One transfer per patient suffices: Structural insights about patient-to-room assignment

Christina Büsing[a,1,2,*], Sigrid Knust[b,*], Tabea Krabs[a,1,*]

[a] *Lehrstuhl II für Mathematik, RWTH Aachen University, Pontdriesch 10-12, 52062 Aachen, Germany, {buesing,krabs}@math2.rwth-aachen.de*

[b] *Institute of Computer Science, Osnabrück University, Wachsbleiche 27, 49090 Osnabrück, Germany, sigrid@informatik.uni-osnabrueck.de*

Abstract

Assigning patients to rooms is a fundamental task in hospitals and, especially, within wards. For this so-called patient-to-room assignment problem (PRA) many heuristics have been proposed with a large variety of different practical constraints. However, a thorough investigation of the problem's structure itself has been neglected so far. Therefore, in this paper, we present insights about the basic, underlying combinatorial problem of PRA. At first we consider the problem with patient-room assignment restrictions, where not every patient can be assigned to any room. In this case, the problem is strongly $\mathcal{NP}$-complete even if all rooms are double bedrooms. This also settles an open question concerning the red-blue transportation problem. Afterwards, we study the problem without such restrictions and focus on minimizing the number of patients who have to change rooms during their stay. Particularly, we prove that in the case of double bedrooms, each patient has to be transferred at most once.

Keywords: patient-to-room assignment, complexity, red-blue transportation
2010 MSC: 90C27,68Q17,68Q25

---

1. Introduction

Hospital beds are next to medical trained staff one of the most important resources in hospitals [21]. Whether or not a patient can be treated in an hospital often depends on the existence of a free bed. Assigning patients to rooms as efficiently as possible is therefore very important in a well-functioning health care system [14].

There are two ways patients enter a hospital for hospitalization: as emergency patients or as elective patients. The allocation of emergency patients to beds is organized centrally for the complete hospital, depending on the allocation of present patients, planned elective patients and the diagnosis made in the emergency room. In contrast to the management of emergency patients, the allocation of elective patients is organized by each ward independently: Hospitals are traditionally divided into several, mostly independent, clinical departments [17, 28]. Although they may share medical equipment or operating rooms, the appointment scheduling is usually organized by each department individually [15]. When an appointment for a patient

* Corresponding author

1 This work was supported by the Freigeist-Fellowship of the Volkswagen Stiftung and by the German research council (DFG) Research Training Group 2236 UnRAVeL.

2 This work was partially supported by the German Federal Ministry of Education and Research (grant no. 05M16PAA) within the project "HealthFaCT - Health: Facility Location, Covering and Transport".

is fixed and hospitalization is necessary, the patient is directly assigned to a specific ward. Afterwards, the actual assignment of these patients to rooms is performed on ward level by experienced nurses or the ward's head [7, 23, 22, 25].

Therefore, we consider the planning of elective patients to be independent of the accommodation of emergency patients. This is impractical for specialities with a high rate of emergency patients, however, there do exist specialities with nearly 100 percent elective patients. Thus, this assumption is not only helpful to gain first structural insights but also has a direct practical application. In the following, we focus on the problem of assigning elective patients only.

Apart from ensuring that all patients always have a bed during their stay, it is important to avoid reassigning patients to different rooms. These so-called *patient transfers* always mean additional work for the medical staff and inconvenience for the patient.

Formally, in the *patient-to-room assignment problem* (PRA), a set $\mathcal{P} = \{1, \ldots, P\}$ of patients, divided into a set $\mathcal{P}^f$ of female and a set $\mathcal{P}^m$ of male patients, is given. Furthermore, we have a set of rooms $\mathcal{R} = \{1, \ldots, R\}$, where room r has capacity C_r, i.e., r contains C_r beds. We denote by T the length of the planning horizon and by $\mathcal{T} = \{1, \ldots, T\}$ the set of all time periods. As discharge and admittance of patients happens at some time during the day, we avoid the use of the term "day". For each patient $p \in \mathcal{P}$, an arrival period $a_p \in \mathcal{T}$, and a departure period $b_p \in \mathcal{T}$ are specified. Additionally, a bipartite graph $G = (\mathcal{P} \cup \mathcal{R}, E)$ models patient-room assignment restrictions by indicating which patient can be assigned to which room. We have an edge $\{p, r\} \in E$ if and only if patient p can be assigned to room r. The goal of PRA is to assign patients to rooms so that

- each patient $p \in \mathcal{P}$ is assigned to a feasible room $r \in \mathcal{R}$ with $\{p, r\} \in E$ in each time period $t \in \{a_p, \ldots, b_p\}$,
- in each time period, no more than C_r patients are assigned to every room,
- no gender-mixed rooms occur.

Depending on the data, it may be possible that no feasible solution respecting all these constraints exists. In order to have more flexibility in the assignments, we allow the option of patient transfers, i.e., patients may be moved among rooms during their stay. However, such transfers are not desirable, for staff as well as patients. Therefore, we consider the two objectives of minimizing

- f^{Σ}: the total number of patient transfers,
- $f^{\max}$: the maximum number of transfers per patient.

As we experience that more and more hospitals switch to an all-double-bedrooms policy, we focus on PRA in wards where all rooms are double bedrooms, i.e., we assume $C_r = 2$ if not stated otherwise.

In literature, the optimization problem of assigning patients to beds has only recently gained attention. The commonly used problem definition introduced by Demeester et al. is based on the situation in Belgium hospitals [6]. They focus on identifying the most suitable room for each patient using a rating that considers different room equipment, ward specialisation as well as patient needs and wishes.

Many aspects of hospital-bed management have already been investigated using a variety of methodologies such as capacity dimensioning or allocation [15]. Already in 1995, Roth and van Dierdonck suggested a centralised bed management [19] and Huang introduced a model for computing needed capacity of emergency beds under consideration of week-day dependent demand fluctuation [13]. Simulation models for evaluating different planning strategies have been proposed in [10, 16, 11, 12, 5].

However, the operational task of assigning patients to beds has been considered and formalized for the first time in 2010 by Demeester et al. [6]. To solve this basic patient-to-room assignment, a tabu search algorithm [6], a multi-neighborhood local search procedure [2], a column generation approach [18], a population-based metaheuristic [9], and a matheuristic [8] have been proposed.

Since 2010, the basic problem definition has been extended to a dynamic context [3], to include uncertainties in the patients' length-of-stay [21], or to include capacities of other hospital resources [27, 4]. A detailed comparison of different versions can be found in [8]. However, the patient-to-room assignment as it was proposed by Demeester et al. and also all of its published versions tackle PRA with non-equal rooms [6]. This, in our experience, may fit for managing the allocation of emergency patients, but is a too general approach for PRA on ward level. Schäfer et al. made a similar observation, however, the decision support system they developed is precisely tailored to the needs of their cooperating hospital and no general approach is presented [20]. Furthermore, in all published literature on patient-to-room assignment no structural examination of the problem is presented.

In this paper, we start with a detailed analysis of complexity and prove that single-day PRA is $\mathcal{NP}$-complete if all rooms have capacity two. This also extends a result for a related combinatorial problem, the so-called red-blue transportation problem which corresponds to single-day patient-to-room assignment [26]. Vancroonenburg et al. proved that single-day patient-to-room assignment is $\mathcal{NP}$-complete if all rooms have capacity three [26].

We further prove that $f^{\max}$ and f^{Σ} can be optimized simultaneously without trade-offs between the wishes of patients and nurses being necessary. More precisely, we prove that, in the case of double-bedrooms, there is always an optimal solution with respect to f^{Σ} for which $f^{\max} \leq 1$ holds. This justifies the restriction of the search space in the heuristic proposed by Ceschia and Schaerf to solutions in which every patient is transfered at most once [2]. Furthermore, we provide upper bounds on f^{Σ} which are helpful for designing efficient algorithms as well as in strategic and tactical decision-making regarding staff capacity and staff scheduling.

Although we focus on a problem setting in health care, the underlying combinatorial optimization problem of the patient-to-room assignment has also other applications, e.g., in container storages, ships or warehouses. Here, items have to be assigned to positions in stacks of limited height, corresponding to rooms in hospitals with limited capacity. The items have different attributes as a weight class, their destination or retrieval time (cf. e.g., [24]). If only items having the same attribute may be stored together, due to space restrictions transfers of items between different stacks may be necessary but should be avoided.

The remainder of this paper is organized as follows. In Section 2, we discuss the case with patient-room assignment restrictions and prove that in this situation PRA is already $\mathcal{NP}$-complete for only double bed rooms, which strengthens the complexity result of [26]. In the remainder of the paper, we consider PRA for elective patients without patient-room restrictions, focusing on avoiding gender-mixed rooms. In Section 3, we prove that there is always an optimal solution with respect to f^{Σ} where every patient has to be transferred at most once. In Section 4 we show that in an optimal solution, there is no need to transfer a patient who arrives in the first time period. In Section 5 we use the previous results to derive upper bounds on f^{Σ}.

2. NP-completeness for arbitrary patient-room assignment restrictions

In this section, we show that the decision version of PRA is strongly $\mathcal{NP}$-complete if patient-room assignment restrictions are given. For this purpose, we consider single-day PRA, (i.e., $T = 1$) which corresponds to the red-blue transportation problem [26]. We use a reduction from the satisfiability problem variant (3,B2)-SAT, which is known to be strongly $\mathcal{NP}$-complete [1].

Theorem 1. *Single-day PRA is strongly $\mathcal{NP}$-complete, even for room capacities $C_r = 2$.*

Proof. An (3,B2)-SAT ψ instance is a 3SAT instance in which every literal is present exactly two times. Like in 3SAT, the task is to decide whether or not there exists a satisfying truth assignment for ψ. Remark that in every (3,B2)-SAT instance, the number of variables is a multiple of three.

Let ψ be a (3,B2)-SAT instance, $\mathcal{X} = \{x_1, \ldots, x_{|\mathcal{X}|}\}$ the set of variables and $\mathcal{C}$ the set of clauses. Then, we construct a corresponding single-day PRA instance as follows. We introduce four female patients for every variable $x \in \mathcal{X}$, i.e.,

$$\mathcal{P}^f = \{v_i^x \mid x \in \mathcal{X},\ i \in \{1, \ldots, 4\}\}$$

that we use to model the variables' interpretation. We introduce three male patients for every clause $c \in \mathcal{C}$, two male patients for every variable, and $\frac{2}{3}|\mathcal{X}|$ additional male patients i.e., $\mathcal{P}^m = \mathcal{P}_1^m \cup \mathcal{P}_2^m$ with

$$\mathcal{P}_1^m = \{u_i^c \mid c \in \mathcal{C},\ i \in \{1, \ldots, 3\}\} \quad \text{and}$$
$$\mathcal{P}_2^m = \left\{d_i^x \;\middle|\; x \in \mathcal{X},\ i \in \{1, 2\}\right\} \cup \left\{e_i \;\middle|\; i \in \left\{1, \ldots, \frac{2}{3}|\mathcal{X}|\right\}\right\}.$$

For every clause c, the male patients u_1^c, u_2^c, and u_3^c represent the literals in c. The male patients d_i^x and e_i are auxiliary nodes that are needed to ensure that a feasible assignment of patients to rooms always exists if the formula is satisfiable. Further, we have four rooms for every variable and one room for every clause, i.e., $\mathcal{R} = \mathcal{R}^1 \cup \mathcal{R}^2$ with

$$\mathcal{R}^1 = \{r_i^x \mid x \in \mathcal{X}, i \in \{1, \ldots, 4\}\} \quad \text{and} \quad \mathcal{R}^2 = \{C^c \mid c \in \mathcal{C}\}.$$

The rooms in $\mathcal{R}^1$ are used together with all female patients to model the variables' interpretation. The rooms in $\mathcal{R}^2$ represent the clauses and model the formula's structure together with the male patients in $\mathcal{P}_1^m$. Note that $|\mathcal{P}^f| + |\mathcal{P}_1^m| + |\mathcal{P}_2^m| = \frac{32}{3}|\mathcal{X}| = 2(|\mathcal{R}^1| + |\mathcal{R}^2|)$ holds, i.e., in every feasible assignment all rooms must be full. We define the set of edges E for the graph $G = (\mathcal{P}^f \cup \mathcal{P}^m \cup \mathcal{R}, E)$ by $E = E_1 \cup E_2 \cup E_3$ as follows:

E_1 determines the feasible rooms for female patients, so that only two possible assignments exist, which models the variable's interpretation:

$$E_1 = \left\{\{v_i^x, r_i^x\}, \{v_i^x, r_{(i-1) \bmod 4}^x\} \mid x \in \mathcal{X},\ i \in \{1, \ldots, 4\}\right\}.$$

Note that here we use $\{1, 2, 3, 4\}$ as representatives for the residue classes of the mod-operator. E_2 determines the feasible rooms for male patients in $\mathcal{P}_1^m$ and we use this to model the structure of ψ. For every clause $c \in \mathcal{C}$, the patients u_1^c, u_2^c, and u_3^c can either be assigned to the corresponding clause room C^c or to rooms in $\mathcal{R}^1$ that model the corresponding literal's interpretation. These two possibilities correspond to whether the corresponding literal is satisfied or not. We therefore

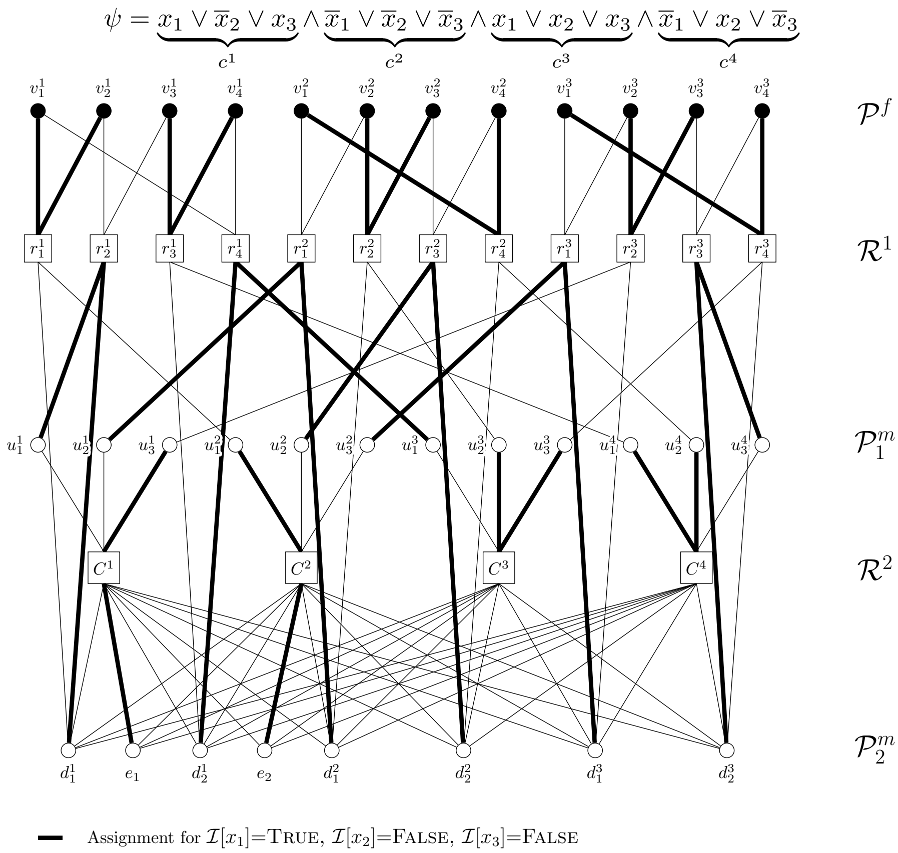

Figure 1: Graph G corresponding to ψ

define

$$E_2 = \big\{\{u_i^c, C^c\} \mid c \in \mathcal{C}, i \in \{1,2,3\}\big\}$$
$$\cup \left\{ \{u_i^c, r_1^x\} \;\middle|\; \begin{array}{l} x \in \mathcal{X} \text{ and literal number } i \text{ in clause} \\ c \in \mathcal{C} \text{ is the first occurence of } \overline{x} \end{array} \right\}$$
$$\cup \left\{ \{u_i^c, r_2^x\} \;\middle|\; \begin{array}{l} x \in \mathcal{X} \text{ and literal number } i \text{ in clause} \\ c \in \mathcal{C} \text{ is the first occurence of } x \end{array} \right\}$$
$$\cup \left\{ \{u_i^c, r_3^x\} \;\middle|\; \begin{array}{l} x \in \mathcal{X} \text{ and literal number } i \text{ in clause} \\ c \in \mathcal{C} \text{ is the second occurence of } \overline{x} \end{array} \right\}$$
$$\cup \left\{ \{u_i^c, r_4^x\} \;\middle|\; \begin{array}{l} x \in \mathcal{X} \text{ and literal number } i \text{ in clause} \\ c \in \mathcal{C} \text{ is the second occurence of } x \end{array} \right\}.$$

Remember that we use the auxiliary patients $\mathcal{P}_2^m$ and rooms $\mathcal{R}^2$ to ensure the feasibility of the patient-to-room assignment. We thus define the feasible rooms for patients in $\mathcal{P}_2^m$ as follows: All patients in $\mathcal{P}_2^m$ can be assigned to all rooms in $\mathcal{R}^2$ and every patient $d_i^x \in \mathcal{P}_2^m$ can additionally be assigned to two distinct rooms in $\mathcal{R}^1$, which is modeled by

$$E_3 = \big\{\{d_1^x, r_1^x\}, \{d_1^x, r_2^x\}, \{d_2^x, r_3^x\}, \{d_2^x, r_4^x\} \mid x \in \mathcal{X}\big\} \cup \big\{\{p, r\} \mid p \in \mathcal{P}_2^m,\ r \in \mathcal{R}^2\big\}.$$

We illustrate the construction of G for $\psi = x_1 \vee \overline{x}_2 \vee x_3 \wedge \overline{x}_1 \vee \overline{x}_2 \vee \overline{x}_3 \wedge x_1 \vee x_2 \vee x_3 \wedge \overline{x}_1 \vee x_2 \vee \overline{x}_3$ in Fig. 1. Remark that G is bipartite with node partitions $\mathcal{P}^f \cup \mathcal{P}^m$ and $\mathcal{R}$.

We now prove that there exists a satisfying interpretation $\mathcal{I}$ for ψ if and only if we can find a feasible patient-to-room assignment for the corresponding single-day PRA instance.

"$\Leftarrow$": On the one hand, let M encode a feasible patient-to-room assignment. Then, for every variable $x \in \mathcal{X}$ either $M(v_1^x) = M(v_2^x) = r_1^x$ and $M(v_3^x) = M(r_4^x) = r_3^x$ holds or $M(v_2^x) = M(v_3^x) = r_2^x$ and $M(v_1^x) = M(v_4^x) = r_4^x$. We define an interpretation $\mathcal{I}$ for ψ as follows.

$$\mathcal{I}[x] := \begin{cases} \text{TRUE} & \text{if } M(v_1^x) = M(v_2^x) = r_1^x \text{ and } M(v_3^x) = M(v_4^x) = r_3^x, \\ \text{FALSE} & \text{otherwise.} \end{cases}$$

The assignment depicted in Fig. 1 corresponds for example to the interpretation $\mathcal{I}[x_1] = \text{TRUE}$, and $\mathcal{I}[x_2] = \mathcal{I}[x_3] = \text{FALSE}$.

Assume to the contrary that $\mathcal{I}$ is not a satisfying interpretation for ψ. Then there exists a clause c that is not satisfied which means that the patients u_1^c, u_2^c and u_3^c are not assigned to rooms in $\mathcal{R}^1$. By construction, the only other feasible room for these three patients is C^c. However, since at most two patients can be assigned to room C^c, M cannot be a feasible patient-to-room assignment.

"$\Rightarrow$": On the other hand, assume that there exists a satisfying interpretation $\mathcal{I}$ for ψ. We define the assignment $M : \mathcal{P}^f \cup \mathcal{P}_1^m \to \mathcal{R}$ as follows: For every variable $x \in \mathcal{X}$, we assign the female patient v_2^x together with v_1^x to room r_1^x if $\mathcal{I}[x] = \text{TRUE}$. If $\mathcal{I}[x] = \text{FALSE}$, we assign v_2^x and v_3^x to r_2^x. Thus, for $x \in \mathcal{X}$ and $i \in \{1, \ldots, 4\}$ we define

$$M(v_i^x) := \begin{cases} r_1^x & \text{if } i \in \{1,2\} \text{ and } x \in \mathcal{X},\ \mathcal{I}[x] = \text{TRUE}, \\ r_3^x & \text{if } i \in \{3,4\} \text{ and } x \in \mathcal{X},\ \mathcal{I}[x] = \text{TRUE}, \\ r_2^x & \text{if } i \in \{2,3\} \text{ and } x \in \mathcal{X},\ \mathcal{I}[x] = \text{FALSE}, \\ r_4^x & \text{if } i \in \{1,4\} \text{ and } x \in \mathcal{X},\ \mathcal{I}[x] = \text{FALSE}. \end{cases}$$

Every clause c is represented by three male patients, representing the clause's literals, and a room C^c. Following a satisfying interpretation, all patients corresponding to not-satisfied literals can be assigned to the room C^c, as at least one literal is satisfied. All patients corresponding to satisfied literals can be assigned to their feasible room in $\mathcal{R}^1$. We therefore define

$$M(u_i^c) := \begin{cases} r_j^x \in \mathcal{R}^1 \cap N(u_i^c) & \text{if literal number } j \text{ in clause } c \text{ is satisfied} \\ C^c \in \mathcal{R}^2 \cap N(u_i^c) & \text{otherwise,} \end{cases}$$

where $N(v)$ denotes the neighborhood, i.e., the set of adjacent nodes, of node v.

For every male patient $d_i^x \in \mathcal{P}_1^m$, $i \in \{1, 2\}$ there are exactly two feasible rooms in $\mathcal{R}^1$. To exactly one of these rooms, female patients are assigned and to the other one, exactly one male patient is assigned. We now assign d_i^x to this room as well:

$$M(d_i^x) := \begin{cases} r_1^x & \text{if } i = 1 \text{ and } x \in \mathcal{X},\ \mathcal{I}[x] = \text{FALSE}, \\ r_2^x & \text{if } i = 1 \text{ and } x \in \mathcal{X},\ \mathcal{I}[x] = \text{TRUE}, \\ r_3^x & \text{if } i = 2 \text{ and } x \in \mathcal{X},\ \mathcal{I}[x] = \text{FALSE}, \\ r_4^x & \text{if } i = 2 \text{ and } x \in \mathcal{X},\ \mathcal{I}[x] = \text{TRUE}. \end{cases}$$

Now, only the $\frac{2}{3}|\mathcal{X}|$ patients $e_i \in \mathcal{P}_2^m$ are left unassigned, and all rooms in $\mathcal{R}^1$ and all but $\frac{2}{3}|\mathcal{X}|$ rooms in $\mathcal{R}^2$ are filled. Thus, we can extend the assignment M to a feasible patient-to-room assignment by greedily assigning the patients e_i to one of those rooms in $\mathcal{R}^2$. In this way, we obtain a feasible patient-to-room assignment. $\square$

We can use the result that single-day PRA is strongly $\mathcal{NP}$-complete, to settle an open question concerning the so-called red-blue transportation problem (Red-Blue TP) stated in [26]. In this variant of the transportation problem, there is a set $S = R \cup B$ of red and blue supply nodes with supplies a_i, $i \in S$, a set of demand nodes D with demands b_j, $j \in D$ fulfilling $\sum_{i \in S} a_i = \sum_{j \in D} b_j$, and a bipartite graph $(S \cup D, E)$ with edge costs c_{ij} for all $\{i, j\} \in E$. In the decision version of Red-Blue TP, it is asked whether a feasible flow exists which sends all flow from the supply to the demand nodes and where all supply nodes that actually supply a demand node have the same color. In the optimization version additionally the total costs have to be minimized.

In [26] it was shown that the decision version of Red-Blue TP is strongly $\mathcal{NP}$-complete, even if $a_i = 1$ for all $i \in S$, and $b_j = 3$ for all $j \in D$. However, the complexity status for the even more restricted situation with $a_i = 1$ for all $i \in S$, and $b_j = 2$ for all $j \in D$ remained open. Since this case corresponds to single-day PRA with female patients R, male patients B, rooms D and room capacities b_j for all $j \in S = D$, the following corollary directly follows from Theorem 1.

Corollary 1. *The decision version of Red-Blue TP is strongly $\mathcal{NP}$-complete, even if $a_i = 1$ for all $i \in S$, and $b_j = 2$ for all $j \in D$.*

3. PRA without patient-room assignment restrictions

As stated in Section 2, the feasibility version of PRA is strongly $\mathcal{NP}$-complete if arbitrary patient-room assignment restrictions have to be taken into account. In the following, we assume that we do not have such restrictions, i.e., every patient can be assigned to every room. In practice, this assumption is usually satisfied for well-equipped rooms in modern wards. The main result of this section is the proof that in an optimal solution w.r.t. f^Σ every patient has to be moved at most once. This means that in real life both the wishes of staff (minimizing the total number of transfers) and that of the patients (being moved as few as possible) can be taken into account simultaneously.

We start again with the feasibility question. If an unlimited number of patient transfers is allowed, checking feasibility is easy, even for arbitrary room capacities $C_r = C \geq 1$: Let F_t, M_t be the number of female and male patients which need a room in time period $t \in \mathcal{T}$. Then, a necessary and sufficient condition for the existence of a feasible assignment is

$$\max_{t \in \mathcal{T}} \left\{ \left\lceil \frac{F_t}{C} \right\rceil + \left\lceil \frac{M_t}{C} \right\rceil \right\} \leq R. \tag{1}$$

However, in a corresponding solution several patient transfers may be necessary to achieve a feasible solution. In the following, we assume that (1) is satisfied, i.e., a feasible solution always exists. We study in more detail the question how many transfers (per patient and in total) are necessary. At first, we show that between the first two time periods no transfers are necessary.

Lemma 1. *For PRA with constant room capacities $C_r = C$, in an optimal solution with respect to f^Σ no transfer between the first two time periods is performed.*

Proof. We prove this lemma by contradiction. Assume that a solution with minimal total number of transfers is given where a patient is transferred between the first two time periods from one room to another. Then there must be another patient who leaves after the first time period. If these two patients simply switch beds in the first time period, this transfer is avoided and feasibility of the solution is not affected. Therefore, the original solution cannot be optimal with respect to f^Σ. □

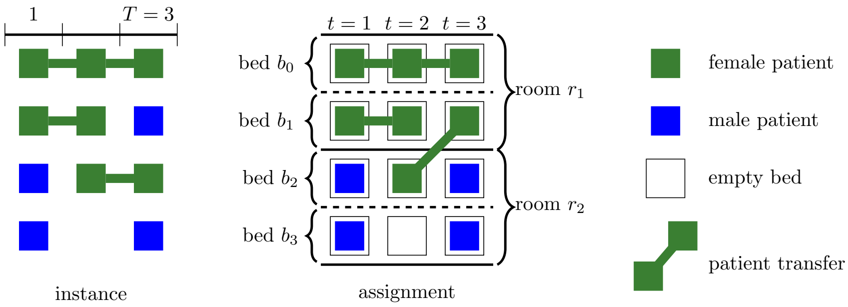

Figure 2: Example for $T = 3$ where patient transfers are necessary for feasibility

This result implies that for a planning horizon of $T = 2$ time periods no patient transfer is necessary. For $T = 3$, this statement is no longer valid which is shown by the example represented in Fig. 2: In the first time period, two female patients arrive, one who stays for three periods and one for two. In the second time period, another female patient arrives who stays for two periods as well. Both in the first and third time period, two male patients arrive who stay for one period each. In this case, one patient transfer is necessary to achieve a feasible assignment, cf. Fig. 2.

In the following, we concentrate on the case $C_r = 2$. We study the questions how many transfers are needed in total (objective f^{Σ}) and how often an individual patient is transferred during their stay (objective $f^{\max}$).

We start by proving that in an optimal solution w.r.t. f^{Σ}, every patient has to be transferred at most once, i.e., we have $f^{\max} \leq 1$. Theoretically, a patient can be transferred at any point during their stay. However, a patient's transfer is only necessary if the patient is alone in a room or the corresponding room-mate leaves the hospital at the end of the current time period. We call this the *transfer condition* (TC). In this situation, two possibilities for transfers exist: the patient who stays in hospital is moved to another room or another patient is moved into the room of the patient. By a (rather technical) interchange argument it can easily be shown that any optimal solution w.r.t. f^{Σ} can be transformed into a solution respecting TC without creating more transfers. Thus, from the last preceding discussion it follows directly that

Lemma 2. *There exists an optimal solution minimizing f^{Σ} which respects TC.*

For any solution respecting TC, we define the *transfer graph* as the digraph $D = (V, A)$ with nodes $V = \mathcal{P}$ and arcs

$$A = \{(p, q) \mid p \text{ is transfered to the room in which } q \text{ already is}\}.$$

This graph represents all performed transfers as no patient is ever transferred into an empty room. Furthermore, remark that neither feasibility nor optimality is affected if a transfer is performed the other way round. In Fig. 3 we illustrate the reversal of a transfer between two patients a and c. While in Fig. 3a patient a is moved to patient c between the 3rd and the 4th time period, in Fig. 3b patient c is moved to a. Remark that in all figures we ignore isolated nodes.

Before we can show that every patient has to be moved at most once in an optimal solution w.r.t. f^{Σ}, we prove the following property for the transfer graph.

Lemma 3. *The transfer graph is a forest.*

Proof. It suffices to prove that no induced circles exist in D. Assume to the contrary that $C = (p_1, \ldots, p_k)$ is a circle in D and w.l.o.g. let (p_1, p_2) be the earliest transfer. The next transfer that involves p_1 or p_2 takes place after p_1 or p_2 has left the hospital,

(a) Initial solution and digraph

(b) Revised solution

Figure 3: Transfers and their representation

according to TC. Thus, only one of p_1 and p_2 can be involved in another transfer, i.e., only one of the arcs (p_n, p_1) and (p_2, p_3) can exist. This is a contradiction to the existence of the circle C. $\square$

By using that a transfer graph is a forest, we can easily transform every given solution that respects TC into a solution in which every patient is transferred at most once, without increasing the total number of transfers:

Theorem 2. *There exists an optimal solution w.r.t. f^{Σ} such that every patient is transferred at most once.*

Proof. Assume that an optimal solution w.r.t. f^{Σ} is given which respects TC. We can transform it into an equivalent optimal solution in which every patient is transferred at most once if we can reorientate the arcs in D such that every node has at most one outgoing arc.

As D is a forest according to Lemma 3, such an orientation is achieved by orientating each subtree from the leaves to the root. For example, in Fig. 3a we choose patient b as root and reorientate D accordingly. $\square$

Theorem 2 implies that the objective functions f^{Σ} and $f^{\max}$ are not conflicting. Moreover, they can be optimized simultaneously by adding the condition that every patient is transferred at most once to the optimization problem for f^{Σ}. In the following, we call a solution that respects TC and in which every patient is transferred at most once, a *proper* solution. According to Lemma 2 and Theorem 2, always a proper optimal solution exists.

Note that for room capacities $C_r = 3$, in an optimal solution w.r.t. f^{Σ} it may be not sufficient to transfer every patient at most once. This can be seen from the example in Fig. 4. If we restrict our considerations to solutions where every patient is transferred at most once, in total at least $f^{\Sigma} = 3$ transfers are necessary (cf. Fig. 4a). However, if we allow that a patient may be transferred twice, a better solution with $f^{\Sigma} = 2$ transfers is possible (cf. Fig. 4b).

Also, Theorem 2 is no longer valid if rooms have different capacities. If we have two rooms $\mathcal{R} = \{r_1, r_2\}$ with capacities $C_{r_1} = 1$ and $C_{r_2} = 2$, even to achieve a feasible solution, it may happen that a patient has to be transferred between every two consecutive time periods, cf. Fig. 5.

Although we know that a proper optimal solution must exist, we do not know yet how to compute one. However, if we only want to compute a proper feasible solution, this can be done in polynomial time. For this, we can use the procedure described in the proof of Theorem 2 to transform any given solution that respects TC into a proper solution without increasing the total number of transfers. Therefore, we can also use it to compute a proper feasible solution as we only need a TC-respecting feasible solution to start from.

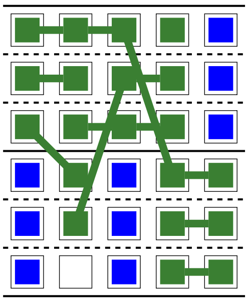

(a) $f^{\max} = 1$ and $f^{\Sigma} = 3$

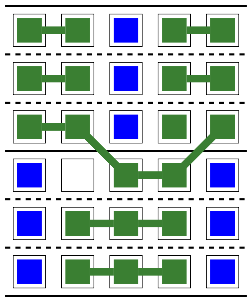

(b) $f^{\max} = 2$ and $f^{\Sigma} = 2$

Figure 4: Optimal solution w.r.t. f^{Σ} for $C_r = 3$ where a patient has to be transferred twice

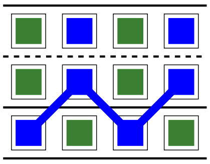

Figure 5: Multiple transfers per patient

Lemma 4. *A solution that respects TC can be computed in polynomial time in $\mathcal{O}(PR)$.*

Proof. A simple greedy heuristic designs a solution that respects TC: We start by sorting the patients according to their arrival periods. Then greedily assign patients to rooms as follows: Assume the next patient in line is a female patient p. If there exists a room r with only one female patient in time period a_p, assign p to r as well. If no such room but an empty room in time period a_p exists, assign p to it. Otherwise, select two rooms r_1, r_2 with one male patient in time period a_p each, transfer the male patient from r_2 to r_1, and assign p to r_2. Proceed for male patients analogously. □

Corollary 2. *A proper feasible solution can be computed in polynomial time in $\mathcal{O}(PR)$.*

4. No need to transfer patients who arrive in the first time period

In Section 3, we proved that there always exists an optimal solution with respect to f^{Σ} in which no patient is transferred more than once. Furthermore, $f^{\max} = 0$ holds if and only if $f^{\Sigma} = 0$. In the following we concentrate on the question how many transfers are needed in total and we assume all considered solutions to be proper.

In this section, we prove that in an optimal solution patients who arrive in the first time period are never transferred. We will use this to derive upper bounds on f^{Σ} in Section 5. We call a patient who arrives in the first time period a t_0-*patient*. However, we first need some more insights about the time dependencies between transfers in a proper solution as well as the corresponding transfer graph.

Lemma 5. *Let $D = (V, A)$ be the transfer graph corresponding to a proper solution. Then, for every path in D the earliest transfer is represented by its first or last arc.*

Proof. To prove this statement, it suffices to show that of every three consecutive arcs, the middle one cannot correspond to the earliest transfer of these three: Let

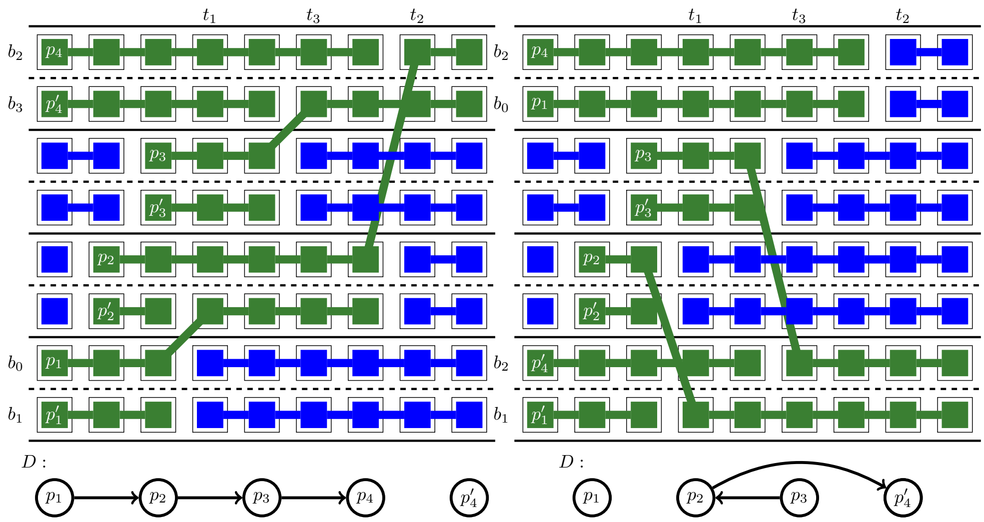

(a) initial solution

(b) revised solution

Figure 6: Example for an assignment revision such that no t_0-patients are transferred with $t_1 < t_3 < t_2$ and beds b_0, b_1, b_2, b_3

(a, b, c, d) be a directed path of length three and assume the transfer (b, c) is the first of them. In order to perform the transfer (a, b) afterwards, patient c is required to leave the hospital according to TC. Therefore, the transfer (c, d) cannot be performed anymore. Thus, the transfer (b, c) cannot be the earliest. As the orientation of the arcs/transfers does not affect feasibility nor optimality, the statement holds for all directed and nondirected paths of length three. □

Remark that Lemma 5 especially holds for every subpath. Therefore, every path can be divided into two chronologically sorted sequences of transfers: Let $D = (V, A)$ be the transfer graph corresponding to a proper solution. Let $W = (p_1, p_2, \ldots, p_k)$ be a directed path in D. Without loss of generality, let (p_1, p_2) represent the earliest transfer in W and $(p_\tau, p_{\tau+1})$ with $\tau \in \{2, \ldots, k-1\}$ the latest transfer in W. Furthermore, we denote by t_j the time period in which the transfer (p_j, p_{j+1}) is performed, $1 \leq j < k$. Then the transfers corresponding to the subpath $(p_1, \ldots, p_\tau, p_{\tau+1})$ are sorted chronologically increasing, i.e., $t_1 < t_2 < \ldots < t_\tau$, and the transfers corresponding to the subpath $(p_\tau, p_{\tau+1}, \ldots, p_k)$ are sorted chronologically decreasing, i.e., $t_\tau > \ldots > t_{k-2} > t_{k-1}$. An example instance that corresponds to the path (p_1, p_2, p_3) with $t_1 < t_3 < t_2$ is illustrated in Fig. 6a.

To be able to alter a given solution such that it remains feasible, we do not only need information about the sequence of transfers but also about the time periods when the patients are discharged. As we assume that TC is respected by our solution, we can easily derive when each patient leaves the hospital:

D1 In time period t_j patients p_{j-1} and p'_j are discharged for $1 \leq j < \tau$, where p'_j denotes the room-mate of p_j before p_j is involved in its first transfer and $p_0 := p'_1$;

D2 In time period t_τ patients $p_{\tau-1}$ and $p_{\tau+2}$ are discharged where, in case of $\tau = k - 1$ we define $p_{k+1} := p'_k$;

D3 In time period t_j patients p_{j+2} and p'_j are discharged for $\tau < j < k$, where p'_j denotes the room-mate of p_j before p_j is involved in its first transfer and $p_{k+1} := p'_k$.

Using the information about the sequence of transfers and about the discharge of patients, we can alter a given solution such that no t_0-patients are transferred without increasing the total number of transfers. First, we prove that no two t_0-patients are connected via a path in D.

Lemma 6. *Let $D = (V, A)$ be the transfer graph corresponding to a proper optimal solution. Then all t_0-patients are in different connected components of D.*

Proof. We prove this lemma by contradiction. Therefore, assume D' to be a connected component of D that contains two t_0-patients p and q. Without loss of generality, the transfer involving p takes place earlier than the transfer involving q. Choose q as root of D' and reorientate D' accordingly from leaves to root, if q is not already its root. Further, let $W = (p = p_1, p_2, \ldots, p_{k-1}, p_k = q)$ be the unique directed path connecting p and q in the revised transfer graph. Let b_0 denote the bed to which p is assigned in the first time period, b_1 that of its room-mate p'_1. Analogously, let b_2 denote the bed to which q is assigned in the first time period and b_3 that of its room-mate p'_k. Remark that in the original solution b_0 and b_1 are located in the same room as well as b_2 and b_3. For our revised solution we may change this pairing later.

We construct a revised solution by assigning all t_0-patients to their beds for their complete stay. The other patients p_j, for $2 \leq j \leq k-1$, start off in their original beds as well, and we now define their transfers. We denote again with t_τ the time period of the latest transfer in W. For $1 \leq j \leq \tau - 1$, we transfer p_{j+1} in time period t_j to bed $b_{j \bmod 2}$. For $\tau < j < k$, we transfer p_j in time period t_j to bed $b_{2+(k-j) \bmod 2}$. All other bed assignments stay as before. Lastly, pair the two beds to which p_τ and $p_{\tau+1}$ are assigned to the same room and the remaining two to another room. An example of such a revision for t_0-patients p_1 and p_4 and $W = (p_1, p_2, p_3, p_4)$, with $t_1 < t_3 < t_2$, is illustrated in Fig. 6. Remark that in Fig. 6a patients p_1 and p'_1 are room-mates in the first time period whereas in Fig. 6b patients p_1 and p_4 are room-mates.

The revised solution is feasible according to observations D1 to D3, and needs one transfer less than the initial solution which is a contradiction to its optimality. Thus, no connected component of the transfer graph corresponding to a proper optimal solution contains two t_0-patients. □

Finally, we prove that there always exists a proper optimal solution in which no t_0-patients are transferred.

Theorem 3. *Let $D = (V, A)$ be the transfer graph corresponding to a proper optimal solution. Then there exists an equivalent proper optimal solution in which patients who arrive in the first time period are never transferred.*

Proof. For every patient p who arrives in the first time period and is transfered we perform the following revision. Let D' denote the connected component of the corresponding transfer graph that contains p. According to Lemma 6, patient p is the only t_0-patient in D'. Therefore choose p as new root and reorientate D' accordingly from leaves to root. As no additional transfers are needed, the revised transfer graph corresponds then to a proper optimal solution in which p is not transferred and also no other t_0-patient is transferred instead. □

Combining these insights with the greedy procedure introduced in the proof of Lemma 4, allows to compute in polynomial time a proper schedule where t_0-patients are never transferred.

Lemma 7. *A feasible proper solution in which patients who arrive in the first time period are never transfered can be computed in $\mathcal{O}(PR)$ time.*

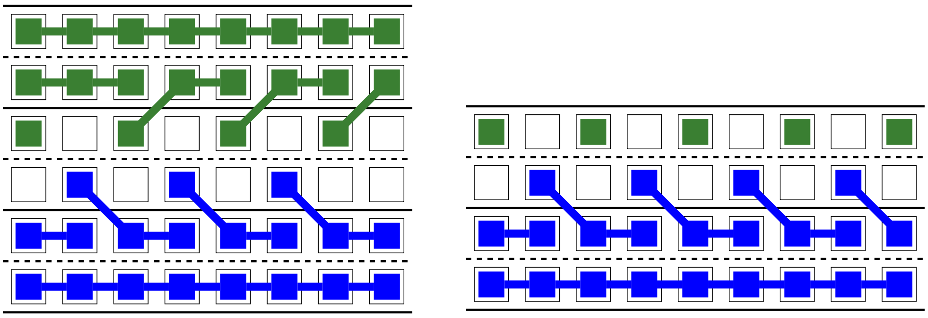

(a) Example for $R = 3$ where $P - (2R - 1)$ patient transfers are necessary, cf. Theorem 4

(b) Example for $R = 2$ where $\frac{P-3}{2}$ patient transfers are necessary, cf. Conjecture 1

Figure 7: Examples for tightness of bounds with R rooms and P patients

Proof. According to Corollary 2 we can compute a feasible proper solution in $\mathcal{O}(PR)$ time and use the construction from Theorem 3 to alter it in $\mathcal{O}(P)$ time so that no t_0-patients are transferred. □

5. Bounds on the number of patient transfers

We proved in Section 3 that the objective function $f^{\max}$ is bounded from above by 1. In this section we will use this insight to determine upper bounds on the optimal value of f^{Σ} for rooms with capacity $C_r = 2$. As parameters, we use the number of patients P, the number of rooms R and the length of the planning horizon T. Bounds depending on P and R are helpful in operational decision-making as they take the actual situation into account. They are also helpful in tactical decision-making regarding one medical speciality department, as they reflect some of the specialities characteristics, e.g., the patients' length of stay. Bounds depending only on the length of the planning horizon however are helpful in general tactical or strategic decision-making as they hold for all wards regardless of their speciality.

First, we use the structural insights from Section 4 to determine an upper bound that uses the number of patients P and rooms R as parameters.

Theorem 4. *For R rooms, an optimal solution with at most $P - (2R - 1)$ patient transfers exists and this bound is tight for $R \geq 3$.*

Proof. We consider a proper optimal solution. In order to make a transfer between time periods two and three necessary, at least all beds but one need to be occupied in the first time period, i.e., $2R - 1$ patients need to arrive in the first time period. According to Theorem 3, patients who arrive in the first time period are never transferred. Therefore, more than $P - (2R - 1)$ transfers are never necessary.

The example in Fig. 7a shows that this bound is tight in case of three rooms. This example can be extended to any larger number of rooms by adding the required amount of rooms and patients who stay for the complete length of the planning horizon. Those patients are obviously never transferred but, as they arrive in the first time period, the bound of $P - (2R - 1)$ is still valid. □

Although we know that the bound from Theorem 4 is tight for $R \geq 3$, we do not know if this also holds true in the case $R = 2$. We suspect, however, that a tight upper bound for f^{Σ} in the case $R = 2$ is much smaller. The worst example that we were able to construct needs one transfer every two days resulting in a total of $\frac{P-3}{2}$ transfers, cf. Fig. 7b.

Conjecture 1. *For $R = 2$ at most $\left\lfloor \frac{P-3}{2} \right\rfloor$ patient transfers are necessary.*

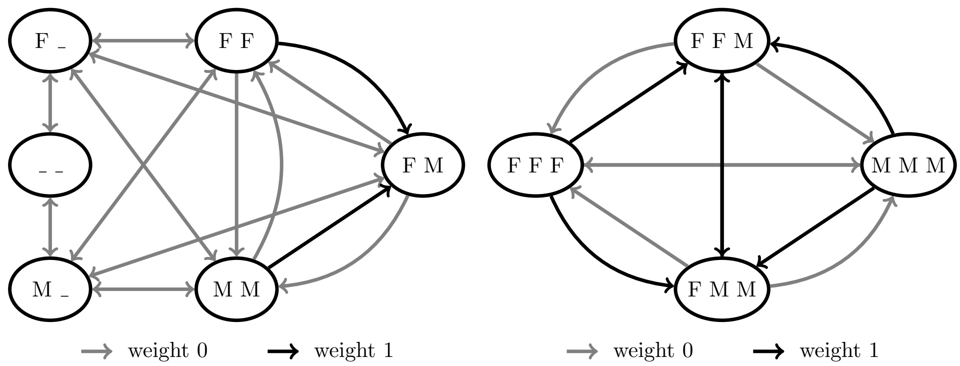

(a) State graph for $R = 2$

(b) Reduced state graph for $R = 3$

Figure 8: State graphs

Second, we use insights about the practical occurrence of patient transfers to determine tight upper bounds for f^{Σ} depending only on the length of the planning horizon for $R \in \{2, 3\}$.

We start by considering the situation of $R = 2$ rooms. In every room either no patient is present, one or two female, or one or two male patients are present. We call such rooms empty, female or male rooms according to the present patients' gender and mark them with _, F, and M, respectively.

Hence, in each time period $t \in \mathcal{T}$ there are six possible states of the two rooms: There are either two empty rooms (_ _), or one room with female patients and the other room empty (F _), or one room with male patients and the other room empty (M _), or two rooms with female patients (F F), or two rooms with male patients (M M), or one room with female patients and one room with male patients (F M).

We introduce a directed state graph having the possible states as nodes and arcs denoting a change from one state to another in consecutive time periods, cf. Fig. 8a. Each arc is weighted with the maximal number of transfers which may be necessary to get from one state to the other. It is easy to see that these weights are 0 or 1, and that weight 1 only occurs on the arcs (F F, F M) and (M M, F M).

Lemma 8. *For $R = 2$, an optimal solution with at most $\lfloor \frac{T-1}{2} \rfloor$ patient transfers exists. Furthermore, there are instances achieving this bound.*

Proof. Since among the states _ _, F _, and M _obviously no patient transfers occur, it is sufficient to consider the subgraph containing states F F, M M, and F M. Each solution of PRA with T time periods induces a directed path with $T - 1$ arcs in the state graph. If we are interested in a solution with a maximum number of patient transfers, we have to look for a path with $T - 1$ arcs having total maximum weight. Since according to Lemma 1, always an optimal solution exists where no transfers occur between the first two time periods, and since in the state graph for any directed path each arc with weight 1 must be followed by an arc with weight 0, each solution needs at most $\lfloor \frac{T-1}{2} \rfloor$ transfers. For example, the cycle (F F, F M, F F, F M,. . .) has this weight, cf. Fig. 8a. That this bound is tight, can be seen from the example in Fig. 7b, where $\lfloor \frac{T-1}{2} \rfloor$ patient transfers are necessary to obtain a feasible solution. □

For $R = 3$, we define the states as triples of _, F and M, analogously to the case $R = 2$. We consider the corresponding state graph limited to nodes where all rooms are occupied, as shown in Fig. 8b.

Lemma 9. *For $R = 3$, an optimal solution with at most $T - 2$ patient transfers exists. Furthermore, there are instances achieving this bound.*

Proof. As in the case $R = 2$, each solution of PRA with T time periods induces a directed path with $T - 1$ arcs in the state graph. Since according to Lemma 1, always an optimal solution exists where no transfers occur between the first two time periods, each solution needs at most $T - 2$ transfers. This corresponds to a path where each arc has weight 1. For example, the cycle alternating between the states F F M and F M M has this weight, cf. Fig. 8b. That this bound is tight, can be seen from the example in Fig. 7a, where $T - 2$ patient transfers are necessary to obtain a feasible solution. □

Unfortunately, for $R \geq 4$ the state graphs and especially the computation of the arcs' weights get much more complex. Up to now, we were not able to derive upper bounds parametrized by T for a larger number of rooms.

6. Conclusion

Assigning patients to beds is an every-day task in hospital wards, but it has gained more and more attention in literature only recently [25]. In this paper, we defined the underlying basic problem without tailoring it to the situation in one specific ward or hospital. We proved that if there are patient-room assignment restrictions, the problem is strongly $\mathcal{NP}$-complete even if all rooms contain exactly two beds. This result also settles an open question concerning the complexity status of a special case of the red-blue transportation problem [26]. Further, we showed that if no such patient-room assignment restrictions exist and all rooms contain exactly two beds, we can find an assignment that minimizes the total number of transfers and, at the same time, every patient is transferred at most once. Hence, no compromises between nurses' and patients' wishes need to be made.

In literature, e.g. [2], heuristics were already developed without discussion of potential limitations of the solution space. Our results prove that for double bedrooms optimality is not affected if we limit the solution space to assignments where every patient is transferred at most once. However, for three or more beds per room optimality with respect to the total number of transfers is lost. Apart from providing structural insights about PRA, we provided a constructive algorithm that converts any feasible assignment into one in which every patient is transferred at most once without changing the total number of transfers.

In strategic or tactical capacity planning in hospitals, the number of transfers determines one important factor, e.g., on the work load of nurses. Our bounds provide estimates on this value and, hence, improve long term decision-making. In literature such bounds were missing so far or were just provided for a weighted sum approach combining work force needs and patient satisfaction. These are in general not suitable as input in the latter case.

Although, our definition already fits the real-life task in most wards, sometimes additionally needs and wishes of patients are considered in the determination of room-mates [23]. In future work, our aim is to use the structural insights presented in this paper to develop and test solution approaches that are applicable in real life and adaptable to different situations. Further, we keep working on the conjecture and on finding an upper bound on f^{Σ} depending on T for $R \geq 4$.

References

[1] P. Berman, M. Karpinski, and A. Scott. Approximation hardness of short symmetric instances of MAX-3SAT. *Electronic Colloquium on Computational Complexity (ECCC)*, 2003.

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

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

[4] S. Ceschia and A. Schaerf. Dynamic patient admission scheduling with operating room constraints, flexible horizons, and patient delays. *Journal of Scheduling*, 19(4):377–389, 2016.

[5] J. K. Cochran and A. Bharti. Stochastic bed balancing of an obstetrics hospital. *Health Care Management Science*, 9(1):31–45, 2006.

[6] P. Demeester, W. Souffriau, P. De Causmaecker, and G. Vanden Berghe. A hybrid tabu search algorithm for automatically assigning patients to beds. *Artificial Intelligence in Medicine*, 48(1):61–70, 2010.

[7] B. Epping. Es muss nicht immer Case Management drauf stehen, wenn es Case Management ist - Interview mit Sabine Peer. *Orthopädie und Unfallchirurgie aktuell*, 12(01):11–14, 2017.

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

[9] A. I. Hammouri and B. Alrifai. Investigating biogeography-based optimisation for patient admission scheduling problems. *Journal of Theoretical & Applied Information Technology*, 70(3), 2014.

[10] P. R. Harper. A framework for operational modelling of hospital resources. *Health Care Management Science*, 5(3):165–173, 2002.

[11] P. R. Harper, A. K. Shahani, J. E. Gallagher, and C. Bowie. Planning health services with explicit geographical considerations: A stochastic location-allocation approach. *Omega*, 33(2):141–152, 2005.

[12] G. W. Harrison, A. Shafer, and M. Mackay. Modelling variability in hospital bed occupancy. *Health Care Management Science*, 8(4):325–334, 2005.

[13] X.-M. Huang. A planning model for requirement of emergency beds. *Mathematical Medicine and Biology: A Journal of the IMA*, 12(3-4):345–353, 1995.

[14] A. Hübner, H. Kuhn, and M. Walther. Combining clinical departments and wards in maximum-care hospitals. *OR Spectrum*, 40(3):679–709, 2018.

[15] P. J. H. Hulshof, N. Kortbeek, R. J. Boucherie, E. W. Hans, and P. J. M. Bakker. Taxonomic classification of planning decisions in health care: A structured review of the state of the art in OR/MS. *Health Systems*, 1(2):129–175, 2012.

[16] B. J. Masterson, T. G. Mihara, G. Miller, S. C. Randolph, M. E. Forkner, and A. L. Crouter. Using models and data to support optimization of the military health system: A case study in an intensive care unit. *Health Care Management Science*, 7(3):217–224, 2004.

[17] T.O. Paulussen, N. R. Jennings, K.S. Decker, and A. Heinzl. Distributed patient scheduling in hospitals. In *18th International Joint Conference on Artificial Intelligence (01/01/03)*, pages 1224–1229, 2003.

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

[19] A. V. Roth and R. van Dierdonck. Hospital resource planning: Concepts, feasibility, and framework. *Production and Operations Management*, 4(1):2–29, 1995.

[20] F. Schäfer, M. Walther, A. Hübner, and H. Kuhn. Operational patient-bed assignment problem in large hospital settings including overflow and uncertainty management. *Flexible Services and Manufacturing Journal*, 2019.

[21] R. Schmidt, S. Geisler, and C. Spreckelsen. Decision support for hospital bed management using adaptable individual length of stay estimations and shared resources. *BMC Medical Informatics and Decision Making*, 13(1):1–19, 2013.

[22] N. Stab and W. Hacker. *Gute Stationsorganisation. Ein Leitfaden für Pflegeeinrichtungen.* Bundesanstalt für Arbeitsschutz und Arbeitsmedizin, Dortmund, 1. edition, 2016.

[23] N. Stab and W. Hacker. *Stationsorganisation im Krankenhaus. Entwicklung und Erprobung eines kriteriengeleiteten Bewertungs- und Gestaltungsverfahrens.* Bundesanstalt für Arbeitsschutz und Arbeitsmedizin, Dortmund, 1. edition, 2016.

[24] D. Steenken, S. Voß, and R. Stahlbock. Container terminal operation and operations research - a classification and literature review. *OR Spectrum*, 26(1):3–49, 2004.

[25] T. Huke and B. Spieß. Belegungsmanagement in Akutkrankenhäusern. Technical report, 2009.

[26] W. Vancroonenburg, F. Della Croce, D. Goossens, and F. C. R. Spieksma. The red–blue transportation problem. *European Journal of Operational Research*, 237(3):814 – 823, 2014.

[27] W. Vancroonenburg, P. De Causmaecker, F. C. R. Spieksma, and G. Vanden Berghe. Scheduling elective patient admissions considering room assignment and operating theatre capacity constraints. *Proceedings of the 5th International Conference on Applied Operational Research, Lecture Notes in Management Science*, 5:153–158, 2013.

[28] A. Zöller, L. Braubach, A. Pokahr, F. Rothlauf, T. O. Paulussen, W. Lamersdorf, and A. Heinzl. Evaluation of a Multi-Agent System for Hospital Patient Scheduling. *International Transactions on Systems Science and Applications*, 1(4):375–380, 2006.