper, we proposed a technique that integrates two optimization problems: the scheduling of patients at the cath labs, and the logistics of patients at the wards where patients need preparative and aftercare. By linearizing the overshoot at cath labs and wards, we can formulate an Integer Linear Program that outputs within a reasonable amount of time clear blueprints in a practical context.

We finally note that linearization of overtime is well usable in practice since the accuracy can be increased as long as the computation time permits. However, the assumption that σ^2_{dr} can be expressed as a linear function of μ_{dr}, i.e., $\sigma^2_{dr} = Q_r \cdot \mu_{dr}$, could be a too strong assumption in practice if there are procedure of which the duration can be well estimated, while other procedures have a large variation. In such cases, there is room for future research to incorporate other types of relationships in de model (e.g., logarithmic), depending on the diversity of the procedures in a room.

References

1. I. Adan, J. Bekkers, N. Dellaert, J. Vissers, & X. Yu, *Patient mix optimisation and stochastic resource requirements: A case study in cardiothoracic surgery planning*, Health Care Management Science, Vol 12(2), 129 - 141, 2009.
2. R. Bekker & A. M. de Bruin, *Time-dependent analysis for refused admissions in clinical wards*, Annals of Operations Research, Vol 178(1), 45 - 65, 2010.
3. R. Bekker & P. M. Koeleman, *Scheduling admissions and reducing variability in bed demand*, Health Care Management Science, Vol 14(3), 237 - , 2011.
4. J. Beliën & E. Demeulemeester, *Building cyclic master surgery schedules with leveled resulting bed occupancy*, European Journal of Operational Research, Vol 176, 1185 - 1204, 2005.
5. B. Cardoen, E. Demeulemeester & J. Beliën, *Operating room planning and scheduling: A literature review*, European Journal of Operational Research, Vol 201(3), 921 - 932, 2010.
6. T. Cayirly & E. Veral, *Outpatient Scheduling in Health Care, A Review of Literature*, Production and Operations Management, Volume 12, No. 4, 2003.
7. E. Demeulemeester, J. Beliën, B. Cardoen & M. Samudra, *Operating room planning and scheduling*, in Handbook of Healthcare Operations Management (pp. 121-152), Springer, NY, 2013.
8. S. Gallivan, & M. Utley, *Modelling admissions booking of elective in-patients into a treatment centre*, IMA Journal of Management Mathematics, Vol 16(3), 305 - 315, 2005.
9. D. Gupta & B. Denton, *Appointment scheduling in health care: Challenges and opportunities*, IIE Transactions, Volume 40, 2008.
10. D. Gupta, M. K. Natarajan, A. Gafni, L. Wang, D. Shilton, D. Holder & S. Yusuf, *Capacity planning for cardiac catheterization: a case study*, Health policy, Vol 82(1), 1 - 11, 2007.

11. J. E. Helm, & M. P. van Oyen, *Design and optimization methods for elective hospital admissions*, Operations Research, Vol 62(6), 1265 - 1282, 2014.
12. P. J. Hulshof, N. Kortbeek, R. J. Boucherie, E. W. Hans, & P. J. Bakker, *Taxonomic classification of planning decisions in health care: a structured review of the state of the art in OR/MS*, Health systems, Vol 1(2), 129 - 175, 2012.
13. O. Levtzion-Korach, A. Reitman, P. Jansen & S. Madden, *Doing More with Less Overtime: Improving Patient Flow Through the Cath Lab*, Cath Lab Digest, Vol 16(10), 2008.
14. E. Litvak & M. Long, *Cost and Quality Under Managed Care: Irreconcilable Differences?*, The American Journal of Managed Care, Vol 6, 305 - 312, 2000.
15. S. Mohan, Q. Li, M. Gopalakrishnan, J. Fowler & A. Printezis, *Improving the process efficiency of catheterization laboratories using simulation*, Health Systems, Vol 6(1), 41 - 55, 2017.
16. J. Papin, *A Suggested Approach for Improving Flow in the Cardiac Catheterization Laboratory*, Cath Lab Digest, Vol 21(7), 2013.
17. A. Rais & A. Viana, *Operations research in healthcare: a survey*, International Transactions in Operational Research, Vol 18(1), 1 - 31, 2011.
18. P. S. Stepaniak, M. A. Soliman Hamad, L. R. C. Dekker & J. J. Koolen, emphImproving the efficiency of the cardiac catherization laboratories through understanding the stochastic behavior of the scheduled procedures, Cardiology Journal, Vol 21, 343 - 349, 2014.
19. J. M. van Oostrum, M. van Houdenhoven, J. L. Hurink, E. W. Hans, G. Wullink & G. Kazemier, *A master surgical scheduling approach for cyclic scheduling in operating room departments*, OR spectrum, Vol 30(2), 355 - 374, 2008.
20. P. T. Vanberkel, R. J. Boucherie, E. W. Hans, J. L. Hurink, W. A. van Lent, & W. H. van Harten, *An exact approach for relating recovering surgical patient workload to the master surgical schedule*, Journal of the Operational Research Society, Vol 62(10), 1851 - 1860, 2011.