Integrated patient-to-room and nurse-to-patient assignment in hospital wards

Tabea Brandt[1], Tom Lorenz Klein[2], Melanie Reuter-Oppermann[3,4], Fabian Schäfer[5*], Clemens Thielen[2,6], Maartje van de Vrugt[3,7] and Joe Viana[8]

[1]Combinatorial Optimization, RWTH Aachen University, Ahornstr. 55, 52074 Aachen, Germany.

[2]TUM Campus Straubing for Biotechnology and Sustainability, Weihenstephan-Triesdorf University of Applied Sciences, Am Essigberg 3, 94315 Straubing, Germany.

[3]Center for Healthcare Operations Improvement & Research, University of Twente, Drienerlolaan 5, 7522 NB Enschede, The Netherlands.

[4]Information Systems, Software & Digital Business Group, Technical University of Darmstadt, Hochschulstr. 1, 64289 Darmstadt, Germany.

[5]Chair of Supply and Value Chain Management, Technical University of Munich, Am Essigberg 3, 94315 Straubing, Germany.

[6]Department of Mathematics, School of Computation, Information and Technology, Technical University of Munich, Bolzmannstr. 3, 85748 Garching bei München, Germany.

[7]Department of Care Support, Strategy and Innovation, Amsterdam University Medical Centers, De Boelelaan 1117, 1081 HV Amsterdam, The Netherlands.

[8]Department of Accounting and Operations Management, BI Norwegian Business School, Nydalsveien 37, N-0484 Oslo, Norway.

*Corresponding author(s). E-mail(s): fab.schaefer@tum.de;
Contributing authors: brandt@combi.rwth-aachen.de;
tom.klein@tum.de; m.n.reuter-oppermann@utwente.nl;
clemens.thielen@tum.de; n.vandevrugt@amsterdamumc.nl;
joe.viana@bi.no;

1

Abstract

Assigning patients to rooms and nurses to patients are critical tasks within hospitals that directly affect patient and staff satisfaction, quality of care, and hospital efficiency. Both patient-to-room assignments and nurse-to-patient assignments are typically agreed upon at the ward level, and they interact in several ways such as jointly determining the walking distances nurses must cover between different patient rooms. This motivates to consider both problems jointly in an integrated fashion. This paper presents the first optimization models and algorithms for the integrated patient-to-room and nurse-to-patient assignment problem. We provide a mixed integer programming formulation of the integrated problem that considers the typical objectives from the single problems as well as additional objectives that can only be properly evaluated when integrating both problems. Moreover, motivated by the inherent complexity that results from integrating these two NP-hard and already computationally challenging problems, we devise an efficient heuristic for the integrated patient-to-room and nurse-to-patient assignment problem. To evaluate the running time and quality of the solution obtained with the heuristic, we conduct extensive computational experiments on both artificial and real-world instances. The artificial instances are generated by a parameterized instance generator for the integrated problem that is made freely available.

Keywords: Integrated Planning, Hospital, Patient-to-room assignment, Nurse-to-patient assignment, Heuristic

1 Introduction

For many years, an ever-rising demand for healthcare and increasing healthcare expenditures challenge hospitals to increase the efficiency of their operations [18]. This results in a need for advanced managerial planning approaches that help to use the available scarce resources as efficiently as possible. Consequently, a wide range of methods and approaches have been developed in the Operations Research (OR) literature that aims at improving resource utilization through efficient planning [30, 31, 43]. In particular, quantitative decision support has been proposed for important resources such as operating rooms [22, 55], intensive care units [3], inpatient beds [26], physicians [19], and nurses [5, 14].

While efficient planning of single resources can already lead to improved resource utilization and large efficiency gains, it ignores the inherent complex interactions between different resources [30] and, as a consequence, might lead to suboptimal decisions on a system level. Therefore, a need for OR models and methods for *integrated planning* of several resources has been identified [30, 32, 56]. This need is particularly apparent in hospitals, where different resources

are typically required and used for treating patients. A recent literature review on integrated planning of multiple resources in hospitals is provided in [42].

Rooms and beds are critical assets of hospitals since they account for a considerable part of a hospital's infrastructure and large financial investments are necessary for equipping them with medical devices that facilitate patient care [58]. Additionally, a shift in demographics, the growing number of patient admissions, and rising inpatient units costs lead to high overall bed occupancy levels and require an increased focus on efficient bed management [26, 46]. On the operational level, an important planning problem typically referred to as *patient-to-room assignment (PRA)* consists of assigning patients to suitable rooms such that a variety of constraints concerning, for instance, the patient's medical needs (e.g., required medical equipment) and preferences (e.g., concerning age and gender of roommates) are satisfied, while available room capacities are respected and transfers of patients between rooms are avoided [11, 12, 17, 46]. Different variants of the PRA problem have been studied extensively in the literature both in the static setting, when all information about patients and their admission and discharge times is known in advance, and in dynamic settings, when new patients may arrive unexpectedly – see Section 2.1 for a detailed overview.

Medical staff also represent a particularly critical resource in hospitals since (1) medical staff are involved in most patient-related activities in a hospital, and (2) medical staff are a particularly scarce resource due to a general shortage of nurses [2] and physicians [7, 19, 52]. This has led to increasing workloads for the staff over the last decades and makes good planning of medical staff a central concern for hospitals [5]. In particular, aspects such as a fair distribution of workloads among staff members have a large impact on employee satisfaction and the efficient operation of a hospital. Concerning the nursing staff, distributing work fairly among nurses is considered essential for optimal quality of care [38]. Here, the workload of each single nurse is mostly determined by the patients the nurse is assigned to and their care requirements. Consequently, determining a suitable *nurse-to-patient assignment (NPA)* that balances the workloads of the nurses represents an important operational problem that has received considerable attention in the literature – see Section 2.2 for an extensive literature review.

With a few exceptions detailed in Section 2.3, the PRA problem and the NPA problem have mostly been considered separately in the literature – although there are important interactions between them. For instance, studies show that the walking distances that result from traveling between their assigned patients' rooms and other locations such as the nearest nursing station have a substantial impact on a nurse's workload during a shift [1, 10]. These walking distances, however, can only be determined and optimized when considering PRAs and NPAs jointly. Moreover, it has been observed that assigning the minimum possible number of nurses to patients in the same room can help to minimize negative effects such as the spread of infections between

rooms by nurses or the disturbance of patients by other nurses entering their room [15, 16, 20, 24]. This provides a strong motivation for integrating the PRA problem and the NPA problem by considering them jointly in one optimization model.

While several publications motivate and discuss the integration of the two assignment problems (see Section 2.3), this paper explicitly considers decisions on PRAs and NPAs in one integrated optimization problem for the first time. Besides the objectives classically considered in the two separate problems, this integrated problem also allows the evaluation of additional objectives that rely on the interaction of PRAs and NPAs. Based on existing studies on nurse workloads [1, 10], these objectives include the walking distances of nurses between assigned patients' rooms and additional relevant rooms such as the nearest nursing station. Moreover, also motivated by findings from the literature [15, 20, 24], assigning the minimum possible number of nurses to patients in the same room is considered as an objective in order to mitigate negative effects such as the spread of infections between rooms by nurses or the disturbance of patients by other nurses entering their room.

In order to solve the computationally challenging integrated PRA and NPA problem, we provide a formulation of the problem as a mixed integer program (MIP) as well as an efficient heuristic. The heuristic extends the heuristic for the PRA problem presented in [46] to the integrated problem and additionally employs a new heterogeneity check between patient admission and discharge times for the room assignment part. Both the MIP and the heuristic are evaluated in extensive experimental results on real-world instances obtained from a ward of our partner hospital Amsterdam University Medical Centers (Amsterdam, The Netherlands) as well as artificial instances. The artificial instances are generated by a parameterized instance generator for the integrated problem that is made freely available. While the MIP only addresses the static version of the integrated problem in which all information about patients is known in advance, the heuristic can be easily adapted to dynamic settings where new patients arrive after some assignments have already been fixed.

The remainder of the paper is structured as follows. Section 2 summarizes the related literature on the PRA and the NPA problems as well as existing work related to their integration. Afterward, Section 3 introduces the integrated PRA and NPA problem, while Section 4 presents our MIP formulation of the problem. Afterward, Section 5 presents a sequential solution approach based on a natural decomposition of the MIP formulation as well as our heuristic for the problem. Section 6 then describes both the developed instance generator for generating artificial test instances as well as the real-world instances obtained from a ward of our partner hospital. Our experimental results obtained on both types of instances are presented in Section 7. The paper concludes in Section 8 with a summary and an outlook on future research.

2 Related literature

In this section, we first summarize the state of the art concerning PRA and NPA separately before discussing existing work on the integration of the two problems.

2.1 Patient-to-room assignment

The static version of the PRA problem has been formally introduced by Demeester et al. [17] under the name of patient admission scheduling problem.[1] In this version, all information about patients is known in advance, and the task is to assign patients to suitable rooms such that room capacity and gender policy are respected while minimizing both patient transfers and penalty costs for undesirable PRAs. One important characteristic in this definition is that not every patient can be assigned to every room and patients may also have preferences towards specific rooms based on, e.g., available equipment or the number of beds.

For the static PRA problem, mostly heuristic solution approaches are proposed in the literature, e.g., a tabu search algorithm [17], a local search algorithm [11], a destroy and repair matheuristic [23], and algorithms based on the Hungarian algorithm [8], column generation [44], or MIP [54]. Currently the best solutions for the benchmark instances provided by Demeester et al. [17] are found by the heuristic of Guido et al. [23] and by the exact, MIP-based solution approach proposed by Bastos et al. [4]. However, the exact approach uses considerably more computation time.

The complexity of the static PRA problem is studied by Vancroonenburg et al. [57] using its correspondence to the red-blue-transportation problem. They show that the PRA problem is $\mathcal{NP}$-hard in general and even if all rooms have a capacity of three. This result is strengthened by Ficker et al. [21], who prove that the PRA problem is also $\mathcal{NP}$-hard if all rooms have capacity two.

Ceschia and Schaerf [11, 12] propose a dynamic version of the PRA problem that includes the handling of emergency patients and uncertainty in the patients' length-of-stay. They propose a metaheuristic based on simulated annealing and an instance generator as well as a set of benchmark instances. Vancroonenburg et al. [58] studied a similar problem version using two online integer linear programming (ILP) models. Lusby et al. [35] propose a large neighborhood search heuristic for the dynamic PRA problem as proposed by Ceschia and Schaerf.

A different approach for incorporating emergency patients is taken by Schäfer et al. [45–47], who use a rolling horizon approach, i.e., recomputation of the solution whenever a new event occurs, using a fast heuristic. In their problem definition, they consider objectives for three stakeholders (patients,

---

[1]We refer to the problem exclusively as the PRA problem in the following since this is the most common term used in the recent literature. Moreover, the original term *patient admission scheduling problem* is also used with a different meaning in the literature.

nurses, and physicians) and they are also the first to consider interdependencies between patients in the same room.

For a more detailed overview of the different versions and solution approaches for the PRA problem, we refer to Zhu et al. [59], who study the compatibility of short-term and long-term objectives in the context of dynamic PRA.

2.2 Nurse-to-patient assignment

The NPA problem is also considered by many different authors, where the most common objective is balancing the workloads of the nurses. For instance, Mullinax and Lawley [38] use this objective in the daily assignment of newborn infants to nurses in an intensive care nursery that is divided into several zones (rooms). Here, each infant might yield a different workload depending on their acuity and each nurse can be assigned a certain maximum number of infants to take care of at the beginning of the shift, but all of these need to be from the same zone. Since they find the problem to be too hard to solve using an integer program, they present a two-step heuristic approach that exploits the subdivision into zones by first computing the number of nurses allocated to each zone before assigning patients to nurses in each zone independently.

The problem introduced by Mullinax and Lawley [38] is also considered in [33, 36, 48, 50] – each time with a slightly different objective function. Based on the formulation proposed in [38], Sir et al. [50] formulate four MIP models for NPA that model the workloads of nurses by either the patient acuity indicators from a patient classification system (PCS), survey-based nurse-specific workload scores, or a combination of the two. Ku et al. [33] focus on minimizing the variance of the nurses' workloads using mixed integer quadratic programming (MIQP) and constraint programming (CP), while Schaus and Régin [48] minimize the variance using a two-step decomposition approach that first computes the number of nurses allocated to each zone (which is done optimally by solving a resource allocation problem) before assigning patients to nurses in each zone independently using CP. Finally, Marzouk and Kamoun [36] formulate a binary integer program that assigns nurses to zones and individual patients with the objective of minimizing the total number of nurses used in a shift.

Other work on NPA includes Punnakitikashem et al. [40], who present a stochastic integer programming model with the objective of minimizing excess workload for nurses and compare their approach to several other assignment policies (random assignment without considering workload, a simple heuristic, and solving the mean value problem using a deterministic integer program). Sundaramoorthi et al. [51] then use three of the assignment policies from [40] as well as a clustered assignment policy to test their developed simulation model for evaluating NPAs.

While most of the literature on NPA mentioned above uses patient acuity as the main factor influencing nursing workloads, Acar and Butt [1] perform

a detailed study in order to identify the activities that comprise a nurse's workload. They find that nurses spend a substantial part of their time traveling (walking) between locations, where travel between patient rooms and the nursing station is the most common type of travel. Here, according to Butt et al. [10], the distance traveled by nurses is correlated to their assigned patient load and location, and key distances influencing the total travel distance of a nurse are the distances between assigned patients' rooms and (1) the nearest nursing station, (2) the nearest supply room, and (3) other assigned patients' rooms. Still, walking distances of nurses have not yet been considered explicitly as an objective in the literature on NPA since their minimization requires the simultaneous optimization of PRAs. This also holds for the objective of assigning the minimum possible number of nurses to patients in the same room, although it is known that assigning all patients in the same room to the same nurse or a small pool of nurses helps to avoid the transfer of hospital-acquired infections [15, 20, 24], in particular *Methicillin-resistant Staphylococcus aureus* (MRSA) [16].

2.3 Integration of PRA and nurse-to-patient assignment

While we are not aware of any papers that explicitly integrate decisions concerning PRA and NPA in one optimization model, there still exists some literature considering the interplay between the two problems or between related problems. For instance, Thomas et al. [53] present a mixed-integer goal programming model for the PRA problem that takes nurses into account via constraints on the required nurse-to-patient ratio in each unit of a hospital. Bilgin et al. [6] develop a general, high-level hyper-heuristic approach that can be used for both the PRA problem and the nurse rostering problem. Pesant [39] addresses the integration of the nurse staffing problem (assigning an appropriate number of nurses to each unit within a ward given a nurse roster) and the NPA problem in a neonatal intensive care unit by solving CP models for the two problems consecutively, and Punnakitikashem et al. [41] extend the stochastic programming model from [40] by integrating nurse staffing decisions over multiple units into the NPA problem. Moreover, several recent papers consider patient appointment planning in outpatient chemotherapy clinics while simultaneously assigning nurses to patients or taking constraints on nurse availability into account [9, 27–29, 34]. Schmidt et al. [49] integrate patient appointment planning into a patient-to-ward assignment problem, where they assume that rooms in the same ward are equal. They propose a binary integer program for this problem and compare exact and heuristic solution approaches. Ceschia and Schaerf [13] consider an integrated planning of PRA with operating room constraints where they allow the postponement of patient appointments, for which they also provide an instance generator and a set of benchmark instances.

3 Problem definition

In this section, we formally introduce the *integrated patient-to-room and nurse-to-patient assignment (IPRNPA)* problem as well as the sets and parameters that are used to represent the input of the problem.

The IPRNPA problem integrates PRA and NPA on the ward level. Thus, the problem consists of assigning patients to rooms and nurses to patients on a hospital ward over a given planning period (typically one or several weeks). This section describes the static version of the problem where, similar to the static version of the PRA problem [17], all information about the patients (in particular, each patient's admission and discharge times) are known at the beginning of the planning period. The dynamic version of the problem differs from the static version in that information about new patients only becomes known either when they are admitted or a fixed time span before admission, which is analogous to the existing literature on dynamic PRA [11, 12].

We are given a set $\mathcal{P}$ of patients, a set $\mathcal{N}$ of nurses, and a set $\mathcal{R}$ of rooms. Moreover, there exists a (usually small) set $\mathcal{A}$ of additional rooms (such as the nursing station) with $\mathcal{R} \cap \mathcal{A} = \emptyset$. These additional rooms cannot be used for assigning patients and will only be relevant when computing the walking distances of the nurses.

The considered planning period consists of a set $\mathcal{S} = \{1, \ldots, S\}$ of shifts, which is partitioned into the subsets $\mathcal{S}^{\text{early}}$ of early shifts, $\mathcal{S}^{\text{late}}$ of late shifts, and $\mathcal{S}^{\text{night}}$ of night shifts. The shifts are numbered chronologically starting with an early shift and ending with a night shift. Hence, the first early, late, and night shift are numbered 1, 2, and, 3, respectively, and belong to the first day, whereas the second early, late, and night shift belong to the second day and so on.

A feasible PRA demands that each patient $p \in \mathcal{P}$ is assigned to exactly one room $r \in \mathcal{R}$ during each shift between their admission shift ad_shift$(p) \in \mathcal{S}^{\text{early}}$ and their discharge shift di_shift$(p) \in \mathcal{S}^{\text{night}}$, which denotes the first and the last shift, respectively, of the patient's stay on the ward. In particular, this means that patients are always admitted and discharged in the morning between a night shift and the following early shift, as is the case in many real hospital wards. If patient p has already been on the ward during the last shift of the previous planning period, the value ad_shift(p) is set to $0 \notin \mathcal{S}$, and if patient p will still be on the ward after the last shift S of the current planning period, the value di_shift(p) is set to $S+1 \notin \mathcal{S}$. Patient transfers between rooms are possible and are assumed to take place at most once a day for each patient between a night shift and an early shift. Transfers are, however, undesirable for both patients and nurses, and should, thus, be minimized. If patient p has already been on the ward during the last shift of the previous planning period, the room $y^{\text{prev}}(p) \in \mathcal{R}$ that the patient has been assigned to during this shift is also given as an input. This allows the evaluation of transfers that happen

between the last (night) shift of the previous planning period and the first (early) shift of the current planning period.

Requirements concerning the PRA include respecting the capacity of each room $r \in \mathcal{R}$, which is given shift-independent number of available beds denoted by num_beds(r) (usually between 1 and 4) that defines the maximum number of patients that can be assigned to room r during any single shift. Moreover, depending on their specific condition, a patient might benefit from certain types of equipment in their room during certain shifts, so, during these shifts, they should be assigned to a room that features this type of equipment if possible. The set of possible equipment types is denoted by $\mathcal{E}$. The types of equipment that are present in room r are represented by the subset $\mathcal{E}(r) \subseteq \mathcal{E}$, and the shift-specific types of desired equipment of patient p during shift s are represented by the subset $\mathcal{E}(p,s) \subseteq \mathcal{E}$. Additionally, gender-mixed rooms should be avoided if possible. To this end, the set of patients is partitioned into the subsets $\mathcal{F}$ of female patients and $\mathcal{M}$ of male patients (i.e., $\mathcal{P} = \mathcal{F} \dot{\cup} \mathcal{M}$) and the number of gender-mixed rooms should be minimized across all shifts. Finally, each patient $p \in \mathcal{P}$ has an associated age group computed as age_group(p) = $\lfloor \frac{\text{age}(p)}{10} \rfloor$ with age(p) denoting the age of the patient in years. Age groups are relevant since large age differences between patients who are simultaneously assigned to the same room are known to result in inconvenience for the patients and should, thus, be avoided.

Concerning the NPA part of the problem, the nurse roster for the planning period is given as an input. Here, for each nurse $n \in \mathcal{N}$, we are given the subset $\mathcal{S}(n) \subseteq \mathcal{S}$ of shifts that the nurse is assigned to. A feasible NPA should assign each patient $p \in \mathcal{P}$ to exactly one nurse $n \in \mathcal{N}$ with $s \in \mathcal{S}(n)$ during each shift $s \in \mathcal{S}$ between the patient's admission shift and discharge shift. Since nurses work at most one shift per day, any patient staying on the ward for at least two shifts must necessarily be assigned to different nurses during different shifts. In order to improve continuity of care, however, the number of different nurses who treat a single patient should be minimized.

Further requirements on the NPA include respecting nurse skill level requirements of the patients. Each nurse $n \in \mathcal{N}$ has a skill level skill_level(n) and each patient $p \in \mathcal{P}$ requires a certain minimum skill level skill_req(p,s) during each shift $s \in \mathcal{S} \setminus \mathcal{S}^{\text{night}}$ between their admission shift and their discharge shift. The set of possible skill levels is denoted by $\mathcal{L} = \{1,2,3\}$, where 1 = trainee, 2 = regular, and 3 = experienced. While an experienced nurse (skill level 3) can take care of any patient, the assignment should ensure that regular nurses (skill level 2) and trainees (skill level 1) are only assigned patients whose required skill level during a specific shift does not exceed the nurse's skill level.

Moreover, patients can induce different, shift-dependent workloads for nurses, and these workloads are to be distributed fairly among the nurses - both during each single shift and overall. The workload resulting from taking care of patient $p \in \mathcal{P}$ during shift $s \in \mathcal{S}$ is expressed by a nonnegative

number $\text{w\_load}(p, s)$ and depends on the age group of the patient, on their specific condition, on the time since admission, and on whether the shift is a day shift or a night shift. A fair distribution of workload among the nurses is then achieved by defining a maximum desired workload $\text{max\_load}(n, s)$ for each nurse $n \in \mathcal{N}$ during each shift $s \in \mathcal{S}(n)$ and ensuring that this maximum workload is not exceeded if possible while at the same time making sure that the assigned workloads relative to the respective maxima do not differ too much between nurses during single shifts as well as overall.

Important considerations that are influenced by both the PRA as well as the NPA involve trying to assign all patients in a room to the same nurse during each shift (which minimizes the spread of infections across rooms by nurses and also leads to patients being disturbed less by other nurses entering their room) as well as the minimization of the walking distances of nurses between different rooms. For the evaluation of the walking distances, we are given a nonnegative number $\text{dist}(r, r')$ for each two rooms $r, r' \in \mathcal{R}$ that specifies the walking distance between rooms r and r'. Similarly, for each additional room $a \in \mathcal{A}$ and each regular room $r \in \mathcal{R}$, there is a nonnegative number $\text{dist}(a, r)$ that specifies the walking distance between additional room a and room r. The actual walking distances of nurses during a shift $s \in \mathcal{S}$ then depend on the patients they are assigned to during shift s and on how frequently nurses are expected to walk in a *circular pattern* directly from patient to patient during shift s (e.g., when rounds are made during early shifts) and on how frequently they are expected to walk in a *star-like pattern* directly from additional rooms (such as the nursing station) to patients and back (e.g., when a patient calls for a nurse). These expected (absolute) frequencies are represented by two nonnegative parameters $\text{walk\_pat}^{\circ}(s)$ (*circular*) and $\text{walk\_pat}^{\star}(s)$ (*star-like*), respectively.

4 Mixed integer programming formulation

In order to model the IPRNPA problem introduced in Section 3 mathematically, we now present a formulation of the problem as a MIP.

The following lists summarize the sets and parameters introduced in the previous section as well as the decision variables used in the MIP:

Sets:

$\mathcal{P}$	set of patients (index p)
$\mathcal{F}$	subset of female patients
$\mathcal{M}$	subset of male patients
$\mathcal{N}$	set of nurses (index n)
$\mathcal{N}^{\text{prev}}(p)$	subset of nurses that patient $p \in \mathcal{P}$ has already been assigned to during at least one shift in a previous planning period
$\mathcal{R}$	set of rooms (index r)
$\mathcal{A}$	set of additional rooms (e.g., the nursing station) (index a)

$\mathcal{S} = \{1, \dots, S\}$	set of shifts (index s)
$\mathcal{S}^{\text{early}}$	subset of early shifts
$\mathcal{S}^{\text{late}}$	subset of late shifts
$\mathcal{S}^{\text{night}}$	subset of night shifts
$\mathcal{S}(n)$	subset of shifts that nurse $n \in \mathcal{N}$ is assigned to
$\mathcal{E}$	set of possible equipment types in the rooms (index e)
$\mathcal{E}(r)$	subset of equipment types present in room $r \in \mathcal{R}$
$\mathcal{E}(p,s)$	subset of desired equipment types of patient $p \in \mathcal{P}$ during shift $s \in \mathcal{S}$
$\mathcal{L} = \{1,2,3\}$	set of possible skill levels of nurses (index l), where $1 =$ trainee, $2 =$ regular, $3 =$ experienced

Parameters:

ad_shift(p)	shift $s \in \mathcal{S}^{\text{early}}$ during which patient $p \in \mathcal{P}$ is admitted (first shift in which a bed is required for patient p). The value is set to 0 if patient p has already been on the ward during the last shift of the previous planning period
di_shift(p)	shift $s \in \mathcal{S}^{\text{late}}$ during which patient $p \in \mathcal{P}$ is discharged (last shift in which a bed is required for patient p). The value is set to $S + 1$ if patient p will still be on the ward after the last shift S of the planning period
$y^{\text{prev}}(p)$	room $r \in \mathcal{R}$ that patient $p \in \mathcal{P}$ with ad_shift$(p) = 0$ has been assigned to during the last shift of the previous planning period
num_beds(r)	nonnegative integer (most likely from $\{1,2,3,4\}$) specifying the number of beds in room $r \in \mathcal{R}$
age_group(p)	age group of patient $p \in \mathcal{P}$ computed as age_group$(p) = \lfloor \frac{\text{age}(p)}{10} \rfloor$ with age(p) denoting the age of the patient in years
skill_level(n)	skill level of nurse $n \in \mathcal{N}$ (possible values are $1,2,3$)
skill_req(p,s)	minimum skill level of a nurse required by patient $p \in \mathcal{P}$ during shift $s \in \mathcal{S} \setminus \mathcal{S}^{\text{night}}$ (possible values are $0,1,2$)
w_load(p,s)	nonnegative number specifying the workload resulting from taking care of patient $p \in \mathcal{P}$ during shift $s \in \mathcal{S}$
max_load(n,s)	nonnegative number specifying the maximum workload allowed for nurse $n \in \mathcal{N}$ during shift $s \in \mathcal{S}(n)$
dist(r,r')	nonnegative number specifying the walking distance between rooms $r, r' \in \mathcal{R}$ (where dist$(r,r') =$ dist(r',r) for all $r, r' \in \mathcal{R}$, i.e., distances are symmetric)
dist(a,r)	nonnegative number specifying the walking distance between additional room $a \in \mathcal{A}$ and room $r \in \mathcal{R}$
walk_pat$^{\circ}(s)$	nonnegative weight for different walking patterns depending on the shift $s \in \mathcal{S}$. A high value of walk_pat$^{\circ}(s)$ indicates that most nurses walk directly from patient to patient during shift s (circular pattern)
walk_pat$^{\star}(s)$	nonnegative weight for different walking patterns depending on the shift $s \in \mathcal{S}$. A high value of walk_pat$^{\star}(s)$

indicates that most nurses walk directly from additional rooms such as the nursing station to patient rooms and back during shift s (star-like pattern)

Decision variables:

$y_{p,r,s}$	binary variable indicating whether patient $p \in \mathcal{P}$ is assigned to room $r \in \mathcal{R}$ during early shift $s \in \mathcal{S}^{\text{early}}$ (only defined if $\text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p)$, i.e., if patient p is on the ward during early shift s)
$\text{f\_in\_room}_{r,s}$	binary variable indicating whether at least one female patient is assigned to room $r \in \mathcal{R}$ during early shift $s \in \mathcal{S}^{\text{early}}$
$\text{m\_in\_room}_{r,s}$	binary variable indicating whether at least one male patient is assigned to room $r \in \mathcal{R}$ during early shift $s \in \mathcal{S}^{\text{early}}$
$\text{vio}^{\text{gender}}_{r,s}$	binary variable indicating whether more than one gender is accommodated in room $r \in \mathcal{R}$ during shift $s \in \mathcal{S}^{\text{early}}$
$\text{trans}_{p,s}$	binary variable indicating whether patient $p \in \mathcal{P}$ is transferred to a different room after night shift $s \in (\mathcal{S}^{\text{night}} \cup \{0\}) \setminus \{S\}$ (and before early shift $s+1$) (only defined if $\text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p) - 1$)
$\text{age\_group}^{\max}_{r,s}$	nonnegative fractional variable representing the maximum age group among all patients $p \in \mathcal{P}$ assigned to room $r \in \mathcal{R}$ during early shift $s \in \mathcal{S}^{\text{early}}$
$\text{age\_group}^{\min}_{r,s}$	nonnegative fractional variable representing the minimum age group among all patients $p \in \mathcal{P}$ assigned to room $r \in \mathcal{R}$ during early shift $s \in \mathcal{S}^{\text{early}}$
$x_{p,n,s}$	binary variable indicating whether patient $p \in \mathcal{P}$ is assigned to nurse $n \in \mathcal{N}$ during shift $s \in \mathcal{S}$, (only defined if $s \in \mathcal{S}(n)$ and $\text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p)$, i.e., if if nurse n is assigned to shift s and patient p is on the ward during shift s)
$\text{vio}^{\text{skill}}_{p,s}$	binary variable indicating whether patient $p \in \mathcal{P}$ is assigned to a nurse with a lower skill level than required during shift $s \in \mathcal{S} \setminus \mathcal{S}^{\text{night}}$ (only defined if $\text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p)$ and $\text{skill\_req}(p,s) \geq 2$, i.e., if patient p is on the ward during shift s and requires at least an experienced nurse during shift s)
$\text{ever\_assigned}_{p,n}$	binary variable indicating whether patient $p \in \mathcal{P}$ is assigned to nurse $n \in \mathcal{N}$ during at least one shift $s \in \mathcal{S}$
$\text{vio}^{\text{load}}_{n,s}$	nonnegative fractional variable representing the excess load assigned to nurse $n \in \mathcal{N}$ during shift $s \in \mathcal{S}$ (only defined if $s \in \mathcal{S}(n)$, i.e., if nurse n is assigned to shift s)

$\text{vio}^{\text{fair}}_{n,n',s}$	nonnegative fractional variable representing the excess in relative workload (relative to the desired maximum) of nurse $n \in \mathcal{N}$ compared to nurse $n' \in \mathcal{N}$ during shift $s \in \mathcal{S}$
$\text{vio}^{\text{fair}}_{n,n'}$	nonnegative fractional variable representing the overall excess in relative workload (relative to the desired maximum) of nurse $n \in \mathcal{N}$ compared to nurse $n' \in \mathcal{N}$
$\text{in\_room}_{n,r,s}$	binary variable indicating whether nurse $n \in \mathcal{N}$ is assigned at least one patient in room $r \in \mathcal{R}$ during shift $s \in \mathcal{S}$ (only defined if $s \in \mathcal{S}(n)$, i.e., if nurse n is assigned to shift s)
$\text{dist}_{n,s}$	nonnegative fractional variable representing the total walking distance for nurse $n \in \mathcal{N}$ during shift $s \in \mathcal{S}$ (only defined if $s \in \mathcal{S}(n)$, i.e., if nurse n is assigned to shift s)
$\text{both\_rooms}_{n,r,r',s}$	binary variable indicating whether nurse $n \in \mathcal{N}$ is assigned patients in both room $r \in \mathcal{R}$ and room $r' \in \mathcal{R}$ during shift $s \in \mathcal{S}$ (only defined if $\text{assign}(n,s) = 1$, i.e., if nurse n is assigned to shift s)

The objective function of the MIP that is to be minimized consists of a weighted sum of several separate objectives. These objectives include those classically considered in the PRA problem (objectives (1)–(4)) and the NPA problem (objectives (5)–(6)). Moreover, two additional objectives are considered that rely explicitly on the interaction of the two problems: Objective (7) considers the assigning the minimum number of nurses per room during each shift, while objective (8) minimizes the walking distances of nurses. The weights of the different objectives in the weighted sum are chosen based on the existing literature and discussions with our partner hospital (see Section 7 for the concrete values used in our computational experiments).

Patient transfers objective

(1) Minimization of the number of patient transfers across all patients and shifts (could also be weighted differently for different patients and / or different shifts):

$$\min \sum_{\substack{p \in \mathcal{P}, s \in (\mathcal{S}^{\text{night}} \cup \{0\}) \setminus \{S\}: \\ \text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p) - 1}} \text{trans}_{p,s}$$

Patient inconvenience objective

(2) Minimization of the age group difference across all rooms and shifts:

$$\min \sum_{r \in \mathcal{R}, s \in \mathcal{S}^{\text{early}}} (\text{age\_group}^{\text{max}}_{r,s} - \text{age\_group}^{\text{min}}_{r,s})$$

Gender mixing objective

(3) Minimization of gender mixing across all rooms and shifts:

$$\min \sum_{r \in \mathcal{R}, s \in \mathcal{S}^{\text{early}}} \text{vio}_{r,s}^{\text{gender}}$$

Equipment violation objective

(4) Minimization of required equipment violation across all rooms and shifts:

$$\min \sum_{\substack{p \in \mathcal{P}, r \in \mathcal{R}, s \in \mathcal{S}^{\text{early}}: \\ \text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p) \\ \text{and } \mathcal{E}(p,s) \setminus \mathcal{E}(r) \neq \emptyset}} y_{p,r,s}$$
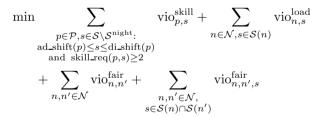
Continuity of care objective

(5) Minimization of the number of different nurses that treat each patient across all patients (could also be weighted differently for different patients):

$$\min \sum_{p \in \mathcal{P}, n \in \mathcal{N} \setminus \mathcal{N}^{\text{prev}}(p)} \text{ever\_assigned}_{p,n}$$

Penalization of skill level requirements and undesired workload distributions objective

(6) Minimization of violations of skill level requirements of patients and undesired workload distributions for nurses:

$$\min \sum_{\substack{p \in \mathcal{P}, s \in \mathcal{S} \setminus \mathcal{S}^{\text{night}}: \\ \text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p) \\ \text{and } \text{skill\_req}(p,s) \geq 2}} \text{vio}_{p,s}^{\text{skill}} + \sum_{n \in \mathcal{N}, s \in \mathcal{S}(n)} \text{vio}_{n,s}^{\text{load}}$$

$$+ \sum_{n,n' \in \mathcal{N}} \text{vio}_{n,n'}^{\text{fair}} + \sum_{\substack{n,n' \in \mathcal{N}, \\ s \in \mathcal{S}(n) \cap \mathcal{S}(n')}} \text{vio}_{n,n',s}^{\text{fair}}$$

Assigning the minimum number of nurses per room objective

(7) Minimization of the number of nurses assigned to rooms across all shifts:

$$\min \sum_{n \in N, r \in R, s \in S} \text{in\_room}_{n,r,s}$$

Walking distances objective

(8) Minimization of the walking distances across all nurses and shifts

$$\min \sum_{n \in \mathcal{N}, s \in \mathcal{S}(n)} \text{dist}_{n,s}$$

The constraints of the MIP can be formulated as follows:

(I) Assignment of patients to rooms

(9) Each patient $p \in \mathcal{P}$ is assigned to exactly one room $r \in \mathcal{R}$ during each early shift $s \in \mathcal{S}^{\text{early}}$ between their admission and discharge:

$$\sum_{r \in \mathcal{R}} y_{p,r,s} = 1 \quad \forall p \in \mathcal{P},\, s \in \mathcal{S}^{\text{early}} : \text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p)$$

(10) No room $r \in \mathcal{R}$ can be assigned more than $\text{num\_beds}(r)$ patients during any early shift $s \in \mathcal{S}^{\text{early}}$:

$$\sum_{p \in \mathcal{P}:\text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p)} y_{p,r,s} \leq \text{num\_beds}(r) \quad \forall r \in \mathcal{R},\, s \in \mathcal{S}^{\text{early}}$$

(11) The variable $\text{f\_in\_room}_{r,s}$ ($\text{m\_in\_room}_{r,s}$) is set to one if at least one female (male) patient is assigned to room r during early shift $s \in \mathcal{S}^{\text{early}}$:

$$y_{p,r,s} \leq \text{f\_in\_room}_{r,s} \quad \forall p \in \mathcal{F},\, r \in \mathcal{R},\, s \in \mathcal{S}^{\text{early}} : \text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p)$$

$$y_{p,r,s} \leq \text{m\_in\_room}_{r,s} \quad \forall p \in \mathcal{M},\, r \in \mathcal{R},\, s \in \mathcal{S}^{\text{early}} : \text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p)$$

(12) No room $r \in \mathcal{R}$ should be assigned both female and male patients during any early shift $s \in \mathcal{S}^{\text{early}}$:

$$\text{f\_in\_room}_{r,s} + \text{m\_in\_room}_{r,s} \leq 1 + \text{vio}_{r,s}^{\text{gender}} \quad \forall r \in \mathcal{R},\, s \in \mathcal{S}^{\text{early}}$$

(II) Patients transfers

(13) The patient transfer variables $\text{trans}_{p,s}$ are set correctly for each patient $p \in \mathcal{P}$ and each night shift $s \in \mathcal{S}^{\text{night}}$ between their admission and discharge:

$$y_{p,r,s+1} - y_{p,r,s-2} \leq \text{trans}_{p,s} \quad \forall p \in \mathcal{P},\, r \in \mathcal{R},\, s \in \mathcal{S}^{\text{night}} \setminus \{S\} :$$
$$\text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p) - 1$$

(14) The patient transfer variables $\text{trans}_{p,0}$ that indicate a transfer between the last shift of the previous planning period (shift 0) and the first (early) shift of the current planning period (shift 1) are set correctly for each patient $p \in \mathcal{P}$ with $\text{ad\_shift}(p) = 0$:

$$y_{p,r,1} \leq \text{trans}_{p,0} \quad \forall p \in \mathcal{P},\, r \in \mathcal{R} \setminus \{y^{\text{prev}}(p)\} : \text{ad\_shift}(p) = 0$$

(III) **Inconvenience of patients**

(15) The variable $\text{age\_group}_{r,s}^{\max}$ is restricted by the maximum age group of patients in room $r \in \mathcal{R}$ during early shift $s \in \mathcal{S}^{\text{early}}$:

$$\text{age\_group}_{r,s}^{\max} \geq \text{age\_group}(p) \cdot y_{p,r,s} \quad \forall p \in \mathcal{P},\, r \in \mathcal{R},\, s \in \mathcal{S}^{\text{early}} :$$
$$\text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p)$$

(16) The variable $\text{age\_group}_{r,s}^{\min}$ is restricted by the minimum age group of patients in room $r \in \mathcal{R}$ during early shift $s \in \mathcal{S}^{\text{early}}$:

$$\text{age\_group}_{r,s}^{\min} \leq \text{age\_group}(p) + 12 \cdot (1 - y_{p,r,s}) \quad \forall p \in \mathcal{P},\, r \in \mathcal{R},\, s \in \mathcal{S}^{\text{early}} :$$
$$\text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p)$$

Note that the coefficient 12 on the right-hand side needs to be increased if patients with ages 130 or older (age group 13 or higher) are present.

(17) The variable $\text{age\_group}_{r,s}^{\min}$ is set to zero if no patients are assigned to room $r \in \mathcal{R}$ during early shift $s \in \mathcal{S}^{\text{early}}$:

$$\text{age\_group}_{r,s}^{\min} \leq 12 \cdot \sum_{p \in \mathcal{P}:\text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p)} y_{p,r,s} \quad \forall r \in \mathcal{R},\, s \in \mathcal{S}^{\text{early}}$$

Note that the coefficient 12 on the right-hand side needs to be increased if patients with ages 130 or older (age group 13 or higher) are present.

(18) The value of the variable $\text{age\_group}_{r,s}^{\max}$ must not be smaller than the value of the variable $\text{age\_group}_{r,s}^{\min}$ for any room $r \in \mathcal{R}$ and any early shift $s \in \mathcal{S}^{\text{early}}$:

$$\text{age\_group}_{r,s}^{\min} \leq \text{age\_group}_{r,s}^{\max} \quad \forall r \in \mathcal{R},\, s \in \mathcal{S}^{\text{early}}$$

Note that this constraint is not required for the correctness of the model and is only used to improve solution times.

(IV) **Assignment of patients to nurses**

(19) Each patient $p \in \mathcal{P}$ is assigned to exactly one nurse $n \in \mathcal{N}$ with $s \in \mathcal{S}(n)$ during each shift $s \in \mathcal{S}$ between their admission and discharge:

$$\sum_{n \in \mathcal{N}:s \in \mathcal{S}(n)} x_{p,n,s} = 1 \quad \forall p \in \mathcal{P},\, s \in \mathcal{S} : \text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p)$$

(20) For each patient $p \in \mathcal{P}$ and each shift $s \in \mathcal{S} \backslash \mathcal{S}^{\text{night}}$ with skill_req$(p, s) \geq 2$, the variable vio$_{p,s}^{\text{skill}}$ is set to one if the patient is not assigned to a nurse with the required skill level:

$$\sum_{\substack{n \in \mathcal{N}: s \in \mathcal{S}(n) \text{ and} \\ \text{skill\_level}(n) \geq \text{skill\_req}(p,s)}} x_{p,n,s} = 1 - \text{vio}_{p,s}^{\text{skill}} \quad \forall p \in \mathcal{P}, \, s \in \mathcal{S} \setminus \mathcal{S}^{\text{night}} :$$

$$\text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p)$$
$$\text{and } \text{skill\_req}(p, s) \geq 2$$

(21) The variable ever_assigned$_{p,n}$ is set to one if and only if $n \in \mathcal{N}^{\text{prev}}(p)$ or patient $p \in \mathcal{P}$ is assigned to nurse $n \in \mathcal{N}$ during at least one shift:

$$x_{p,n,s} \leq \text{ever\_assigned}_{p,n} \quad \forall p \in \mathcal{P}, \, n \in \mathcal{N}, \, s \in \mathcal{S} : s \in \mathcal{S}(n)$$
$$\text{and } \text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p)$$
$$\text{ever\_assigned}_{p,n} = 1 \quad \forall p \in \mathcal{P}, \, n \in \mathcal{N}^{\text{prev}}(p)$$

$$\text{ever\_assigned}_{p,n} \leq \sum_{\substack{s \in \mathcal{S}: s \in \mathcal{S}(n) \text{ and} \\ \text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p)}} x_{p,n,s} \quad \forall p \in \mathcal{P}, \, n \in \mathcal{N} \setminus \mathcal{N}^{\text{prev}}(p)$$

(V) Workload of nurses

(22) For each nurse $n \in \mathcal{N}$ and each shift $s \in \mathcal{S}(n)$, any workload exceeding max_load(n, s) leads to a corresponding increase of the variable vio$_{n,s}^{\text{load}}$:

$$\sum_{p \in \mathcal{P}: \text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p)} x_{p,n,s} \cdot \text{w\_load}(p, s) \leq \text{max\_load}(n, s) + \text{vio}_{n,s}^{\text{load}}$$

$$\forall n \in \mathcal{N}, \, s \in \mathcal{S}(n)$$

(23) Fair distribution of workload during each shift: For any two nurses $n, n' \in \mathcal{N}$ assigned to a shift $s \in \mathcal{S}$, if the relative workload (relative to the desired maximum) of nurse n during shift s exceeds the relative workload of nurse n' during shift s, the variable vio$_{n,n',s}^{\text{fair}}$ must be increased correspondingly:

$$\sum_{p \in \mathcal{P}: \text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p)} x_{p,n,s} \cdot \frac{\text{w\_load}(p, s)}{\text{max\_load}(n, s)}$$

$$\leq \sum_{p \in \mathcal{P}: \text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p)} x_{p,n',s} \cdot \frac{\text{w\_load}(p, s)}{\text{max\_load}(n', s)} + \text{vio}_{n,n',s}^{\text{fair}}$$
$$\forall n, n' \in \mathcal{N}, \, s \in \mathcal{S}(n) \cap \mathcal{S}(n')$$

(24) Fair distribution of workload overall: For any two nurses $n, n' \in \mathcal{N}$, if the relative workload (relative to the desired maximum) of nurse n exceeds the relative workload of nurse n', the variable $\mathrm{vio}_{n,n'}^{\mathrm{fair}}$ must be increased correspondingly:

$$\sum_{s \in \mathcal{S}(n)} \sum_{p \in \mathcal{P} : \mathrm{ad\_shift}(p) \leq s \leq \mathrm{di\_shift}(p)} x_{p,n,s} \cdot \frac{\mathrm{w\_load}(p,s)}{\mathrm{max\_load}(n,s)}$$

$$\leq \sum_{s \in \mathcal{S}(n')} \sum_{p \in \mathcal{P} : \mathrm{ad\_shift}(p) \leq s \leq \mathrm{di\_shift}(p)} x_{p,n',s} \cdot \frac{\mathrm{w\_load}(p,s)}{\mathrm{max\_load}(n',s)} + \mathrm{vio}_{n,n'}^{\mathrm{fair}}$$

$$\forall n, n' \in \mathcal{N}$$

(VI) **Assignment of all patients in the same room to the same nurse**

(25) The variables $\mathrm{in\_room}_{n,r,s}$ are set correctly for each nurse $n \in \mathcal{N}$, each room $r \in \mathcal{R}$, and each shift $s \in \mathcal{S}(n)$:

$$\mathrm{in\_room}_{n,r,s} \geq x_{p,n,s} + y_{p,r,s} - 1 \qquad \forall p \in \mathcal{P}, \, n \in \mathcal{N}, \, r \in \mathcal{R}, \, s \in \mathcal{S}(n) \cap \mathcal{S}^{\mathrm{early}} :$$
$$\mathrm{ad\_shift}(p) \leq s \leq \mathrm{di\_shift}(p)$$

$$\mathrm{in\_room}_{n,r,s} \geq x_{p,n,s} + y_{p,r,s-1} - 1 \qquad \forall p \in \mathcal{P}, \, n \in \mathcal{N}, \, r \in \mathcal{R}, \, s \in \mathcal{S}(n) \cap \mathcal{S}^{\mathrm{late}} :$$
$$\mathrm{ad\_shift}(p) \leq s \leq \mathrm{di\_shift}(p)$$

$$\mathrm{in\_room}_{n,r,s} \geq x_{p,n,s} + y_{p,r,s-2} - 1 \quad \forall p \in \mathcal{P}, \, n \in \mathcal{N}, \, r \in \mathcal{R}, \, s \in \mathcal{S}(n) \cap \mathcal{S}^{\mathrm{night}} :$$
$$\mathrm{ad\_shift}(p) \leq s \leq \mathrm{di\_shift}(p)$$

(VII) **Walking distances of nurses**

(26) The variables $\mathrm{both\_rooms}_{n,r,r',s}$ are set correctly for each nurse $n \in \mathcal{N}$, each two rooms $r, r' \in \mathcal{R}$, and each shift $s \in \mathcal{S}(n)$:

$$\mathrm{both\_rooms}_{n,r,r',s} \geq \mathrm{in\_room}_{n,r,s} + \mathrm{in\_room}_{n,r',s} - 1 \; \forall n \in \mathcal{N}, \, r, r' \in \mathcal{R}, \, s \in \mathcal{S}(n)$$

(27) The walking distance variables $\mathrm{dist}_{n,s}$ are set correctly for each nurse $n \in \mathcal{N}$ and each shift $s \in \mathcal{S}(n)$:

$$\mathrm{dist}_{n,s} = \mathrm{walk\_pat}^{\circ}(s) \cdot \frac{1}{2} \cdot \sum_{r \in \mathcal{R}} \sum_{r' \in \mathcal{R} \setminus \{r\}} \mathrm{both\_rooms}_{n,r,r',s} \cdot \mathrm{dist}(r,r')$$

$$+ \, \mathrm{walk\_pat}^{\star}(s) \cdot \sum_{a \in \mathcal{A}} \sum_{r \in \mathcal{R}} \mathrm{in\_room}_{n,r,s} \cdot \mathrm{dist}(a,r) \; \forall n \in \mathcal{N}, \, s \in \mathcal{S}(n)$$

5 Solution methods

Addressing PRA and NPA challenges in hospitals reveals complexities that hinder timely optimal solutions. The PRA problem, even in isolation, is known to be NP-hard and only unrealistically small instances can be solved to optimality in reasonable time. Thus, it not surprising that the integrated problem

is very difficult to solve on instances of a realistic size and solving the MIP provided in the previous section takes a prohibitive amount of time on such instances (see also Section 7). Therefore, we now present two different methods for generating good solutions in reasonable computation times. The first method, which is mainly presented as a point of comparison to the completely integrated MIP from Section 4, is a sequential approach based on a natural decomposition of the integrated MIP into models for the two interacting subproblems. The second method is an efficient greedy heuristic for the integrated problem, which also allows easy adaptions to various dynamic problem versions.

5.1 Sequential solution approach

As a point of comparison to the completely integrated MIP presented in the previous section and to potentially achieve faster computation times, we consider the sequential solution of the two submodels for the PRA part and the NPA part. Here, we first solve the PRA part to determine the PRA and then solve the NPA part to determine the NPA based on the given PRA. Note that this approach of solving the two assignment problems sequentially will of course not yield optimal solutions for the integrated problem in general.

The PRA part of the model consists of the patient transfers objective (1), the patient inconvenience objective (2), the gender mixing objective (3), and the equipment violations objective (4) as well as the constraints concerning the assignment of patients to rooms (I), the patients transfers (II), and the inconvenience of patients (III) and the associated decision variables.

The NPA part of the model contains the continuity of care objective (5), the penalization of skill level requirements and undesired workload distributions objective (6), the objective for assigning the minimum number of nurses per room (7), and the walking distance objective (8). As constraints, the NPA part contains those concerning the assignment of patients to nurses (IV), the workload of nurses (V), the assignment of all patients in the same room to the same nurse (VI), and the walking distance of nurses (VII). Concerning variables, the NPA model contains all variables appearing in these objectives and constraints, but the $y_{p,r,s}$ are not decision variables anymore since their values are carried over from the solution obtained for the PRA part of the model.

5.2 Heuristic solution approach

Due to the computationally challenging nature of the PRA problem, heuristic solution methods are often used to tackle this problem in the recent literature [23, 46]. Naturally, the computational challenges become even greater when considering PRAs and NPAs jointly in an integrated optimization problem as we do in this paper. Therefore, we now introduce a heuristic solution approach for the IPRNPA problem, which extends the heuristic for the PRA

problem suggested in [46]. The presented heuristic balances computational effort and solution quality with the goal of obtaining a practically applicable procedure for timely, satisfactory outcomes.

Our heuristic is initially tailored to the static version of the IPRNPA problem, but its scope extends beyond the static setting. In fact, the heuristic does not require to have all information about patients available in advance and can easily be extended to a variety of dynamic settings. As an illustrative example, instead of restricting patient transfers between rooms to happen between a night shift and the following early shift as in the static IPRNPA problem, the heuristic could handle patient transfers whenever new information about patients is obtained. The same applies for NPAs, which could alternatively be allowed also during shifts whenever new information becomes available. Thus, even complex dynamic problem settings based on dynamic PRA problems such as the one considered by Ceschia and Schaerf [11, 12] can be handled without difficulties.

5.2.1 Algorithm description

The idea of the heuristic is to assign patients to rooms and nurses to patients in a greedy fashion. To this end, the heuristic considers the days of the planning period chronologically, where each day is represented by the corresponding early shift. For each day, the heuristic iteratively fixes both the PRA and the NPA for only this day jointly for a single patient in a way that yields the lowest current contribution to the objective function. After each such assignment, the current objective function contributions of the remaining possible assignments for the day are updated before the assignments for the next patient are fixed. Once the room and nurse assignments have been fixed for all patients that are on the ward during the considered day, the heuristic moves on to the next day. During this iterative process, we have to take into account that decisions for the two kinds of assignments are made on different time scales. Concerning PRAs, a decision about the room the patient is assigned to is made only once per day before the start of the early shift, whereas, concerning NPAs, three different nurse must be assigned to a patient for each day (for the early, late, and night shift) since each nurse works at most one shift per day.

We now describe our heuristic, whose pseudocode is shown in Algorithm 1, more formally. Here, we first describe the algorithm without the heterogeneity check between patient admission and discharge times that is represented by the heterogeneity matrix HetMatrix in the pseudocode and is motivated and explained afterwards in Section 5.2.2.

Upon initialization of the heuristic, all relevant assignment variables $x_{p,n,s}$ and $y_{p,r,s}$ are set to 0 to indicate that no assignments have been made so far. Afterwards, the days of the planning period are considered in chronological order represented by their early shifts. When considering a day of the planning period represented by the corresponding early shift $s \in \mathcal{S}^{\text{early}}$, the set of relevant patients for which room and nurse assignments are to be made for

this day is denoted by $\mathcal{P}(s) := \{p \in \mathcal{P} : \text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p)\}$. Note that $\mathcal{P}(s)$ also includes patients that have been on the ward already on the previous day if they have not been discharged yet since the room and nurse assignments for such patients have so far only been fixed up to the previous day. At each point during the assignment process for the day, the set of rooms that are still available (i.e., not yet fully occupied) is denoted by $\mathcal{R}(s) \subseteq \mathcal{R}$. The set of nurses who are on duty on the corresponding day is partitioned according to their assigned shifts into $\mathcal{N}^{\text{early}}(s)$, $\mathcal{N}^{\text{late}}(s)$, and $\mathcal{N}^{\text{night}}(s)$. Consequently, the Cartesian product $\mathcal{N}^{\text{comb}}(s) := \mathcal{N}^{\text{early}}(s) \times \mathcal{N}^{\text{late}}(s) \times \mathcal{N}^{\text{night}}(s)$ corresponds to all possible ordered triples of nurses that can potentially be assigned to a patient during the three shifts of the day that starts with the early shift $s \in \mathcal{S}^{\text{early}}$. The potential contributions to the objective function for each triple (p, n^{comb}, r) consisting of a patient $p \in \mathcal{P}(s)$, a triple $n^{\text{comb}} \in \mathcal{N}^{\text{comb}}(s)$ of nurses, and an available room $r \in \mathcal{R}(s)$ are stored in a contribution table denoted by ContribTable.

Whenever the heuristic starts considering some day of the planning period starting with early shift $s \in \mathcal{S}^{\text{early}}$, the contribution table is filled with the current contributions that correspond to the potential assignments of patients $p \in \mathcal{P}(s)$ to nurse combinations $n^{\text{comb}} \in \mathcal{N}^{\text{comb}}(s)$ and rooms $r \in \mathcal{R}(s)$. These current contributions are computed based on the changes to the objective function value that would currently be induced by the corresponding assignments. Here, the objective function is the same weighted sum of the objectives (1)–(8) as in the MIP described in Section 4.

The heuristic then fixes the assignments of rooms and nurse combinations for the day represented by early shift s for all patients in $\mathcal{P}(s)$ in a greedy fashion. This is done by iteratively identifying the triple (p, n^{comb}, r) that has the lowest contribution value in ContribTable and setting the values of the corresponding assignment variables $x_{p,n,s}$ and $y_{p,r,s}$ to 1, which means that patient p is assigned to the nurse combination n^{comb} and the room r. Subsequently, the allocated patient p is removed from $\mathcal{P}(s)$ and the room r is removed from $\mathcal{R}(s)$ in case that the assignment has resulted in full occupancy of room r on the corresponding day. Moreover, the contribution table ContribTable needs to be updated by removing all entries corresponding to patient p and possibly room r. In addition, the update involves adjusting the objective contributions for the remaining possible assignments of the day to ensure that they accurately represent the new state of PRA and NPAs. For example, the previously-calculated contribution to the gender mixing objective for room r is updated for all remaining patients in $\mathcal{P}(s)$ with a different gender than the patient p just assigned. After the updated contribution table has been computed, the heuristic continues with the next iteration based on the updated contribution table until $\mathcal{P}(s) = \emptyset$ (i.e., all assignments for the current day have been fixed) and the procedure continues with the next day.

The heuristic terminates once the last day of the planning period has been considered, i.e., once all PRA and NPAs have been fixed for the whole planning period.

5.2.2 Heterogeneity check between patients

A point that is not taken into account in the heuristic as described so far is the coordination of arrivals and discharges of patients that are assigned to the same room. If occupancy levels are high, this could lead to the unnecessary creation of gender-mixed rooms or avoidable patient transfers. For example, assigning several male patients with very similar arrival and discharge shifts to different rooms instead of the same room might leave no other rooms available for female patients admitted shortly afterward unless some of the male patients are transferred. In order to avoid such unfavorable incidents explicitly, we next describe a new patient-to-patient heterogeneity measure, whose values are stored in a heterogeneity matrix. This heterogeneity matrix is used in the heuristic to foster the assignment of patients who are discharged at similar times to the same room by explicitly taking discharge shifts into account during the computation of possible objective function contributions.

The proposed patient-to-patient heterogeneity measure is based on the difference between two patients' discharge shifts. Formally, for two patients $p, p' \in \mathcal{P}$, we let $\text{di\_shift\_diff}(p, p')$ denote the absolute difference between the discharge shifts $\text{di\_shift}(p)$ and $\text{di\_shift}(p')$. The heterogeneity for the two patients is then calculated by the following formula[2]:

$$\text{het}(p, p') \coloneqq \ln \text{di\_shift\_diff}(p, p')$$

This means that the heterogeneity value $\text{het}(p, p')$ is the natural logarithm of the difference between the $\text{di\_shift}(p)$ and $\text{di\_shift}(p')$. Here, the logarithm is used in order to limit the growth of the values in cases where the planning period is long, where large values would otherwise occur. Note that, since the logarithm is a strictly increasing function, a lower heterogeneity value for two patients signifies more similar discharge dates and, thus, a better fit between the two patients for being assigned to the same room.

The heterogeneity values of all patient pairs are calculated upon initialization of the heuristic and stored in a heterogeneity matrix denoted by HetMatrix. Here, since the heterogeneity values are symmetric (i.e., $\text{het}(p, p') = \text{het}(p', p)$ for all $p, p' \in \mathcal{P}$), it suffices to compute the values above the main diagonal of the matrix in order to improve efficiency. The heterogeneity values of patient pairs are then used whenever the current objective contributions are computed for an assignment that involves assigning a patient to a room to which other patients have already been assigned on the same day. In this case, the appropriately weighted maximum of the heterogeneity values between the new patient and the already assigned patients is added as an additional summand to the current objective contributions.

---

[2]In case that $\text{di\_shift}(p) = \text{di\_shift}(p')$, where the argument of the logarithm equals zero, we set $\text{het}(p, p')$ to zero. In all other cases, the value is strictly positive.

Algorithm 1 IPRNPA heuristic

Input: An instance of the IPRNPA problem
Output: patient-to-nurse assignments x, patient-to-room assignments y
1: $x_{p,n,s} \leftarrow 0$ for all $p \in \mathcal{P}, n \in \mathcal{N}, s \in \mathcal{S}$
2: $y_{p,r,s} \leftarrow 0$ for all $p \in \mathcal{P}, r \in \mathcal{R}, s \in \mathcal{S}^{\text{early}}$
3: HetMatrix $\leftarrow$ calculateHeterogeneityMatrix
4: **for** $s \in \mathcal{S}^{\text{early}}$ in increasing order **do**
5:     $\mathcal{P}(s) \leftarrow \{p \in \mathcal{P} : \text{ad\_shift}(p) \leq s \leq \text{di\_shift}(p)\}$
6:     $\mathcal{N}^{\text{comb}}(s) \leftarrow \mathcal{N}^{\text{early}}(s) \times \mathcal{N}^{\text{late}}(s) \times \mathcal{N}^{\text{night}}(s)$
7:     $\mathcal{R}(s) \leftarrow \mathcal{R}$
8:     ContribTable $\leftarrow$ calculateContributionTable($\mathcal{P}(s), \mathcal{N}^{\text{comb}}, \mathcal{R}(s),$ HetMatrix)
9:     **while** $\mathcal{P}(s) \neq \emptyset$ **do**
10:         $(p, n^{\text{comb}}, r) \leftarrow \text{argmin}(\text{ContribTable})$
11:         $y_{p,r,s} \leftarrow 1$
12:         $(n^{\text{early}}, n^{\text{late}}, n^{\text{night}}) \leftarrow n^{\text{comb}}$
13:         $x_{p,n^{\text{early}},s}, x_{p,n^{\text{late}},s}, x_{p,n^{\text{night}},s} \leftarrow 1$
14:         $\mathcal{P}(s) \leftarrow \mathcal{P}(s) \setminus \{p\}$
15:         **if** room r is fully occupied **then**
16:             $\mathcal{R}(s) \leftarrow \mathcal{R}(s) \setminus \{r\}$
17:         **end if**
18:         ContribTable $\leftarrow$ updateContributionTable(ContribTable, $p, n^{\text{comb}}, r,$ HetMatrix)
19:     **end while**
20: **end for**

---

6 Instance generator and case study

We now present a detailed description of the parameterized instance generator that we developed for creating realistic test instances of the IPRNPA problem. Afterwards, we present the structure and key data of the real-world instances obtained from our partner hospital. The concrete parameter values used for generating the artificial instances as well as the numbers of instances considered will be described later in Section 7.

6.1 Instance generator

In order to generate larger numbers of realistic test instances, we developed a parameterized instance generator for the IPRNPA problem. The source code of the instance generator is publicly available on GitHub at https://github.com/TabeaBrandt/instance_generation_integrated_beds_and_staff_planning.git.

The instance generator offers the possibility to create a specifiable number of test instances based on user-defined parameters. These parameters include the number of instances to be created, as well as the option to specify the length of the planning period in weeks. The remaining input parameters can be divided into two main categories: room-related parameters and nurse-related parameters. In the context of room-related parameters, aspects such as the number of patient rooms, their capacity (single, double, triple, or quadruple rooms), occupancy rate, presence of additional rooms, and possible types of room equipment play a central role. In the context of nurse-related parameters, significant options for fine-tuning are available. This includes setting the maximum desired workloads for nurses, as well as specifying the skill levels to be considered. Concerning the total number of nurses that are available and the nurse roster, the generator operates in two modes: manual and automatic.

In the manual mode, the user specifies the number of nurses explicitly, based on which the generator strives to create a feasible nurse roster using the binary integer program (BIP) outlined in Appendix A. In the automatic mode, on the other hand, the generator increases the number of nurses until a feasible nurse roster can be generated using the BIP.

The output of the generator is a set of instances of the specified cardinality. In each instance, a defined number of rooms is described, which is divided into single, double, triple, and quadruple rooms according to the specified distribution of room sizes. The rooms are assigned randomly selected equipment from the specified types of possible equipment. Additional rooms such as nursing stations or storage rooms can be added, and a minimum of one such room is required to calculate walking distances based on the star-like walking pattern. The weighting factors that determine the importance of the circular versus the star-like walking pattern depend on the type of shift. The circular pattern is favored during early shifts, a more equal split is favored during late shifts, and the star-like pattern is favored during night shifts.

The number of nurses is derived from the number of patient rooms, the distribution of the number of beds per room, and the number of nurses of each skill level that is required per shift, all of which are specified when creating an instance. In instances where three skill levels are specified, we assume that 20% are experienced nurses (skill level 3), 60% are regular nurses (skill level 2) and that 20% are trainees (skill level 1). For instances with two skill levels, we assumed that 80% are skill level 2 and 20% are skill level 1. Additionally, the maximum desired workload associated with each skill level is set to 10 for skill level 1, 12.5 for skill level 2, and 15 for skill level 3. This represents an average nurse-to-patient ratio of 1:4, 1:5, and 1:6, respectively, during each shift.

The patients are generated based on the room configuration and the desired occupancy level. Each patient is assigned a ten-year age group uniformly sampled from 20–30 to 90–100, and, an admission shift, based on the number of rooms and the rooms' capacities. A patient's discharge shift is set as the minimum of the admission shift plus a sampled length of stay (LOS) in days, drawn from a discrete uniform distribution on $\{1, \ldots, 5\}$, or the last shift of the planning period. Gender is currently assigned based on a 50-50 female-male split. The nurse skill level required for each shift of a patient's LOS is assumed to decrease monotonously. The workload generated by a patient p during each shift is based on a gamma distribution with $\alpha = 3$, $\beta = 0.5 + \text{age\_group}(p)/10$, with a minimum of 1, a maximum of 5, and an exponential smoothing parameter of 0.1 that describes the monotonous decrease. The equipment required by the patient is sampled from the types of possible room equipment and assumes monotonously decreasing requirements over the shifts of a patient's LOS.

6.2 Real-world instances

To investigate the potential of our methods using real-world data, we also test them on real-world instances from a Short Stay Unit of our partner hospital. As with our instance generator, these instances are publicly available on GitHub at https://github.com/TabeaBrandt/instance_generation_integrated_beds_and_staff_planning.git.

The considered ward is not restricted to a single medical specialty, but only patients requiring care that can be delivered according to a strict protocol are admitted. Consequently, there are no acute admissions to this ward and admission and discharge dates of patients as well as their care requirements are known in advance. The ward is closed on weekends, so each patient's LOS is at most five days.

During the week, the nursing staff operates in three shifts (early, late and night) to ensure continuous day and night care. Due to the protocol-based care, the nurses who work on the ward do not specialize in a single medical specialty, but require a broad skill set. There are two nurse skill levels: experienced and trainee. One experienced nurse can take care of four to six patients simultaneously during a shift depending on the patient's care requirements and the nurse's experience, while trainee nurses can take care of about two patients in parallel. There is no particular nurse-to-patient ratio during night shifts, but at least two nurses must always be present.

The ward consists of 17 patient rooms of varying sizes: four single rooms, 10 double rooms, two triple rooms, and one quadruple room. Therefore, the ward has a total capacity of 34 beds. Additionally, there is one nursing station where the nurses are usually located when they are not attending to patients. We were provided with the floor plan to estimate the walking distances between the rooms, which we calculated according to the shortest walking path between the centers of each room pair.

In our numerical experiments, we use real admission data and nurse rosters of 40 weeks (about nine months), from before the COVID-19 pandemic. Because the ward closes on the weekend, each week in the data can be considered as a separate instance. We acknowledge that, despite the comprehensive dataset shared by the hospital, certain input data have been omitted due to privacy concerns in order to safeguard individual patient identities. For example, data on the skill levels required for taking care of individual patients and the resulting patient-specific workloads for nurses have not been provided. These missing data have been generated using the corresponding functions of our instance generator presented in Section 6.1 based on realistic parameter values that have been established in cooperation with Amsterdam University Medical Centers (see Section 6.2). The data provided led to infeasible instances for two weeks. In one case, a patient was assigned to a shift for which no nurse was on duty, while in the second instance, a shortage in bed capacity during a particular shift rendered it infeasible to accommodate the required patient load.

7 Experimental results

This section presents our experimental results obtained by testing the MIP as well as the solution methods presented in Section 5 on both artificial instances generated by our instance generator from Section 6.1 and the real-world instances described in Section 6.2.

All computational experiments were performed on a Linux system running Ubuntu 23.04. The hardware includes an AMD EPYC 7542 processor with 32 CPU cores and 64 threads, operating at a base clock speed of 2.9GHz. The system is equipped with 512 GB of memory. The experiments were implemented using Python 3.11 and Gurobi 10.0.1. To solve the completely integrated MIP presented in Section 4 and the two submodels in the sequential solution approach from Section 5.1, we dedicated 8 threads to each instance when solving with Gurobi.

As described in Section 4, the objective function to be minimized consists of a weighted sum of several separate objectives. The specific weights for these objectives were determined based on the existing literature and discussion with our partner hospital. The weights from the existing literature were taken from Demeester et al. [17] and are as follows: the patient transfers objective (1) was assigned a weight of 11, the gender mixing objective (3) a weight of 5, and the equipment violation objective (4) a weight of 5. Furthermore, in consultation with our partner hospital the following weights were established: the patient inconvenience objective (2) and the continuity of care objective (5) were each given a weight of 1, the penalization of skill level requirements and undesired workload distributions objective (6) a weight of 5, the assignment of the minimum number of nurses per room objective (7) a weight of 2, and the walking distances objective (8) a weight of 0.05. When applying our heuristic, we also considered the heterogeneity values of patient pairs as discussed in Section 5.2.2. Here, we assigned a weight of 1 to appropriately incorporate this factor.

7.1 Artificial instances

We established a structured framework involving two distinct scenarios, each subdivided into three specific variations and two different lengths of the planning period. This results in a total of 12 scenario-variation-planning period combinations. For each of these 12 combinations, we generated 10 artificial instances using our instance generator described in Section 6.1.

The two scenarios encompass configurations of 30 and 60 beds. Within each scenario, Variation 1 comprises exclusively double rooms, Variation 2 exclusively triple rooms, and Variation 3 encompasses a diverse mix of room types, including single, double, triple, and quadruple rooms. For the 30 beds scenario, this allocation translates to 3 single rooms, 5 double rooms, 3 triple rooms, and

2 quadruple rooms. These numbers are doubled for the 60 beds scenario. Moreover, our investigation includes two different lengths of the planning period for each scenario-variation combination, spanning either 2 or 4 weeks.

Throughout our comprehensive analysis, we used input parameters for our instance generator that encompass two distinct equipment types, three possible nurse skill levels, the inclusion of a nursing station as an additional room, and a constant occupancy rate of 85%.

In addressing the given scenarios, we applied three distinct solution methods to solve the artificial instances, which we will refer to as Methods 1–3 in the following: the MIP from Section 4 (Method 1), the sequential solution approach from Section 5.1 (Method 2) and the heuristic solution approach from Section 5.2 (Method 3). For Method 1, the running time was limited to a maximum of 3 hours (10800 seconds). In the case of Method 2, the same total running time limit was evenly distributed between the two addressed submodels. Additionally, we configured all MIPs to be terminated when the MIP gap falls below 5%. The results, including objective values and running time data, are presented in Table 1.

The reported results present the objective values as percentages, with the Method 1 objective value serving as the baseline for each instance. These percentages represent the resulting objective proportions, averaged across the 10 instances for each scenario-variation-planning period combination. It is important to emphasize that values below 100% indicate superior performance to Method 1 since the model aims to minimize the objective function. Notably, Method 2 consistently outperforms Method 1 in all scenario-variation-planning period combinations, while Method 3 solution approach surpasses Method 1 in nearly all instances. When comparing Method 2 to Method 3, it becomes clear that Method 2 excels in the 30 beds scenario but lags behind Method 3 for the 60 beds scenario, except for one case. The standard deviation of the objective values for Method 2 and Method 3 is generally low across most scenario-variation-planning period combinations. However, a few instances with higher standard deviations can be attributed to a handful of strongly deviating instances.

When comparing the running times of the three methods, it is evident that Method 1 consistently exceeded the running time limit in all instances, while Method 2 did so in nearly all instances. In contrast, Method 3 displayed remarkable efficiency, requiring, on average, only 27 to 80 seconds for the 30 beds scenario instances and 871 to 1784 seconds for the 60 beds scenarios. This disparity in running times is primarily due to the non-linear growth in nurse combinations for one-day NPAs, as the heuristic assigns three nurses simultaneously. In total, the 30 beds scenario had 21 nurses available to be on duty, while the 60 beds scenario had 31 nurses. Additionally, for each scenario-variation-planning period combination, the running times remained stable with low standard deviation, highlighting the consistency of Method 3's performance. Doubling the planning horizon from 2 to 4 weeks in both scenarios and

Table 1 Artificial instances: Running times and objective values

Scenario	Variation	Planning horizon	Running time (seconds)		Objective value[1]	
			Avg	Stdev	Avg	Stdev
Method 1: MIP						
	Var. 1 (double rooms)	2 weeks	10800	0	100%	0%
		4 weeks	10800	0	100%	0%
30 beds	Var. 2 (triple rooms)	2 weeks	10800	0	100%	0%
		4 weeks	10800	0	100%	0%
	Var. 3 (mixed rooms)	2 weeks	10800	0	100%	0%
		4 weeks	10800	0	100%	0%
	Var. 1 (double rooms)	2 weeks	10800	0	100%	0%
		4 weeks	10800	0	100%	0%
60 beds	Var. 2 (triple rooms)	2 weeks	10800	0	100%	0%
		4 weeks	10800	0	100%	0%
	Var. 3 (mixed rooms)	2 weeks	10800	0	100%	0%
		4 weeks	10800	0	100%	0%
Method 2: Sequential solution approach						
	Var. 1 (double rooms)	2 weeks	10800	0	61%	2%
		4 weeks	10800	0	46%	14%
30 beds	Var. 2 (triple rooms)	2 weeks	7829	970	73%	4%
		4 weeks	10800	0	36%	10%
	Var. 3 (mixed rooms)	2 weeks	10800	0	64%	2%
		4 weeks	10800	0	38%	13%
	Var. 1 (double rooms)	2 weeks	10800	0	25%	1%
		4 weeks	10800	0	25%	3%
60 beds	Var. 2 (triple rooms)	2 weeks	10800	0	30%	23%
		4 weeks	10800	0	24%	2%
	Var. 3 (mixed rooms)	2 weeks	10800	0	24%	1%
		4 weeks	10800	0	24%	1%
Method 3: Heuristic solution approach						
	Var. 1 (double rooms)	2 weeks	27	1	88%	4%
		4 weeks	57	1	65%	19%
30 beds	Var. 2 (triple rooms)	2 weeks	38	1	116%	6%
		4 weeks	80	1	57%	16%
	Var. 3 (mixed rooms)	2 weeks	30	1	98%	5%
		4 weeks	65	2	58%	21%
	Var. 1 (double rooms)	2 weeks	888	46	17%	1%
		4 weeks	1531	62	16%	0%
60 beds	Var. 2 (triple rooms)	2 weeks	1071	46	27%	20%
		4 weeks	1784	66	17%	1%
	Var. 3 (mixed rooms)	2 weeks	871	45	18%	1%
		4 weeks	1515	101	17%	0%

[1] All objective values are provided as proportions (expressed as percentages) of those obtained using Method 1. Values below 100% indicate superior performance to Method 1.

variations resulted in an approximate doubling of the running time, indicating a linear relationship between running time and planning horizon length.

By factoring in the achieved MIP gaps (Table B1 in appendix) alongside the objective values and running times, we gain insights into the difficulty levels of each scenario-variation-planning period combination. Variation 2, which

exclusively considers triple rooms in the 30 beds scenario, emerges as the easiest to solve for both Method 1 and Method 2, while Variations 1 and 3, focusing on double rooms and a mixed set of room types, respectively, seem equally challenging for both methods. Notably, Method 3 exhibits greater efficiency in Variations 1 and 3 compared to Variation 2 when considering running times. MIP gaps exceeding 100% for Method 1 indicate cases where the initial root relaxation of the MIP model could not be solved within the running time limit, resulting in Gurobi finding suboptimal heuristic solutions. In summary, Method 3 consistently outperforms Method 1 and Method 2 across all cases, excelling in both running time and objective value, particularly in complex scenarios where achieving optimality is notably demanding.

Since the considered objective function is a weighted sum of the separate objective functions (1)–(8), we also consider the values of these separate objective functions for a more comprehensive analysis. The detailed outcomes are presented in Table 2.

Table 2 Artificial instances: Separate objective values differentiated by method

	Objective	Method 1 Obj. value[1]		Method 2 Obj. value[1]		Method 3 Obj. value[1]	
		30 beds	60 beds	30 beds	60 beds	30 beds	60 beds
(1)	Transfers	100%	100%	0%	21%	1%	3%
(2)	Inconvenience	100%	100%	29%	41%	64%	56%
(3)	Gender mixing	100%	100%	2%	2%	40%	28%
(4)	Equipment violation	100%	100%	144%	24%	291%	74%
(5)	Continuity of care	100%	100%	103%	109%	98%	100%
(6)	Skill & workload	100%	100%	26%	17%	68%	9%
(7)	Nurses per room	100%	100%	83%	109%	79%	101%
(8)	Walking distances	100%	100%	59%	36%	53%	26%

[1] All objective values are provided as proportions (expressed as percentages) of those obtained using Method 1 (MIP).

Upon examining each objective separately, several noteworthy observations emerge. Method 2 and Method 3 in both scenarios exhibit a notable advantage over Method 1 in strictly minimizing patient transfers (objective (1)). The objectives related to continuity of care (objective (5)) and the number of nurses per room (objective (7)) exhibit similar behavior across Methods 1 to 3. Gender mixing (objective (3)) is strongly avoided in Method 2 and to a moderate extent in Method 3. Interestingly, it is worth noting that skill and workload violations (objective (4)) are more pronounced in Methods 2 and 3 only in the 30 beds scenario, while this issue does not appear as prominently in the 60 beds scenario. These findings underscore each method's distinct strengths and weaknesses in addressing specific objectives.

7.2 Real-world instances

To test all developed methods using real-world data, we sourced data from a Short Stay Unit of our partner hospital as described in Section 6.2. The dataset comprises 40 individual instances, each spanning a planning period of one week.

To address the absence of certain input data needed for the IPRNPA problem, we generated the missing information using the corresponding functions implemented in our instance generator introduced in Section 6.1 based on realistic parameter values established in cooperation with our partner hospital. Specifically, we utilized the functions to generate data on nurse skill level requirements of patients and nurse workloads induced by patients. This approach allowed us to create complete instances for our analyses while at the same time maintaining patient privacy and confidentiality.

When comparing average (minimum, maximum) parameter values across the considered instances, we observe notable variations. There are approximately 17 (min. 13, max. 23) nurses attending to around 62 (min. 45, max. 76) patients. Patient LOS are around 4.1 (min. 3.6, max. 4.8) shifts, resulting in an average occupancy rate of 61% (min. 46%, max. 72%), while each patient causes a workload of about 2.7 (min. 2.2, max. 3.1) during a day shift, i.e., early and late shift.[3] We selected the 20 instances with the highest occupancy rates from the available dataset for further analysis, taking care to exclude holiday times and instances exhibiting unnatural utilization levels. This careful selection process ensures that the chosen instances accurately represent the most demanding and meaningful scenarios for in-depth examination.

In tackling the real-world instances, we again applied all three solution methods, which are again referred to as Methods 1–3 as in Section 7.1. The termination criteria based on running time limits and MIP gap were also adopted from Section 7.1. The results, encompassing objective values and running time data, are presented in Table 3.

The results in real-world instances parallel those in the artificial ones (Section 7.1) concerning running time limits, where Method 1 consistently reaches the time limit in all cases. In Method 2, the running time limit is only exceeded in the PRA subproblem of instance 20. In stark contrast, Method 3 demonstrates remarkable efficiency, with an average running time of just 10 seconds, vastly outperforming Method 2, which has an average running time of 2535 seconds.

Regarding objective values, Method 1 consistently outperforms Methods 2 and 3 across all instances, with average objective values of 107% and 137%, respectively. The real-world instances are less complex than the artificial 30 beds scenario, evidenced by an average MIP gap of 43% (Table B2 in the appendix). Method 2 optimally solves all PRA subproblems except for

---

[3]The presented LOS and workload numbers are averages over all patients of an instance.

Table 3 Real-world instances: Running times and objective values

Instance #	Method 1: MIP		Method 2: Sequ. sol. app.		Method 3: Heur. sol. app.	
	Running time	Obj. value[1]	Running time	Obj. value[1]	Running time	Obj. value[1]
1	10800	100%	4476	103%	10	124%
2	10800	100%	4346	110%	9	129%
3	10800	100%	3300	106%	9	146%
4	10800	100%	3142	100%	10	134%
5	10800	100%	4593	107%	8	127%
6	10800	100%	3397	102%	10	127%
7	10800	100%	1479	106%	11	138%
8	10800	100%	2792	102%	12	128%
9	10800	100%	1997	112%	10	147%
10	10800	100%	1894	109%	13	140%
11	10800	100%	1746	109%	10	137%
12	10800	100%	2358	110%	11	164%
13	10800	100%	294	111%	9	148%
14	10800	100%	113	123%	6	142%
15	10800	100%	656	109%	8	121%
16	10800	100%	4551	104%	8	144%
17	10800	100%	1920	103%	9	122%
18	10800	100%	403	107%	10	125%
19	10800	100%	1233	109%	11	139%
20	10800	100%	6006	101%	13	161%
Total[2]	10800	100%	2535	107%	10	137%

[1] All objective values are provided as proportions (expressed as percentages) of those obtained using Method 1 (MIP).

[2] Average across all instances

instance 20, where a 100% MIP gap occurs because the lower bound equaled 0. In NPA subproblems, the MIP gap drops below 5% during runtime, prompting the optimization to halt accordingly.

In summary, the evaluation of three methods for healthcare management optimization reveals distinct trade-offs. Method 1, while suitable for small-scale settings such as those provided by our partner hospital, suffers from prohibitively high running time costs, often taking thousands of seconds to achieve results comparable to the heuristic, which only needs a few seconds. For medium-sized problems like the 30 beds scenario, Method 2 delivers superior results compared to Method 3 at the expense of significant running time. In contrast, Method 3 consistently stands out regarding running time efficiency across problem sizes and showcases excellent objective values, particularly in larger scenarios like the 60 beds scenario. Hence, Method 3 emerges as the preferred choice for optimizing healthcare management, striking a favourable balance between computational efficiency and solution quality, especially in more extensive and complex healthcare settings.

8 Conclusion and outlook

Motivated by important interactions of PRA and NPA decisions in hospital wards, this paper explicitly considers both types of assignment decisions in one integrated optimization problem for the first time. We introduce the IPRNPA problem and provide a formal mathematical description as a mixed integer program. Since the PRA problem and the NPA problem are already NP-hard

and very difficult to solve for realistic instance sizes, it is not surprising that the integrated problem is computationally challenging and cannot be solved to (near) optimality in reasonable time using the completely integrated MIP model. Therefore, we also present an efficient heuristic for the integrated problem that is able to compute high-quality solutions quickly on both artificially generated instances and real-world instances obtained from our partner hospital. The heuristic solution approach highlighted its superiority in running time efficiency across various problem sizes, particularly excelling in larger and more complex scenarios. While the integrated MIP struggles with high running time costs the heuristic solution approach strikes a favorable balance between computational efficiency and solution quality. Its superiority in objective values is notably pronounced in medium and large-scale scenarios, making it the preferred choice for optimizing real-world healthcare settings.

We also devise a parameterized instance generator for the problem. This generator is made freely available to other researchers to foster additional investigations on the IPRNPA problem, which we believe represents a challenging and at the same time practically relevant problem to be further investigated in the future. For instance, while our heuristic solution method allows easy adaptions to some dynamic versions of the problem, explicitly investigating different dynamic extensions with increasing degrees of data uncertainty (e.g., no-shows of patients or unexpected changes of patients' care requirements and/or LOS after admission) might represent a fruitful direction for future research. Moreover, while the nurse roster for the planning period is considered as an input of the problem in this paper, integrating rostering decisions into the problem formulation might represent an interesting extension.

Statements and declarations

Data availability

The instance generator and the instances used for the computational experiments are available on GitHub, https://github.com/TabeaBrandt/instance_generation_integrated_beds_and_staff_planning.git.

Funding

Conflicts of interest

The authors have no conflicts of interest to declare that are relevant to the content of this article.

References

[1] Acar I, Butt SE (2016) Modeling nurse-patient assignments considering patient acuity and travel distance metrics. Journal of Biomedical Informatics 64:192–206

[2] Aiken LH, Clarke SP, Sloane DM, et al (2002) Hospital nurse staffing and patient mortality, nurse burnout, and job dissatisfaction. JAMA 288(16):1987–1993

[3] Bai J, Fügener A, Schoenfelder J, et al (2018) Operations research in intensive care unit management: A literature review. Health Care Management Science 21(1):1–24

[4] Bastos LS, Marchesi JF, Hamacher S, et al (2019) A mixed integer programming approach to the patient admission scheduling problem. European Journal of Operational Research 273:831–840

[5] Benazzouz T, Echchatbi A, Bellabdaoui A (2015) A literature review on the nurses' planning problems. International Journal of Mathematics and Computational Science 1(5):268–274

[6] Bilgin B, Demeester P, Misir M, et al (2012) One hyper-heuristic approach to two timetabling problems in health care. Journal of Heuristics 18:401–434

[7] Bodenheimer TS, Smith MD (2013) Primary care: Proposed solutions to the physician shortage without training more physicians. Health Affairs 32(11):1881–1886

[8] Borchani R, Masmoudi M, Jarboui B, et al (2021) Heuristics-based on the hungarian method for the patient admission scheduling problem. In: [37], chap 2, p 33–62

[9] Bouras A, Masmoudi M, Saadani NEH, et al (2021) Multi-stage appointment scheduling for outpatient chemotherapy unit: a case study. RAIRO Operations Research 55(2):589–610

[10] Butt SE, Fredericks TK, Kumar AR, et al (2004) An evaluation of physiological work demands on registered nurses over a 12-hour shift. In: Proceedings of the XVIII Annual International Occupational Ergonomics and Safety Conference (ISOES)

[11] Ceschia S, Schaerf A (2011) Local search and lower bounds for the patient admission scheduling problem. Computers & Operations Research 38(10):1452–1463

[12] Ceschia S, Schaerf A (2012) Modeling and solving the dynamic patient admission scheduling problem under uncertainty. Artificial Intelligence in Medicine 56(3):199–205

[13] Ceschia S, Schaerf A (2016) Dynamic patient admission scheduling with operating room constraints, flexible horizons, and patient delays. Journal of Scheduling 19:377–389

[14] Clark A, Moule P, Topping A, et al (2015) Rescheduling nursing shifts: Scoping the challenge and examining the potential of mathematical model based tools. Journal of Nursing Management 23(4):411–420

[15] Cohen B, Hyman S, Rosenberg L, et al (2012) Frequency of patient contact with health care personnel and visitors: Implications for infection prevention. The Joint Commission Journal on Quality and Patient Safety 38(12):560–565

[16] Dancer SJ (2009) The role of environmental cleaning in the control of hospital-acquired infection. Journal of Hospital Infection 73(4):378–385

[17] Demeester P, Souffriau W, De Causmaecker P, et al (2010) A hybrid tabu search algorithm for automatically assigning patients to beds. Artificial Intelligence in Medicine 48(1):61–70

[18] Drupsteen J, van der Vaart T, van Donk DP (2013) Integrative practices in hospitals and their impact on patient flow. International Journal of Operations & Production Management 33(7):912–933

[19] Erhard M, Schoenfelder J, Fügener A, et al (2018) State of the art in physician scheduling. European Journal of Operational Research, 265(1):1–18

[20] Eveillard M, Hitoto H, Raymond F, et al (2009) Measurement and interpretation of hand hygiene compliance rates: importance of monitoring entire care episodes. Journal of Hospital Infection 72(3):211–217

[21] Ficker AMC, Spieksma FCR, Woeginger GJ (2021) The transportation problem with conflicts. Annals of Operations Research 298(1):207–227

[22] Guerriero F, Guido R (2011) Operational research in the management of the operating theatre: A survey. Health Care Management Science 14(1):89–114

[23] Guido R, Groccia MC, Conforti D (2018) An efficient matheuristic for offline patient-to-bed assignment problems. European Journal of Operational Research 268(2):486–503

[24] Halwani M, Solaymani-Dodaran M, Grundmann H, et al (2006) Cross-transmission of nosocomial pathogens in an adult intensive care unit: incidence and risk factors. Journal of Hospital Infection 63(1):39–46

[25] Haspeslagh S, De Causmaecker P, Schaerf A, et al (2014) The first international nurse rostering competition 2010. Annals of Operations Research 218:221–236

[26] He L, Madathil SC, Oberoi A, et al (2019) A systematic review of research design and modeling techniques in inpatient bed management. Computers & Industrial Engineering 127:451–466

[27] Hesaraki AF, Dellaert NP, de Kok T (2019) Generating outpatient chemotherapy appointment templates with balanced flowtime and makespan. European Journal of Operational Research 275:304–318

[28] Hesaraki AF, Dellaert NP, de Kok T (2020) Integrating nurse assignment in outpatient chemotherapy appointment scheduling. OR Spectrum 42:935–963

[29] Heshmat M, Nakata K, Eltawil A (2018) Solving the patient appointment scheduling problem in outpatient chemotherapy clinics using clustering and mathematical programming. Computers & Industrial Engineering 124:347–358

[30] Hulshof PJH, Kortbeek N, Boucherie RJ, et al (2012) 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

[31] Jha RK, Sahay BS, Charan P (2016) Healthcare operations management: a structured literature review. Decision 43(3):259–279

[32] Jun JB, Jacobson SH, Swisher JR (1999) Application of discrete-event simulation in health care clinics: A survey. Journal of the Operational Research Society 50(2):109–123

[33] Ku WY, Pinheiro T, Beck JC (2014) CIP and MIQP models for the load balancing nurse-to-patient assignment problem. In: Proceedings of the 20th International Conference on Principles and Practice of Constraint Programming (CP), pp 424–439

[34] Liang B, Turkcan A (2016) Acuity-based nurse assignment and patient scheduling in oncology clinics. Health Care Management Science 19:207–226

[35] Lusby R, Schwierz M, Range T, et al (2016) An adaptive large neighborhood search procedure applied to the dynamic patient admission scheduling problem. Artificial Intelligence in Medicine 74:21–31

[36] Marzouk M, Kamoun H (2020) Nurse to patient assignment through an analogy with the bin packing problem: Case of a Tunisian hospital. Journal of the Operational Research Society 72(8):1808–1821

[37] Masmoudi M, Jarboui B, Siarry P (eds) (2021) Operations Research and Simulation in Healthcare. Springer

[38] Mullinax C, Lawley M (2002) Assigning patients to nurses in neonatal intensive care. Journal of the Operational Research Society 53(1):25–35

[39] Pesant G (2016) Balancing nursing workload by constraint programming. In: Proceedings of the 13th International Conference on Integration of Constraint Programming, Artificial Intelligence, and Operations Research (CPAIOR), pp 294–302

[40] Punnakitikashem P, Rosenberger JM, Buckley-Behan DF (2008) Stochastic programming for nurse assignment. Computational Optimization and Applications 40:321–349

[41] Punnakitikashem P, Rosenberger JM, Buckley-Behan DF (2013) A stochastic programming approach for integrated nurse staffing and assignment. IEEE Transactions 45(10):1059–1076

[42] Rachuba S, Reuter-Oppermann M, Thielen C (2023) Integrated planning in hospitals: A review. https://arxiv.org/abs/2307.05258

[43] Rais A, Viana A (2011) Operations research in healthcare: A survey. International Transactions in Operational Research 18(1):1–31

[44] Range T, Lusby R, Larsen J (2014) A column generation approach for solving the patient admission scheduling problem. European Journal of Operational Research 235(1):252–264

[45] Schäfer F, Walther M, Hübner A (2017) Patient-bed allocation in large hospitals. In: Proceedings of the 3rd International Conference on Health Care Systems Engineering (HCSE), pp 299–300

[46] Schäfer F, Walther M, Hübner A, et al (2019) Operational patient-bed assignment problem in large hospital settings including overflow and uncertainty management. Flexible Services and Manufacturing Journal 31:1012–1041

[47] Schäfer F, Walther M, Grimm DG, et al (2021) Combining machine learning and optimization for the operational patient-bed assignment problem. https://doi.org/10.2139/ssrn.3919282

[48] Schaus P, Régin JC (2014) Bound-consistent spread constraint - application to load balancing in nurse-to-patient assignments. Euro Journal on Computational Optimization 2(3):123–146

[49] Schmidt R, Geisler S, Spreckelsen C (2013) Decision support for hospital bed management using adaptable individual length of stay estimations and shared resources. BMC Medical Informatics and Decision Making 13(3)

[50] Sir MY, Dundar B, Barker Steege LM, et al (2015) Nurse–patient assignment models considering patient acuity metrics and nurses' perceived workload. Journal of Biomedical Informatics 55:237–248

[51] Sundaramoorthi D, Chen VCP, Rosenberger JM, et al (2009) A data-integrated simulation model to evaluate nurse–patient assignments. Health Care Management Science 12:252–268

[52] Thielen C (2018) Duty rostering for physicians at a department of orthopedics and trauma surgery. Operations Research for Health Care 19:80–91

[53] Thomas BG, Bollapragada S, Akbay K, et al (2013) Automated bed assignments in a complex and dynamic hospital environment. Interfaces 43:435–448

[54] Thuran A, Bilgen B (2017) Mixed integer programming based heuristics for the patient admission scheduling problem. Computers & Operations Research 80:38–49

[55] Van Riet C, Demeulemeester E (2015) Trade-offs in operating room planning for electives and emergencies: A review. Operations Research for Health Care 7:52–69

[56] Vanberkel PT, Boucherie RJ, Hans EW, et al (2010) A survey of health care models that encompass multiple departments. International Journal of Health Management and Information 1(1):37–69

[57] Vancroonenburg W, Croce FD, Goossens D, et al (2014) The Red–Blue transportation problem. European Journal of Operational Research 237(3):814–823

[58] Vancroonenburg W, De Causmaecker P, Vanden Berghe G (2016) A study of decision support models for online patient-to-room assignment planning. Annals of Operations Research 239:253–271

[59] Zhu YH, Toffolo TAM, Vancroonenburg W, et al (2019) Compatibility of short and long term objectives for dynamic patient admission scheduling. Computers & Operations Research 104:98–112

Appendix A Nurse rostering formulation

We use a simple binary integer programming formulation to generate the nurse rosters that are part of our random instances. The formulation is based on the description presented in the first International Nurse Rostering Competition (INRC) 2010 [25]. As the focus of this work is not on nurse rostering, we use a simple but fast formulation instead of a perfectly detailed one. This formulation includes all constraints that are relevant concerning the use of a nurse roster as an input of the IPRNPA. However, it is very easy to substitute the used nurse rostering formulation in our code.

Similar to the INRC, we determine the roster for the planning period considering one ward. We use a subset of constraints of the INRC in order to compute a simple, yet still realistic nurse roster. These constraints include that the number of required nurses per shift must be met, not more than a given maximum allowed number of shifts can be assigned to any single nurse during the planning period, and that minimum rest times for nurses between shifts are respected.

In addition to some of the notation and parameters introduced in Section 4, we use the following parameters and decision variables:

Parameters:

skill_nurses(s, l) number of nurses with at least skill level $l \in \mathcal{L}$ required during shift $s \in \mathcal{S}$. The sum of this number over all skill levels defines the minimum number of nurses needed per shift in total.

max_shifts maximum allowed number of shifts per nurse within the time horizon.

Decision variables:

assign$_{n,s}$ binary variable indicating whether nurse $n \in \mathcal{N}$ is assigned to shift $s \in \mathcal{S}$

The focus of our formulation is on the generation of a feasible roster to be used as an input for the IPRNPA. Therefore, we consider the minimization of the total number of assigned nurses as our objective in order to prevent unnecessary assignments that are not required in order to fulfill the considered constraints:

Assignment objective

(28) Minimization of the number of assigned nurses:

$$\min \sum_{s \in \mathcal{S}, n \in \mathcal{N}} \text{assign}_{n,s}$$

The following constraints must be met by a feasible nurse roster:

(29) Each nurse can work at most one shift per day:

$$\text{assign}_{n,s} + \text{assign}_{n,s+1} + \text{assign}_{n,s+2} \leq 1 \quad \forall n \in \mathcal{N}, s \in \mathcal{S}^{\text{early}}$$

(30) For each skill level $l \in \mathcal{L}$, at least skill_nurses(s, l) nurses with at least skill level l must be assigned during each shift $s \in \mathcal{S}$:

$$\sum_{n \in \mathcal{N}: \text{skill\_level}(n) \geq l} \text{assign}_{n,s} \geq \text{skill\_nurses}(l, s) \quad \forall l \in \mathcal{L}, s \in \mathcal{S}$$

(31) The minimum total number of nurses must be assigned during each shift $s \in \mathcal{S}$:

$$\sum_{n \in \mathcal{N}} \text{assign}_{n,s} \geq \sum_{l \in \mathcal{L}} \text{skill\_nurses}(l, s) \quad \forall s \in \mathcal{S}$$

(32) No nurse $n \in \mathcal{N}$ can be assigned to more than max_shifts many shifts during the planning period:

$$\sum_{s \in \mathcal{S}} \text{assign}_{n,s} \leq \text{max\_shifts} \quad \forall n \in \mathcal{N}$$

(33) On the day after a night shift, a nurse can only have another night shift (or the day off):

$$\text{assign}_{n,s} + \text{assign}_{n,s+1} + \text{assign}_{n,s+2} \leq 1 \quad \forall n \in \mathcal{N}, s \in \mathcal{S}^{\text{night}}$$

(34) On the day after a late shift, a nurse cannot have a morning shift:

$$\text{assign}_{n,s} + \text{assign}_{n,s+2} \leq 1 \quad \forall n \in \mathcal{N}, s \in \mathcal{S}^{\text{late}}$$

Appendix B   MIP Gap evaluation on artificial and real-world instances

Table B1 Artificial instances: Average MIP gap

Scenario	Variation	Planning horizon	Method 1 MIP gap[1] IPRNPA	Method 2 MIP gap[1]	
				PRA	NPA
30 beds	Var. 1 (double rooms)	2 weeks	70.0%	58.0%	7.6%
		4 weeks	80.7%	63.4%	10.9%
	Var. 2 (triple rooms)	2 weeks	57.6%	82.7%	5.0%
		4 weeks	80.8%	79.7%	6.8%
	Var. 3 (mixed rooms)	2 weeks	65.3%	55.0%	6.1%
		4 weeks	81.9%	68.9%	8.8%
60 beds	Var. 1 (double rooms)	2 weeks	110.9%	75.2%	60.5%
		4 weeks	111.1%	89.3%	62.3%
	Var. 2 (triple rooms)	2 weeks	98.0%	94.6%	51.1%
		4 weeks	109.5%	97.0%	62.1%
	Var. 3 (mixed rooms)	2 weeks	110.0%	74.4%	58.2%
		4 weeks	110.3%	90.5%	61.4%

[1] The MIP gap is evaluated after either a running time limit of 10,800 seconds or 5,400 seconds for each subproblem or when the MIP gap reaches a value less than 5%. The MIP gap is shown summarized as an average value.

Table B2 Real-World instances: MIP gap

Instance #	Method 1 MIP gap[1] IPRNPA	Method 2 MIP gap[1] PRA	NPA
1	53%	0%	5%
2	54%	0%	5%
3	45%	0%	5%
4	48%	0%	5%
5	51%	0%	5%
6	46%	0%	5%
7	45%	0%	5%
8	48%	0%	5%
9	35%	0%	5%
10	38%	0%	5%
11	40%	0%	5%
12	47%	0%	5%
13	33%	0%	5%
14	23%	0%	5%
15	48%	0%	5%
16	53%	0%	5%
17	44%	0%	5%
18	35%	0%	5%
19	30%	0%	5%
20	38%	100%	5%
Total[2]	43%	5%	5%

[1] The MIP gap is evaluated after either a running time limit of 10,800 seconds or 5,400 seconds for each subproblem or when the MIP gap reaches a value less than 5%.

[2] Average across all instances