Patient-to-room assignment planning in a dynamic context

Wim Vancroonenburg · Patrick De Causmaecker · Greet Vanden Berghe

Received: date / Accepted: date

Abstract The present contribution proposes an extension to the patient assignment (PA) planning problem in a dynamic context. Two ILP-models have been developed for optimizing this day- to- day planning problem. The first considers finding the optimal assignment for newly arrived patients, whereas the second also considers future, but planned, arrivals. The performance of both models is compared to each other on a set of benchmark instances. The relative performance with respect to a known lower bound is also presented. Furthermore, the effect of uncertainty on the patients' length of stay is studied, as well as the effect of the percentage of emergency patients. The results show that the second model provides better results under all conditions, while still being computationally tractable.

Keywords Patient assignment problem · Dynamic planning · Integer Linear Programming

1 Introduction

Rooms and beds belong to the critical assets of just any hospital. They account for a considerable part of a hospital's infrastructure, and a large amount of financial resources are invested in equipping them with medical apparatus to facilitate patient care. Furthermore, they also represent the place where most patients will spend a large part of their stay, as they recover from surgery, wait for examinations to take place, etc. In order to improve their comfort, patients are offered a choice between single bed rooms, luxury rooms with

Wim Vancroonenburg
CODeS research group, Computer Science, KAHO Sint-Lieven
Gebroeders De Smetstraat 1, 9000 Ghent, Belgium
Tel.: +32 9 265 87 04
Fax: +32 9 225 62 69
E-mail: Wim.Vancroonenburg@kahosl.be

private showers, and other amenities. As a result, a large variety of hospital rooms exists in terms of capacity, which are equipped with different medical apparatus and amenities. Assigning patients to such a variety of hospital rooms can therefore be challenging, necessitating an efficient plan for making such an assignment.

Bed managers aim at finding an assignment of patients to rooms that strikes a balance between patients' preferences and comfort on the one hand, and patients' clinical conditions and the resulting required room facilities on the other. However, both the availability of rooms and equipment, and hospital policies and standards need to be considered, making it difficult to generate a balanced patient-to-room assignment. A lack of overview on occupied beds and the uncertainty on how long patients will stay in the hospital, further complicate the matter.

Demeester et al (2010) defined and studied the patient assignment (PA) problem in the context just described. They consider a set of patients that arrive at a hospital over a certain period of time. The hospital comprises a set of rooms, each with given capacity and characteristics. The problem is to find an effective assignment of patients to rooms, satisfying room capacity restrictions. Moreover, a perceived cost is associated with each patient to room assignment relating to the *appropriateness* of that assignment. The objective is to minimize the total cost of these assignments. The present contribution focuses on this problem.

1.1 Related work

As pointed out by Rais and Viana (2011) in their survey on operations research in healthcare, a great deal of the considered literature has focussed on scheduling of patients and hospital resources. Notably, nurse rostering and operating theatre (OT) planning and scheduling have received a considerable amount of attention (see e.g. Burke et al 2004; Cardoen et al 2010), which is evident given that personnel and the OT are among the most expensive resources for any hospital.

The PA problem considered in this paper comprises an assignment problem that occurs at the operational level of hospital admission offices. It assumes that patients have already been attributed an admission date, a decision that is made as part of either an intervention scheduling[1] process (see e.g. Riise and Burke 2010) during operational surgery scheduling, or an appointment scheduling process when no surgery is required. The type of patients and the arrival pattern of patients with different pathologies is often also largely influenced by the Master Surgery Schedule (MSS), a timetable that allocates operating rooms and operating time to different medical disciplines. For example, Beliën and Demeulemeester (2007) show how a cyclic MSS can be constructed that results in an expected levelled bed occupancy.

[1] Also known as *advance scheduling.*

Demeester et al (2010) introduced the PA problem to the academic community as a challenging combinatorial optimization problem. In a follow up paper, Bilgin et al (2012) presented a new hyper-heuristic algorithm to the PA problem. They provide new benchmark instances and report test results. Vancroonenburg et al (2011) show that the PA problem is NP-hard.

The problem was also studied by Ceschia and Schaerf (2011), who developed a Simulated Annealing algorithm that improves on the best known results for the benchmark instances. Lower bounds for these instances are provided as well. Furthermore, they argue that the problem definition only assumes patients that are planned in advance (elective patients), and that it does not capture the dynamics of uncertainty on patient arrivals and departures. An extension to the problem definition is proposed where patient admission and discharge dates are revealed a few days before they occur (denoted as the *forecast level*). To this end, Ceschia and Schaerf developed a dynamic version of their algorithm that can be used for day- to- day scheduling. The performance of this algorithm is analysed under an increasingly larger forecast level.

The PA problem where patient transfers are not allowed, is related to the interval scheduling problem: patients can be represented by fixed length intervals/jobs with fixed start and end time, that need to be assigned to a machine (a room) for ‘processing’. The PA problem comprises required jobs and non-identical machines with different capacities, the goal being to find a minimum-cost schedule subject to side-constraints. In the dynamic context, it constitutes an *online* interval scheduling problem with uncertainty on the interval lengths. We refer to Kolen et al (2007) for a review on the subject of (online) interval scheduling problems. Ouelhadj and Petrovic (2009) give a survey of dynamic scheduling problems in manufacturing in general.

1.2 Present contribution

Similarly to Ceschia and Schaerf (2011), we define a new extension to the PA problem in a dynamic context. To this end, registration dates for each patient are added to the problem definition to denote when a patient’s (possibly future) arrival time is revealed. Contrary to Ceschia and Schaerf (2011) however, an estimate of the *length of stay* (LOS) for each patient is also assumed to be available, which in practice often is the case. Special care is taken to accommodate the decision process when patients outstay their estimated length of stay.

This dynamic version of the problem is modelled and solved using Integer Linear Programming (ILP). We discuss the performance of this approach and study the effect of the percentage of emergency cases and the accuracy of the LOS estimate.

2 Problem formulation

2.1 PA in a static context

The PA problem considers a set of patients P that each need to be assigned to one of a set of hospital rooms R over a certain time horizon $H = \{1, \ldots, T\}$. Each room $r \in R$ has a given capacity, denoted by $c(r)$. Each patient $p \in P$ is attributed an arrival time $a(p)$ and a departure time $dd(p)$, with the time interval $[a(p), dd(p))$ representing the patient's stay in the hospital. The length of the patient's stay, $dd(p) - a(p)$, is denoted as $los(p)$.

The problem is to find an assignment $\sigma : P \mapsto R$ of patients to rooms that minimizes a certain cost $w(\sigma)$ related to these assignments. This cost $w(\sigma)$ consists of two parts:

- $w_1(\sigma) = \sum_{p \in P} los(p) \cdot c(p, \sigma(p))$: each patient/room combination is attributed a cost $c(p, r)$, that relates to the *appropriateness* of assigning patient p to room r for one time interval (the lower $c(p, r)$, the better). The goal is to minimize the sum of these assignment costs.
- $w_2(\sigma) = \sum_{r \in R} \sum_{t=1}^{T} Conflict_{\sigma,r,t}$: the sum of all *gender conflicts* in all rooms r over the entire planning horizon H. The goal is to avoid that male and female patients are assigned to the same room at the same time. These conflicts are calculated as follows:

$$\begin{aligned} Conflict_{\sigma,r,t} = \min(\quad & |p \in P_{\sigma,r,t} : p \text{ is } male|, \\ & |p \in P_{\sigma,r,t} : p \text{ is } female|) \end{aligned} \tag{1}$$

with

$$P_{\sigma,r,t} = \{p \in P : a(p) \leq t < dd(p), \sigma(p) = r\} \tag{2}$$

denoting the set of patients assigned to room r at time t. Furthermore, this assigment should respect the room capacities at all times, i.e. :

$$\forall t = 1, \ldots, T,\ r \in R : \quad |P_{\sigma,r,t}| \leq c(r) \tag{3}$$

The definition proposed by Demeester et al (2010) allows for patients to be transferred from one room to another during their stay. The present contribution considers the slightly simpler version of the problem, which does not allow for transfers.

2.2 PA in a dynamic context

In practice, the arrivals and departures of patients are gradually revealed over the planning horizon. The problem definition is therefore extended to account for these dynamics. Each patient p is attributed a *registration* date $r(p)$, at which point the patient becomes known to the system, and an *expected* departure date $ed(p)$, which is an estimate of the patient's departure date. The departure date of the patient $dd(p)$ however, remains hidden until the patients' departure date has passed.

At each point $t' \in H$ of the planning horizon, two sets are revealed:

- $P_{t'}$: the set of patients with $r(p) = t'$, i.e. the patients that are registered at time t'. At this point, for each patient $p \in P_{t'}$ only $a(p)$ and $ed(p)$ are known, $dd(p)$ remains hidden.
- $DP_{t'}$: the set of patients with $dd(p) = t'$, i.e. the patients that leave the hospital at time t'.

Let $A_{t'}$ denote the set of patients that arrived at t', i.e. :

$$A_{t'} = \{p \in P : a(p) = t'\} \tag{4}$$

The goal of the problem is to find at each time t' an assignment

$$\sigma_{t'} : \bigcup_{i=1}^{t'} A_i \mapsto R \tag{5}$$

that maps each *arrived* patient p (i.e. all p for which $t' \geq a(p)$) to a hospital room r. Furthermore, patients who arrived before t' should not be moved, i.e. :

$$\forall p \in \bigcup_{i=1}^{t'-1} A_i : \quad \sigma_{t'}(p) = \sigma_{t'-1}(p) \tag{6}$$

The assignment σ_T denotes the solution at the end of the planning horizon. It contains all the patients' assignments within that period. The solution quality can be assessed by computing $w(\sigma_T)$. It is interesting to compare this value with the quality obtained for the static variant of the problem, which supposes that each patient's departure date is fixed in advance. Any lower bound for (or the optimal solution to) the static version is a lower bound for the dynamic problem.

3 Optimization models

Two models were developed for the dynamic PA planning problem that correspond to the situation at each decision step t'. They extend the previous assignment $\sigma_{t'-1}$ to include available information on newly arrived patients $p \in A_{t'}$.

The first approach is modelled after current practice, namely the assignment decision is made shortly before patient arrival and only current room availability is considered. The model tries to find the optimal assignment for the patients who arrived at the current decision step. Moreover, it uses the estimate of the newly arrived patients' LOS. The second model builds on the previous model by also considering all registered patients at each decision step, therefore anticipating future occupancy and room demand.

Both models are implemented as ILP-models. They are described in Sections 3.1 and 3.2. To simplify the description, the following notation will be used:

- $\mathbf{P_{t'}} = \bigcup_{i=1}^{t'} P_i$, the set of all registered patients up till t',

- $\mathbf{A_{t'}} = \bigcup_{i=1}^{t'} A_i$, the set of all patients who arrived up till t',
- $P^M, P^F \subseteq P$, restricts a set of patients to either *males* or *females* respectively,
- $elos(p) = \max(ed(p) - a(p), t' - a(p))$, the expected length of stay of patient p as it is known at the decision time t'. If the patient's stay has exceeded his or her expected departure date $ed(p)$, he or she is expected to stay at least one day longer.
- $\mathbf{AP_{tt'}} = \{p \in \mathbf{A_{t'}} : a(p) \leq t < a(p) + elos(p)\}$, the set of *arrived* patients that are *expected* to be present at time t,
- $\mathbf{PP_{tt'}} = \{p \in \mathbf{P_{t'}} : a(p) \leq t < a(p) + elos(p)\}$, the set of *registered* patients that are *expected* to be present at time t.

3.1 Model 1

The decision variables are defined as follows:

$$x_{p,r} = \begin{cases} 1 & \text{if patient } p \text{ is assigned to room } r, \\ 0 & \text{otherwise.} \end{cases} \tag{7}$$

$$v_{r,t} = \text{the number of } \textit{gender} \text{ conflicts in room } r \text{ at time } t \tag{8}$$

$$y_{r,t} = \begin{cases} 1 & \text{if the number of } \textit{male} \text{ patients assigned to room } r \\ & \text{at time } t \text{ is larger than or equal to the number of } \textit{female} \text{ patients,} \\ 0 & \text{otherwise.} \end{cases} \tag{9}$$

The optimization problem is then modelled as follows:

$$\text{Min} \sum_{p \in \mathbf{A_{t'}}} \sum_{r \in R} elos(p) \cdot c(p,r) \cdot x_{p,r} + \sum_{r \in R} \sum_{t=1}^{T} w_G \cdot v_{r,t} \tag{10}$$

$s.t.$

$$\sum_{r \in R} x_{p,r} = 1 \quad \forall p \in \mathbf{A_{t'}} \tag{11}$$

$$\sum_{p \in \mathbf{AP_{tt'}}} x_{p,r} \leq c(r) \quad \forall r \in R, t = 1, \ldots, T \tag{12}$$

$$\sum_{p \in \mathbf{AP^M_{tt'}}} x_{p,r} \geq v_{r,t} \quad \forall r \in R, t = 1, \ldots, T \tag{13}$$

$$\sum_{p \in \mathbf{AP^F_{tt'}}} x_{p,r} \geq v_{r,t} \quad \forall r \in R, t = 1, \ldots, T \tag{14}$$

$$\sum_{p \in \mathbf{AP^M_{tt'}}} x_{p,r} \leq v_{r,t} + c(r) \cdot y_{r,t} \quad \forall r \in R, t = 1, \ldots, T \tag{15}$$

$$\sum_{p \in \mathbf{AP^F_{tt'}}} x_{p,r} \leq v_{r,t} + c(r) \cdot (1 - y_{r,t}) \quad \forall r \in R, t = 1, \ldots, T \tag{16}$$

$$x_{p,r} = 1 \quad \forall p \in \mathbf{A_{t'-1}}, r = \sigma_{t'-1}(p) \tag{17}$$
$$x_{p,r} \in \{0,1\} \quad \forall p \in \mathbf{A_{t'}}, r \in R$$
$$v_{r,t} \geq 0 \quad \forall r \in R, t = 1, \ldots, T$$
$$y_{r,t} \in \{0,1\} \quad \forall r \in R, t = 1, \ldots, T$$

The model describes an assignment that minimizes the expected cost of the newly arrived patients (Expression (10)). Constraint (11) specifies that each arrived patient has to be assigned to a room, while constraint (12) expresses that room capacity should be respected at all times. Constraints (13), (14), (15), and (16) relate the variables $v_{r,t}$ and $y_{r,t}$, forcing $v_{r,t}$ to take on the expected value of the minimum number of either males or females in room r at time t. Finally, constraint (17) ensures that the new assignment respects the assignments of previously arrived patients.

3.2 Model 2

The second model defines the same decision variables as Model 1, however it differs in the set of patients for which they are defined. Whereas the $x_{p,r}$ variables are defined for all arrived patients $\mathbf{A_{t'}}$ in the first model, in this model they are defined for all *registered* patients $\mathbf{P_{t'}}$.

Another difference is that patients can be assigned to a *dummy* room, denoted as $\perp$. Only registered patients who have not arrived are allowed in this dummy room, to ensure feasibility of the model under an expected, future, undercapacity. These assignments are attributed a high cost $c(p, \perp)$ in such a way that the model gives priority to a real assignment for each future arrival.

The model is defined as follows:

$$\text{Min} \sum_{p \in \mathbf{P_{t'}}} \sum_{r \in R \cup \perp} elos(p) \cdot c(p,r) \cdot x_{p,r} + \sum_{r \in R} \sum_{t=1}^{T} w_G \cdot v_{r,t} \tag{18}$$

s.t.

$$\sum_{r \in R} x_{p,r} = 1 \quad \forall p \in \mathbf{A_{t'}} \tag{19}$$
$$\sum_{r \in R \cup \perp} x_{p,r} = 1 \quad \forall p \in \mathbf{P_{t'}} \backslash \mathbf{A_{t'}} \tag{20}$$
$$\sum_{p \in \mathbf{PP_{tt'}}} x_{p,r} \leq c(r) \quad \forall r \in R, t = 1, \ldots, T \tag{21}$$
$$\sum_{p \in \mathbf{PP^M_{tt'}}} x_{p,r} \geq v_{r,t} \quad \forall r \in R, t = 1, \ldots, T \tag{22}$$
$$\sum_{p \in \mathbf{PP^F_{tt'}}} x_{p,r} \geq v_{r,t} \quad \forall r \in R, t = 1, \ldots, T \tag{23}$$
$$\sum_{p \in \mathbf{PP^M_{tt'}}} x_{p,r} \leq v_{r,t} + c(r) \cdot y_{r,t} \quad \forall r \in R, t = 1, \ldots, T \tag{24}$$

instance	$\|P\|$	$\|R\|$	$\sum_{r\in R} c(r)$	avg. occupancy (%)	planning horizon
1	652	98	286	59.69	14
5	587	102	325	49.32	14
8	895	148	441	43.90	21
10	1575	104	308	47.76	56

Table 1: Problem characteristics of the instances.

$$\sum_{p\in \mathbf{PP^F_{tt'}}} x_{p,r} \leq v_{r,t} + c(r)\cdot(1-y_{r,t}) \quad \forall r \in R, t=1,\ldots,T \tag{25}$$

$$x_{p,r} = 1 \quad \forall p \in \mathbf{A_{t'-1}}, r = \sigma_{t'-1}(p) \tag{26}$$

$$x_{p,\perp} = 0 \quad \forall p \in \mathbf{A_{t'}} \tag{27}$$

$$x_{p,r} \in \{0,1\} \quad \forall p \in \mathbf{P_{t'}}, r \in R\cup \perp$$

$$v_{r,t} \geq 0 \quad \forall r \in R, t=1,\ldots,T$$

$$y_{r,t} \in \{0,1\} \quad \forall r \in R, t=1,\ldots,T$$

The objective of the model, expression (18), is again to minimize the total assignment cost, including minimizing any possible dummy assignments. Constraints (19) and (20) specify that each arrived and registered patient should be assigned to one room, allowing for dummy assignments for future arrivals. Constraints (21) - (26) are similar to their counterparts in Model 1, this time also considering future arrivals. Constraint (27) ensures that arrived patients are not assigned to dummy rooms.

4 Experimental setup

We tested the sensitivity of the two models to the following problem characteristics:

- Occupancy
- Accuracy of the length of stay estimate (see further)
- Emergency versus planned cases

For this purpose, we used a subset[2] of the benchmark instances for the static PA problem available from the patient admission scheduling website (Demeester 2012). The instances were extended to the dynamic problem by adding a random registration date $rd(p)$ and an expected departure date $ed(p)$ for each patient p over the planning horizon. See Table 1 for the characteristics of these instances. The procedure for enriching the instances is as follows:

- $ed(p)$ is selected uniformly from the interval $[dd(p) - acc, dd(p) + acc]$ for each patient individually. If $ed(p) <= a(p)$, then it is set to $ed(p) = a(p) + 1$. We investigate the effect of acc, i.e. the effect of the accuracy of the expected departure date estimate.

[2] The subset consists of instances 1,5, 8 and 10.

– $rd(p)$ is either selected uniformly from the interval $[a(p) - T, a(p) - 1]$ for planned patients, or is set to $a(p)$ for emergency patients. We investigate the effect of the percentage (denoted em) emergency versus planned cases.

To test the effect of increasing occupancy, we randomly remove (uniformly selected) beds from the instances to increase the projected average occupancy. This procedure is similar to what Ceschia and Schaerf (2011) did. To maintain feasibility, we limit this increase such that the peak occupancy is never above[3] 100%.

A lower bound for each instance was obtained by calculating the linear relaxation of a straightforward MIP model (not included in this paper) on the static version of the problem (i.e. where everything is known *a priori*). For studying the effect of increasing occupancy, we have calculated the lower bound for every occupancy setting since it increases as beds are removed from the instance. In the figures discussed in the following section, this lower bound is denoted as LB.

The ILP models have been implemented using Gurobi 4.5.2 with a free academic license and were solved with a time limit of 300 seconds per decision step. Thus, for example, an instance with a planning horizon of 14 days (14 decision steps) is solved in at most $14 \times 300 = 4200$ seconds.

All experiments were performed on a computer equipped with a 3.0 GHz Core2Quad processor, and 4 GB of ram memory, running Windows XP Professional (Service Pack 3). The solver was configured to use only one processing thread. The supporting code was implemented in Java 1.6.

5 Discussion

5.1 Emergency versus planned cases, and the effect of the LOS estimate

Both models were tested on all combinations of the factors acc (LOS estimate) and em (percentage emergency cases), with acc ranging from 0 time units (perfect estimate) to 5 time units (a poor estimate) and $em \in \{0, 0.25, 0.50, 0.75, 1.0\}$. All tests were performed on 10 randomized instances for each specified combination of the mentioned factors. The subsequent figures and tables report on the *averages* over these 10 runs.

Figure 1 shows the effect of an increasing percentage of emergencies, under a perfect LOS estimate (left column) and a poor LOS estimate (right column), for instances 1 and 5 (top and bottom row). The results show that Model 2 consistently outperforms Model 1. However, in the limit for increasing percentage of emergencies, the result of Model 2 converges to Model 1. Obviously, in the case for 100% emergency cases, no future arrivals can be planned and the model is reduced to Model 1.

[3] Note however, that hospitals do face a 100% occupancy (and higher, using unlisted beds) from time to time.

For a perfect LOS estimate, Model 1 is not sensitive to the percentage of emergency cases, as it only considers arrived patients, whether they are emergency or planned.

Under a poor LOS estimate, both models show a more erratic behaviour. This appears unrelated to the percentage of emergencies. The reason for this change is that a decision (both for Model 1 and Model 2) may turn out good or bad when patients depart earlier or later than estimated. However, Model 2 still outperforms Model 1 for both good and bad estimates.

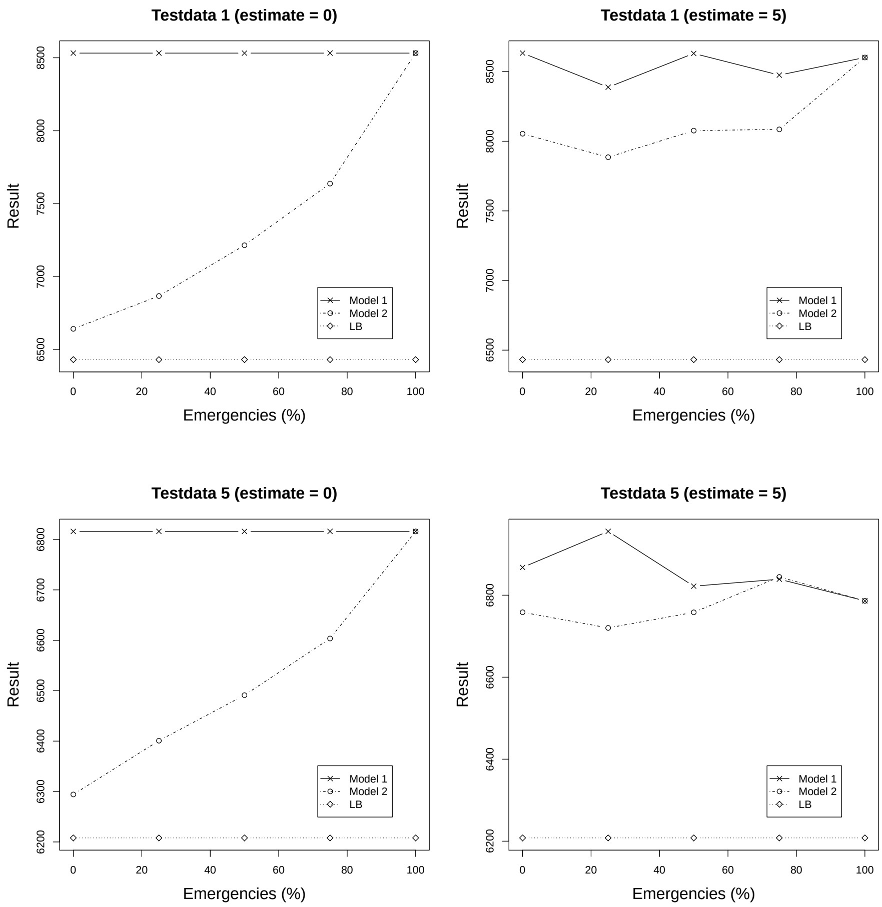

Fig. 1: Model performance for increasing percentage of emergencies, under a perfect LOS estimate (left) and a poor LOS estimate (right). Results shown for instance 1 and 5.

The above conclusions are again confirmed in Figure 2 that shows the effect of the model performance under an increasingly poorer LOS estimate, for a low to high percentage of emergencies. From the results, it follows that the performance of both models deteriorates for an increasingly poorer LOS estimate, while Model 2 always outperforms Model 1. This result was expected, as an increasing inaccuracy of the LOS estimate causes an inaccurate weighing of the patient assignments and thus suboptimal solutions. Furthermore, more patients will (significantly) go over their planned LOS, requiring future arrivals to be replanned. Again, in the limit for 100 % emergency cases, Model 2 converges to Model 1. For an overview of the results, please refer to Table 2.

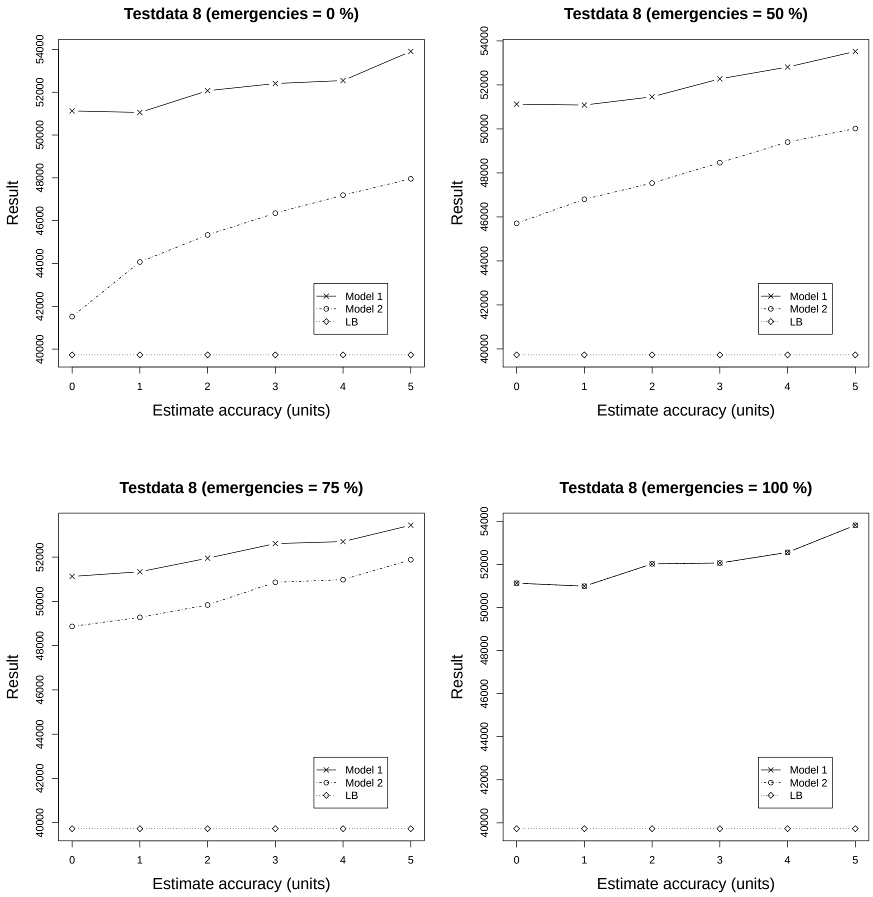

Fig. 2: Model performance for an increasing error on the LOS estimate, under a low percentage of emergencies towards a high percentage of emergencies. Results shown for instance 8.

		Instance 1, increasing *acc*						Instance 5, increasing *acc*					
em (%)	Model	0	1	2	3	4	5	0	1	2	3	4	5
0	1	853.2	815.6	811.1	818.3	851.0	863.3	681.6	673.9	676.9	678.9	685.6	686.8
	2	664.3	705.9	714.2	735.7	760.1	805.4	629.4	649.8	660.4	663.9	669.2	675.8
25	1	853.2	814.2	802.6	816.5	838.4	838.8	681.6	676.2	673.5	685.8	685.5	695.6
	2	686.8	708.3	721.4	729.6	755.6	788.6	640.1	659.8	662.6	668.3	676.2	672.0
50	1	853.2	822.4	820.0	821.4	840.9	863.1	681.6	677.0	676.1	676.4	684.1	682.2
	2	721.5	740.3	744.8	758.9	777.5	807.6	649.1	657.4	667.5	664.5	681.0	675.8
75	1	853.2	817.7	828.8	831.7	847.8	847.5	681.6	679.8	677.6	681.8	690.2	683.9
	2	763.8	766.8	810.0	799.1	821.7	808.5	660.4	666.0	677.0	676.8	677.4	684.4
100	1	853.2	821.0	818.1	810.0	845.9	860.2	681.6	675.6	675.9	677.2	694.2	678.6
	2	853.2	821.0	818.1	810.0	845.9	860.2	681.6	675.6	675.9	677.2	694.2	678.6

		Instance 8, increasing *acc*						Instance 10, increasing *acc*					
em (%)	Model	0	1	2	3	4	5	0	1	2	3	4	5
0	1	5112.6	5105.4	5207.1	5240.5	5254.5	5390.5	10347.0	10329.9	10408.1	10508.8	10688.5	10786.4
	2	4151.5	4407.0	4533.3	4634.9	4719.1	4795.5	8135.1	8705.6	9047.1	9403.4	9688.7	9805.4
25	1	5112.6	5090.4	5207.3	5200.0	5249.7	5390.7	10347.0	10333.0	10443.7	10561.9	10689.0	10782.9
	2	4361.4	4526.9	4610.4	4704.8	4790.6	4898.0	8592.7	8999.7	9299.0	9468.7	9674.5	9893.0
50	1	5112.6	5108.5	5146.0	5227.6	5281.2	5352.2	10347.0	10389.4	10383.9	10548.5	10693.6	10853.0
	2	4571.0	4680.4	4754.0	4846.5	4940.3	5001.8	9386.8	9578.4	9717.7	10003.6	10125.7	10235.2
75	1	5112.6	5133.7	5195.1	5261.0	5270.0	5343.9	10347.0	10347.7	10448.4	10552.5	10626.8	10806.0
	2	4886.9	4927.6	4983.8	5086.2	5097.9	5188.0	9959.4	10093.1	10199.9	10295.8	10442.9	10604.2
100	1	5112.6	5098.7	5202.2	5206.1	5255.6	5381.5	10347.0	10345.0	10427.4	10525.3	10644.7	10743.7
	2	5112.6	5098.7	5202.2	5206.1	5255.6	5381.5	10347.0	10345.0	10427.4	10525.3	10644.7	10743.7

Table 2: Performance comparison of Model 1 and Model 2 under different percentages of emergencies and a worsening LOS estimate. Results shown for instances 1, 5, 8 and 10.

5.2 Effect of increasing occupancy

The effect of an increasing occupancy was tested by artificially forcing a higher, average, occupancy in instances 1 and 5, ranging from 59% to 77% for instance 1 and from 49% to 67% for instance 5. Both instances reach a peak occupancy of 100%. Again, all combinations of factors were tested 10 times to reduce random effects. The following results report on the averages of those 10 runs.

Figure 3 shows the effect of an increasing occupancy on the performance of both models, under an increasing percentage of emergencies (from left-to-right, top-to-bottom) for instance 1. It is clear that both models perform worse under an increasing occupancy. However, the lower bounds of the instances also increase as beds are removed from the instances. Thus, the relative performance of the models compared to the lower bound does not change, indicating that occupancy does not have an effect on what the models can achieve. Figure 4 shows the same effect for instance 5.

6 Conclusion

In the present contribution, a dynamic version of the patient assignment problem that models a day- to- day planning process at hospital admission offices was proposed. The problem definition extended the previous, static, definition to account for the dynamics of *online* patient arrivals, including emergency patients, and explicitly models the LOS of patient as an *estimate*.

Two ILP models were developed: one that is modelled after current practice, namely assigning patients to rooms as they arrive, and one that also accounts for future planned arrivals. The first model improves on current practice by also considering the expected LOS of patients, therefore enabling a proper weighing of patient assignments. The second model also accounts for future, planned, arrivals in order to weigh patient assignments even better.

Experimental results showed that the second model can still be solved efficiently using a commercial MIP solver (under 5 minutes per scheduling step), outperforming the first model as it considers more available information on future arrivals. Furthermore, experimentation with the percentage of emergency patients, poorer LOS estimates and an increasing hospital occupancy indicate that this behaviour does not change under these conditions, advocating the use of model two over model one.

Acknowledgements Research funded by a Ph.D. grant of the Agency for Innovation by Science and Technology (IWT).

References

Beliën J, Demeulemeester E (2007) Building cyclic master surgery schedules with leveled resulting bed occupancy. European Journal of Operational Research 176(2):1185–1204, DOI 10.1016/j.ejor.2005.06.063

Bilgin B, Demeester P, Misir M, Vancroonenburg W, Vanden Berghe G (2012) One hyper-heuristic approach to two timetabling problems in health care. Journal of Heuristics pp 1–34, DOI 10.1007/s10732-011-9192-0

Burke EK, De Causmaecker P, Vanden Berghe G, Van Landeghem H (2004) The state of the art of nurse rostering. Journal of Scheduling 7:441–499, DOI 10.1023/B:JOSH.0000046076.75950.0b

Cardoen B, Demeulemeester E, Beliën J (2010) Operating room planning and scheduling: A literature review. European Journal of Operational Research 201(3):921–932, DOI 10.1016/j.ejor.2009.04.011

Ceschia S, Schaerf A (2011) Local search and lower bounds for the patient admission scheduling problem. Computers & Operations Research 38(10):1452–1463, DOI 10.1016/j.cor.2011.01.007

Demeester P (2012) Patient admission scheduling problem website. Online, URL `http://allserv.kahosl.be/~peter/pas/`

Demeester P, Souffriau W, De Causmaecker P, Vanden Berghe G (2010) A hybrid tabu search algorithm for automatically assigning patients to beds. Artificial Intelligence in Medicine 48(1):61–70, DOI 10.1016/j.artmed.2009.09.001

Kolen AWJ, Lenstra JK, Papadimitriou CH, Spieksma FCR (2007) Interval scheduling: A survey. Naval Research Logistics (NRL) 54(5):530–543, DOI 10.1002/nav.20231

Ouelhadj D, Petrovic S (2009) A survey of dynamic scheduling in manufacturing systems. Journal of Scheduling 12:417–431, DOI 10.1007/s10951-008-0090-8

Rais A, Viana A (2011) Operations research in healthcare: a survey. International Transactions in Operational Research 18(1):1–31, DOI 10.1111/j.1475-3995.2010.00767.x

Riise A, Burke EK (2010) Local search for the surgery admission planning problem. Journal of Heuristics 17(4):389–414, DOI 10.1007/s10732-010-9139-x

Vancroonenburg W, Goossens D, Spieksma FCR (2011) On the complexity of the patient assignment problem. Tech. rep., KAHO Sint-Lieven, Gebroeders De Smetstraat 1, Gent, Belgium, URL `http://allserv.kahosl.be/~wimvc/pas-complexity-techreport.pdf`

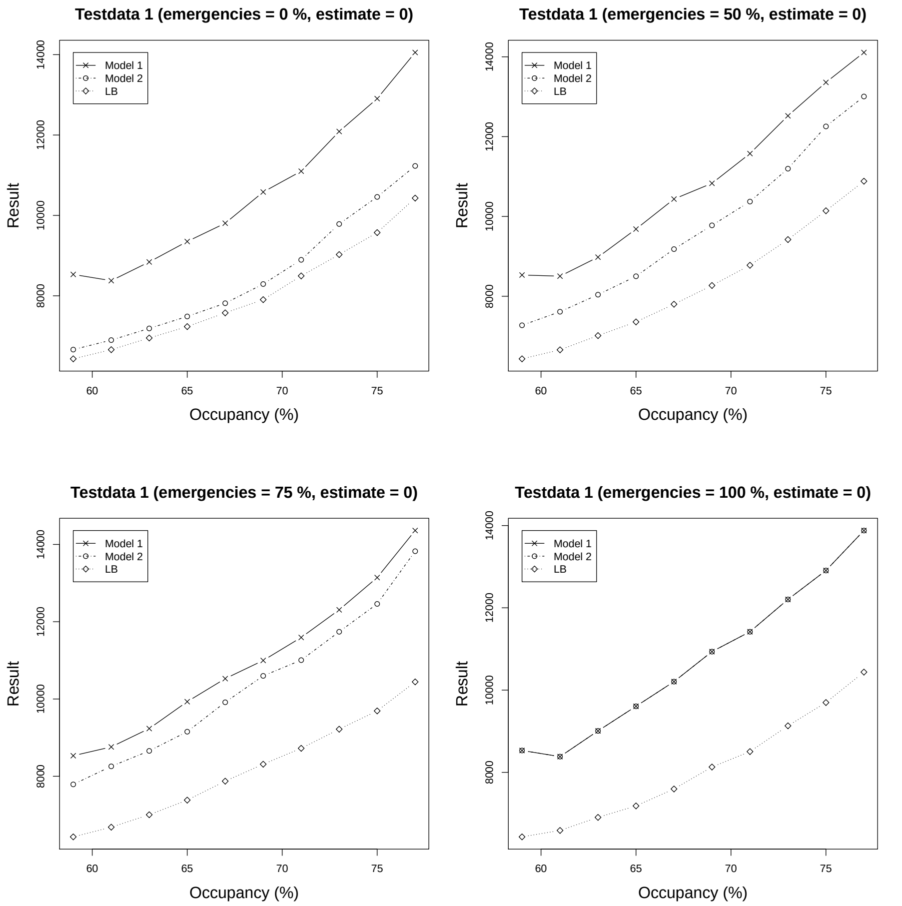

Fig. 3: Model performance for an increasingly higher occupancy rate, under different levels of emergency vs planned patients. Results shown for instance 1 with a perfect estimate.

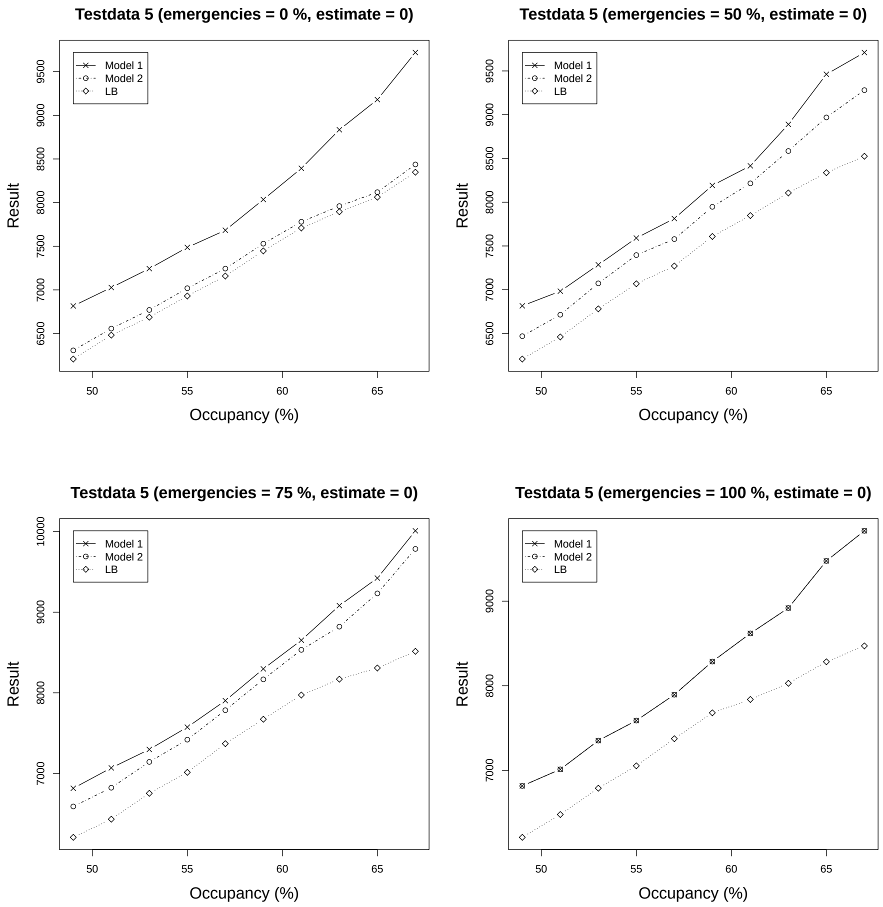

Fig. 4: Model performance for an increasingly higher occupancy rate, under different levels of emergency vs planned patients. Results shown for instance 5 with a perfect estimate.