



**An Adaptive Large Neighbourhood Search Procedure  
Applied to the Dynamic Patient Admission Scheduling  
Problem**

by

**Richard Martin Lusby, Martin Schwierz, Troels Martin Range and Jesper Larsen**

Discussion Papers on Business and Economics  
No. 1/2016

FURTHER INFORMATION  
Department of Business and Economics  
Faculty of Business and Social Sciences  
University of Southern Denmark  
Campusvej 55, DK-5230 Odense M  
Denmark

E-mail: [lho@sam.sdu.dk](mailto:lho@sam.sdu.dk) / <http://www.sdu.dk/ivoe>

# An Adaptive Large Neighbourhood Search Procedure applied to the Dynamic Patient Admission Scheduling Problem

Richard Martin Lusby<sup>1</sup>, Martin Schwierz<sup>2</sup>, Troels Martin Range<sup>3</sup>, and Jesper  
Larsen<sup>1</sup>

<sup>1</sup>Department of Engineering Management, Technical University of Denmark, Kgs.  
Lyngby, Denmark

<sup>2</sup>AMCS, Copenhagen, Denmark

<sup>3</sup>Department of Business and Economics, COHERE, University of Southern  
Denmark, Odense

March 14, 2016

## Abstract

The Patient Admission Scheduling problem involves assigning a set of patients to hospital beds over a given time horizon in such a way that several quality measures reflecting patient comfort, treatment efficiency, and hospital utilization are maximized. Usually it is assumed that all information regarding each patient is known in advance, making it possible to solve a static, offline planning problem. Such an approach, however, often has shortcomings in practice given the dynamic setting in which hospitals operate. An extension of this problem, known as the Dynamic patient Admission Scheduling problem, better reflects reality by attempting to capture, among other things, uncertainty in the length of patient stays as well as the ability to consider emergency patients. In this paper we devise an Adaptive Large Neighbourhood Search procedure, utilizing a Simulated Annealing framework, for this new variant of the problem and test its performance on a set of 450 publicly available problem instances of different size. A comparison with the current state-of-the-art indicates that the proposed methodology provides solutions that are of comparable quality for small and medium sized instances, but in a much shorter time frame. For larger instances the improvement in solution quality is dramatic, approximately 3-14% on average. In such cases, it does, however, take slightly longer.

*Keywords:* Metaheuristic, ALNS, OR in health services, Scheduling

*JEL code:* C61, *MSC 2010 code:* 90B50, 90C27

## 1 Introduction

Providing high quality health care is one of the most important challenges worldwide. To efficiently, and perhaps even safely, manage the day-to-day activities of a hospital requires a thorough

understanding of the hospital’s patient flow. In many countries, the admission of patients to hospitals is administered from a central unit, its responsibility being to facilitate the best possible assignment of patients to beds in the hospital. This is indeed a non-trivial task. No two patients are identical. Each patient’s expected length of stay at the hospital, not to mention the specific treatment they require, is likely to be different. Furthermore, in addition to having a finite capacity, a hospital only has a finite set of resources with which all required services can be performed. Coordinating the diverse set of patient requests with the hospital’s available resources is thus an important, combinatorial optimization problem and is termed the Patient Admission Scheduling (PAS) problem. Despite its complexity, however, the problem, is typically still solved manually by an experienced nurse or planner.

Patient admission scheduling is a problem that arises at all planning levels within a hospital. On a strategic, or long term planning, level, historical data is used to forecast arrival patterns of patients and determines whether the hospital’s available resources and their configuration will perform satisfactorily in the future. Examples of research addressing strategic level variants can be found in Hutzschenreuter et al. (2008) and Chen et al. (2010). The former presents a simulation tool that is used to find good, long term policies, where the aim is to maximize the yearly throughput of the hospital. The latter develops a genetic algorithm for an admission scheduling problem for a single department. This algorithm exploits historical data and optimizes a long-term admission strategy. Other examples include, among others, Kusters and Groot (1996), Harper (2002), and Jittamai and Kangwansura (2011).

The PAS problem also arises at the tactical planning level. This variant of the problem considers a medium to long term planning horizon, i.e. from a few weeks to some months ahead of time, and attempts to assign upcoming patients to beds in such a way that several quality measures reflecting patient comfort, treatment efficiency, and hospital utilization are maximized. Almost always is the data assumed to be known with certainty, i.e. a *static* offline problem is solved. As such, it typically only considers elective patients, whose admission and discharge dates are stated exactly; doctors can, in the absence of complications, usually state with a high degree of accuracy the length of stay needed by a patient for a given treatment. Emergency patients, i.e. those that arrive unexpectedly, are normally implicitly considered by simply reserving a certain number of beds. The work by Demeester et al. (2010) formalized an academic version of this problem, and was accompanied with 14 publicly available data sets, which are based on real-life instances of the problem. The authors proposed a hybrid tabu search algorithm to solve this problem. The same instances have subsequently been considered by Ceschia and Schaerf (2011), Bilgin et al. (2012), and Range et al. (2014). The first two contributions also consider local search based heuristics, while the last describes a column generation based heuristic approach combined with constraint aggregation. The general consensus is that exact methods are only tractable for small problem instances (a two week horizon with around at most 150 rooms and 780 patients).

Due to the dynamic setting in which hospitals operate, it is extremely unlikely that a tactical level plan will be executed unchanged in practice. Modifications will naturally be required whenever unforeseen events occur. In an effort to better reflect the real-life characteristics of the problem more closely, Vancroonenburg et al. (2012), and Ceschia and Schaerf (2012) have independently introduced operational level variants of the problem, collectively known as the Dynamic Patient Admission Scheduling (DPAS) problem *under uncertainty* (henceforth referred to as sim-

ply the DPAS problem). Both studies extend the static PAS problem presented in Demeester et al. (2010) in similar ways and attempt to capture several practical uncertainties encountered in the day-to-day running of a hospital. In particular, the assumption that all patient admission dates are known a priori is no longer made; these dates sporadically reveal themselves over time. In other words, associated with each patient is a *registration date*. This is the time at which the patient’s planned admission date becomes known to the hospital and it does not necessarily need to be the start of the planning horizon. Other dynamic extensions include the possibility to delay a patient’s expected admission date, the possibility for patients to stay longer than expected, and emergency patients. Ceschia and Schaerf (2012) test and compare an exact Integer Linear Programming (ILP) formulation, solved using the commercial solver Cplex, and a Simulated Annealing (SA) approach on a set of 450 self-generated instances, which have since been made public. Results suggest the exact approach is only tractable for small instances of the problem, while the metaheuristic approach can at least find feasible solutions for all instance sizes. As no comparison is possible for the larger instances, it is, however, difficult to vouch for the quality of the solutions obtained in these cases. Vancroonenburg et al. (2012) present two ILPs for dynamic versions of the six smallest data sets given in Demeester et al. (2010), and consider the impact of emergency patients and patient length of stay estimates. More recently, Ceschia and Schaerf (2014) have extended the work of Ceschia and Schaerf (2012) to also consider the scheduling of operating theaters.

In this paper we also focus on the DPAS problem and devise an Adaptive Large Neighborhood Search (ALNS) procedure embedded within a SA framework to solve it. We concern ourselves with the variant proposed by Ceschia and Schaerf (2012). For this there is a large set of diverse instances on which the proposed methodology can be tested. Furthermore, through a direct comparison with the results of Ceschia and Schaerf (2012), not only can we say something about the performance of our algorithm, but we can also comment on the quality of the solutions found in Ceschia and Schaerf (2012). An exhaustive computational study on the 450 available instances shows the proposed methodology is extremely competitive. On small instances solutions of similar quality are obtained, much faster. On medium to large instances we are consistently better. For the large instances the difference in solution quality is sometimes dramatic; the algorithm finds better solutions in all cases, and is, on average, 3-14% better. Such results, to some degree, benchmark the quality of the solutions found in Ceschia and Schaerf (2012).

An outline of this paper is as follows. In Section 2 we introduce some required terminology and provide a more formal definition of the problem at hand. Section 3 gives an overview of the proposed ALNS procedure, while Section 4 details our computational study and, in particular, the comparison with Ceschia and Schaerf (2012). Finally, the main conclusions from this study are drawn in Section 5.

## 2 Problem Definition

In this section we formalize the definition of the DPAS problem we consider. This follows the original problem description given in Ceschia and Schaerf (2012). In any instance of the DPAS problem it is assumed that there is a set of patients  $\mathcal{P} = \{1, 2, \dots, P\}$  requiring admission to a hospital over a pre-specified time horizon of known length. This patient set is partitioned into a set

of male patients,  $\mathcal{M}$ , and a set of female patients,  $\mathcal{F}$ . The planning horizon is divided into a set of consecutive days, denoted by the set  $\mathcal{D} = \{1, 2, \dots, D\}$ , and indicates for how long the admission schedule is required. Associated with each patient  $p \in \mathcal{P}$  are several important dates, and these are known in advance. The first,  $d_p^{reg} \in \mathcal{D}$ , states the day on which patient  $p$  is registered in the hospital’s administrative records. Although a patient’s registration date is known in advance, knowledge of this cannot be used prior to this date. The second,  $d_p^{plan} \in \mathcal{D}$ , indicates the day on which patient  $p$  is expected to be admitted to the hospital. Finally,  $d_p^{max} \in \mathcal{D}$ , gives the the last possible day patient  $p$  can be admitted. Obviously, by setting  $d_p^{plan} = d_p^{reg}$ , the DPAS problem can easily incorporate emergency patients. Additionally, a patient’s age, required treatment(s), length of stay,  $l_p$  (given in days), and whether the patient possesses a so-called *overstay risk* are known. The last of these indicates that there is the risk that the patient will stay *one* day longer. If a patient has an overstay risk, then it is desirable for their assigned room to have an empty bed on their planned discharge date. The problem does not consider overstays of arbitrary length.

A hospital is assumed to consist of several departments, where each department itself further consists of several rooms. Departments typically specialize in the treatment of one or more illnesses. We denote the set of all rooms available at the hospital as  $\mathcal{R} = \{1, 2, \dots, R\}$ . Each room  $r \in \mathcal{R}$  has a certain capacity  $\kappa_r$ ; this indicates the number of beds in the room. Room capacity is typically one, two, or four beds. Each room is assumed to be equipped with set of features. Features indicate the presence of specific equipment or properties, which can be necessary (or only preferable) for patients. Finally, each room is also assumed to have a gender policy. There are four possible gender policies, and these are listed in Table 1.

Policy	Description
SG	Both genders are allowed, but not at the same time.
All	There is no gender restriction.
Ma	Only male patients are allowed.
Fe	Only female patients are allowed.

Table 1: The available gender policies

The same gender policy (SG) is the most used policy in many hospitals. An illustration of a patient admission schedule is given in Figure 1. For this fictitious hospital, the rooms in department 1 have room policy (SG), department 2 has the (All) policy, department 3 has the (Ma) policy, and finally department 4 has the (Fe) policy. Note that the gender policy does not necessarily have to be the same for all rooms of a department. The SG policy rooms are solution wise of interest and we let  $\mathcal{R}^{SG}$  be the subset of rooms having the SG policy.

When finding a solution to the DPAS problem a variety of constraints, some hard and some soft, exist. Hard constraints must be respected and concern things like room capacity, patient age, department specialism, and required room features. Obviously, we cannot assign more patients to a room than the room has beds for. A patient must be treated in a department that is consistent with the patient’s age. Similarly, a patient must be treated in a department which has the specialism the patient requires, and cannot be assigned to a room not equipped with properties the patient needs.

Soft constraints, on the other hand, are used to model preferences or desirable properties of

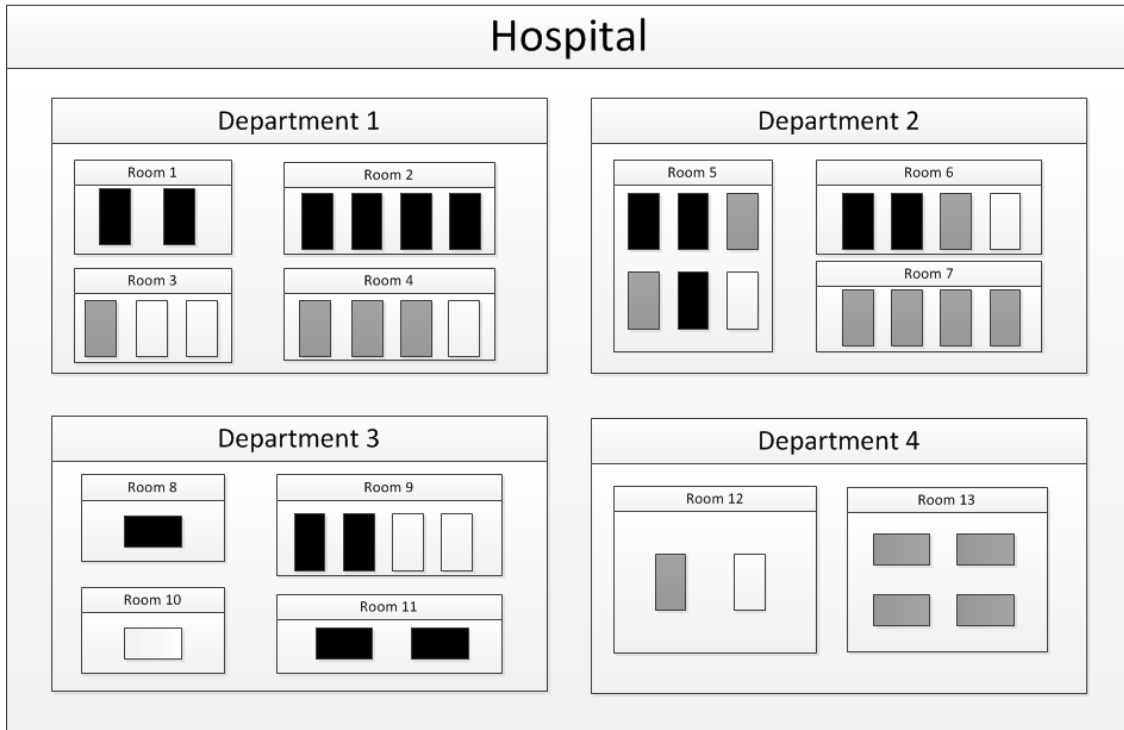


Figure 1: Example schedule for a hospital. Black beds are assigned to male patients, gray to female patients. White beds are empty

the solution. Ideally, they should be enforced, but it is not necessary; they can be violated at a cost. These include things such as: department main specialism, preferred room feature, transfers, gender, room size, admission delay, and overstay risk. To elaborate, ideally a patient should be assigned a department whose main specialism is in the treatment the patient needs; preferred properties for a patient should be in the room they receive; the transferring of a patient between rooms should be avoided; the patients assigned to a room should be consistent with the room's gender policy; a patient should receive a room with capacity in line with their preferences; and if any of the patients assigned to a room have an overstay risk, the room should have an available bed on their expected discharge date (just in case). Table 2 summarizes the different constraints of the problem, their type and, if it is a soft constraint, the penalty incurred when violating the given constraint once. These values are the same as for Ceschia and Schaerf (2012). The highest weight is given to transferring (Tr) the patient. This is in line with common practice, where transfers are really only used as a last option to maintain or gain feasibility. The two last constraints, (De) and (OR), are extensions of the static version of the problem. The first of these penalizes delaying a patient's admission to the hospital; for the admission date this can, however, not be more than is  $d_p^{max}$ . The second concerns the overstay risk and states that a penalty is incurred if a patient, who has the potential to stay longer, cannot be assigned the room they were in on their discharge date for capacity reasons. Such a situation would result in so-called *overcrowding*.

Given the constraints defined above, the objective of the DPAS problem is to find the assignment of patients to rooms that minimizes the sum of the penalties incurred from soft constraint

Constraints	Abbreviation	Type	Weight
Room capacity	RC	hard	-
Patient age	PA	hard	-
Department specialism	DS1	hard	-
Department main specialism	DS2	soft	20
Needed room feature	RF1	hard	-
Preferred room feature	RF2	soft	20
Transfer	Tr	soft	100
Gender	Gr	soft	50
Room size preference	RS	soft	10
Admission delay	De	soft	2
Overcrowd risk	OR	soft	1

Table 2: Constraints and their weight

violations. Due to the fact that the patient’s age and required treatment are known beforehand, and as these are hard constraints, it is possible to eliminate from consideration rooms that do not conform to these requirements for the specific patient.

### 3 Solution Approach

The size and complexity of the DPAS problem makes the application of exact mathematical programming procedures difficult. Several ILP based exact procedures have been proposed in the literature for both the static and dynamic variants of the PAS problem, and all have concluded that only small instances of the problem are tractable with such approaches, see e.g. Ceschia and Schaefer (2011, 2012); Range et al. (2014). Metaheuristic approaches are therefore an attractive alternative as they can overcome such tractability issues. Being heuristic in nature they are, however, unable to independently provide a certificate of quality on the obtained solution. Nonetheless, they make it possible to consider large problem instances and, if designed well, can produce good solutions. In this section we present and describe such an approach for solving the DPAS problem. The developed methodology repeatedly applies a SA-based ALNS algorithm throughout the planning horizon. To aid the exposition, this section has been divided into four parts. To have a concise description of the problem we are solving we begin with a mathematical model in Section 3.1. In Section 3.2 we continue with a short discussion on the dynamic nature of the problem. Section 3.3 provides a general description of a SA-based ALNS algorithm, while Section 3.4 elaborates on the application of the technique to the DPAS problem.

#### 3.1 Mathematical model

We introduce a mathematical formulation of the DPAS problem, not to be able to solve it as such, but rather as a means to state the problem precisely. The mathematical formulation presented is a static model including all patients in the solution, and it does not exclude patients who have not been registered yet. It is, however, possible to apply this to the dynamic setting where some variables are then fixed. We describe this application of the model in Section 3.2.

Each patient  $p \in \mathcal{P}$  has a window of dates in which (s)he can be admitted. We denote this window as  $\mathcal{D}_p = \{d_p^{plan}, \dots, d_p^{max}\} \subset \mathcal{D}$ . A patient  $p \in \mathcal{P}$  can be allocated to a room  $r \in \mathcal{R}$  if the hard constraints PA, DS1, and RF1 are satisfied for the patient. We let  $\mathcal{R}_p$  be the set of feasible rooms for patient  $p \in \mathcal{P}$ . For each patient  $p \in \mathcal{P}$  and room  $r \in \mathcal{R}$  we denote the daily cost of admitting the patient into the room as  $c_{pr}$ . This cost corresponds to the sum of the soft constraint penalties DS2, RF2, RS, and Gr (for the policies Ma and Fe) from Table 2, if any of these are violated when assigning patient  $p$  to room  $r \in \mathcal{R}_p$ . If  $r \notin \mathcal{R}_p$ , then we put  $c_{pr} = \infty$ . This cost is independent of the day on which the patient occupies the room. We use  $c^D$  to denote the cost per day of delaying a patient. This corresponds to the penalty incurred when violating soft constraint De from Table 2. The parameter  $O_p \in \{0, 1\}$  is used to indicate whether or not patient  $p \in \mathcal{P}$  has a risk of staying an additional day, while  $c^O$  denotes the unit overcrowd cost for having potentially more patients in the room than the room's capacity. This unit overcrowd penalty is associated with violation of soft constraint OR in Table 2. Finally,  $c^T$  denotes the cost of a transfer and corresponds to violation of soft constraint Tr from Table 2.

Binary decision variables  $x_{prd} \in \{0, 1\}$  are used to indicate whether or not patient  $p \in \mathcal{P}$  occupies a bed in room  $r \in \mathcal{R}$  on day  $d \in \mathcal{D}$ . An auxiliary set of variables  $\bar{x}_{prd} \in \{0, 1\}$  indicate that the patient  $p$  may occupy a bed in room  $r$  on day  $d$  after discharge due to the overstay risk, and this set is used to calculate the potential overcrowding of a room. In addition, binary variables  $t_{pd}$  indicate whether or not patient  $p \in \mathcal{P}$  is transferred from one room to another, from day  $d - 1$  to day  $d$ . The day of admission for patient  $p \in \mathcal{P}$  is controlled by the binary variable  $y_{pd} \in \{0, 1\}$ ; this equals one if and only if patient  $p \in \mathcal{P}$  has day  $d \in \mathcal{D}$  as their day of admission. For each room and each day we let the binary variables  $f_{rd}, m_{rd} \in \{0, 1\}$  indicate whether or not at least one female or at least one male is present in the room, respectively. The binary variable  $b_{rd} \in \{0, 1\}$  indicates that both genders are present in the room  $r \in \mathcal{R}$  on day  $d \in \mathcal{D}$ . Finally,  $z_{rd} \geq 0$  measures the potential number of patients beyond the capacity of the room  $r \in \mathcal{R}$  on day  $d \in \mathcal{D}$ .

Some of the variables can be fixed at zero. As it is only possible to admit a patient  $p$  during the days in  $\mathcal{D}_p$ ,  $y_{pd} = 0$  for all  $d \notin \mathcal{D}_p$ . Likewise,  $x_{prd} = 0$  if  $r \notin \mathcal{R}_p$  or if the day  $d$  is either before  $d_p^{plan}$  or after  $d_p^{max} + l_p$ .

For a solution to be feasible it has to satisfy the set of hard constraints. These are listed below. First of all, the rooms can only accommodate as many patients as the room has beds for:

$$\sum_{p \in \mathcal{P}} x_{prd} \leq \kappa_r, \quad r \in \mathcal{R}, d \in \mathcal{D}. \quad (1)$$

The individual patients have to be admitted on precisely one day within the planning horizon:

$$\sum_{d \in \mathcal{D}_p} y_{pd} = 1, \quad p \in \mathcal{P}. \quad (2)$$

If patient  $p$  is admitted on day  $\bar{d}$ , the patient must appear in a room the following  $l_p$  days

$$\sum_{r \in \mathcal{R}_p} x_{prd} \geq y_{p\bar{d}}, \quad p \in \mathcal{P}, \bar{d} \in \mathcal{D}_p, d = \bar{d}, \dots, \bar{d} + l_p. \quad (3)$$



Constraints (1)-(3) constitute the core of the model. The remaining constraints serve different purposes. These include indicating the presence of male and female patients in rooms (to calculate the violation of the gender constraints in rooms with the SG policy), identifying when transfers happen, and calculating the potential overcrowding of rooms.

The first set of constraints is used to segregate genders in rooms having the SG gender policy. The presence of a female patient is determined by

$$f_{rd} \geq x_{prd}, \quad p \in \mathcal{F}, r \in \mathcal{R}_p \cap \mathcal{R}^{SG}, d \in \mathcal{D}_p, \quad (4)$$

while the presence of a male patient is determined by

$$m_{rd} \geq x_{prd}, \quad p \in \mathcal{M}, r \in \mathcal{R}_p \cap \mathcal{R}^{SG}, d \in \mathcal{D}_p. \quad (5)$$

This allows us to calculate when both genders are present using the following constraint

$$b_{rd} \geq f_{rd} + m_{rd} - 1, \quad r \in \mathcal{R}^{SG}, d \in \mathcal{D}. \quad (6)$$

Thus,  $b_{rd}$  is calculated for all same-gender policy rooms on all days. The transfer of patients is enforced using the following constraint:

$$t_{pd} \geq x_{prd} - x_{pr,d-1} - y_{pd}, \quad p \in \mathcal{P}, r \in \mathcal{R}_p, d = 2, \dots, D. \quad (7)$$

In other words, a transfer is recorded if patient  $p \in \mathcal{P}$  does not stay in the same  $r \in \mathcal{R}$  on consecutive stays during their stay. Without the  $y_{pd}$  term the constraint would also record a transfer when admitting patient  $p$  to the hospital (since the patient would not have been assigned the same room on the previous day).

If a patient  $p$  has an overstay risk, i.e  $O_p = 1$ , then the patient potentially occupies a bed for an extra day in the last room of their planned stay. This is modelled as follows:

$$y_{pd} + x_{(p,r,d+l_p)} O_p \leq 1 + \bar{x}_{(p,r,d+l_p+1)}, \quad p \in \mathcal{P}, r \in \mathcal{R}_p, d \in \mathcal{D}_p. \quad (8)$$

Having settled the potential occupation of a bed in room  $r$  the day after the planned discharge, the potential overcrowding of any room can be recorded with the following constraint:

$$\sum_{p \in \mathcal{P}} (x_{prd} + \bar{x}_{prd}) \leq c_r + z_{rd}, \quad r \in \mathcal{R}, d \in \mathcal{D}. \quad (9)$$

The objective consists of five terms. The first measures the time-independent penalty of assigning a patient to a room for one night. The second term calculates the cost of all transfers of patients, while the third term determines the penalty given for gender violations of rooms having the same-gender policy. The fourth term penalizes the potential overcrowding, and finally the fifth

term penalizes the admission delay of patients. The objective can be stated as:

$$\begin{aligned} \text{minimize} \quad & \sum_{p \in \mathcal{P}} \sum_{d \in \mathcal{D}_p} \sum_{r \in \mathcal{R}_p} c_{pr} x_{prd} + \sum_{p \in \mathcal{P}} \sum_{d \in \mathcal{D}} c^T t_{pd} + \sum_{d \in \mathcal{D}} \sum_{r \in \mathcal{R}^{SG}} c^G b_{rd} \\ & + \sum_{d \in \mathcal{D}} \sum_{r \in \mathcal{R}} c^O z_{rd} + \sum_{p \in \mathcal{P}} \left( \sum_{d \in \mathcal{D}_p} dy_{pd} - d_p^{plan} \right) c^D \end{aligned} \quad (10)$$

The overall aim of the problem is to minimize the penalties incurred from violations of soft constraints, while satisfying the constraints (1)-(9) and the domains of the variables.

The model presented above corresponds to the problem verbally stated by Ceschia and Schaerf (2012), though it is slightly different to the model the authors present. The most obvious difference is that our model allows transfers due to patients having the possibility of staying in different rooms during their stay. Another difference lies in the calculation of the overcrowd penalty. Our model penalizes per unit of potential overcrowd, while the model presented by Ceschia and Schaerf (2012) just penalizes the presence of potential overcrowd. In practice this model suffers from severe degeneracy and standard integer linear programming solvers are typically not able to solve the problem in reasonable time – see e.g. Ceschia and Schaerf (2011).

### 3.2 The Dynamic Setup

The DPAS problem is a dynamic problem and as such it calls for continual optimization based on the available information. We assume that the information on registered patients is made available once each day, after which an optimization problem is solved to identify which rooms patients will be assigned to in the future. Patients who are already admitted can only stay in the room from the day before or be moved to other rooms, resulting in a transfer, and naturally they cannot be postponed. Other registered patients can freely be moved by the optimization approach. Formally, on each day  $d \in \mathcal{D}$  we must optimize the assignment of the set of patients

$$\mathcal{P}_d = \{p \in \mathcal{P} | d_p^{reg} \leq d, p \text{ not discharged}\}$$

to rooms. We refer to  $\mathcal{P}_d$  as the set of *active* patients on day  $d$  and denote this as  $A := |\mathcal{P}_d|$ . Because each patient has a last possible day of admission, the planning horizon for the optimization on day  $d$  is given by  $\max\{d_p^{max} | p \in \mathcal{P}_d\}$ .

Suppose that we have solved all the days until day  $d - 1$  and want to solve the problem for day  $d$ . In terms of the model from Section 3.1 we exclude all patients (and corresponding variables) from  $\mathcal{P} \setminus \mathcal{P}_d$ , as they have been discharged or have not been registered yet. We obtain the solution from the day before and fix all patients  $p \in \mathcal{P}_{d-1} \cap \mathcal{P}_d$  with  $y_{pd'} = 1$  for a day  $d' \leq d - 1$  to be admitted on day  $d'$ . For these patients we also fix the rooms in which they have stayed up until day  $d - 1$ . For patients who are not fixed we state that their earliest day of admission will be  $\max\{d, d_p^{plan}\}$ . Fixing these variables corresponds to setting up boundary conditions for the optimization of day  $d$  which depends on the irreversible decisions made on previous days.

When optimizing the assignment of patients  $\mathcal{P}_d$  on a given day  $d \in \mathcal{D}$ , we must first construct an initial solution where newly registered patients, i.e. patients having  $d_p^{reg} = d$ , are included. To do this we use a simple greedy insertion algorithm. For the first day there are no previous

daily assignments we must respect; however, for all subsequent days, we must respect the previous day’s final patient to room assignment for the already admitted patients when constructing the solution. This greedy insertion heuristic is explained in more details in Section 3.4 as it is used extensively in the metaheuristic framework.

It is important to note that this constructive heuristic does not guarantee feasibility; a patient’s admission date cannot be delayed indefinitely if infeasibility persists. The admission date of patient  $p \in \mathcal{P}$  can only be delayed up until, and including, either  $d_p^{\max}$ , or the end of the planning horizon, whichever is the earlier. When this does occur, the search for a compatible room for the patient is postponed until the algorithm has assigned all other patients. The admission date of the patient is then set to a date which exceeds the current day and for which a compatible room can be found. This results in an infeasible solution, violating either the requirement that a patient be admitted before his/her maximal admission date or the requirement that there is a fixed number of days in the planning horizon. To be comparable to Ceschia and Schaerf (2014), these infeasibilities receive an extra penalty of 200 per occurrence to ensure that they are quickly remedied during the optimization process.

### 3.3 Simulated Annealing Based Adaptive Large Neighborhood Search

The main principle of local search is to move from one feasible solution  $x$  to another feasible solution  $x'$  by applying small changes. The set of these small changes is denoted a neighborhood. As long as the neighboring feasible solution is better we accept it and continue the process to try to construct an even better feasible solution. The downside with this approach is that a solution that is only locally optimal is ultimately reached, and this halts the exploration of better solutions. Metaheuristics attempt to overcome local optima by allowing the search to move to worse solutions in an attempt to introduce some diversification. One well known approach for this is the SA approach described originally in Metropolis et al. (1953), but first adopted for optimization by Kirkpatrick et al. (1983). In a SA metaheuristic better solutions are always accepted; however, worse solutions are also accepted with a certain probability. This acceptance probability is continually decreased during the execution of the algorithm, ensuring that worse solutions are less likely to be accepted late in the process.

The process of decreasing the acceptance probability is often described as lowering the temperature of an annealing process (hence the name). We start out from an initial temperature  $T := T_{start}$  and then decrease this temperature by putting  $T := \alpha T$  after each iteration, where  $\alpha \in (0, 1)$ . The acceptance probability of a new but worse solution  $x'$  is then calculated as

$$\rho_{acc} := \exp\left(-\frac{f(x') - f(x)}{T}\right)$$

The algorithm terminates when the temperature reaches a lower temperature  $T_{end}$ . For a more detailed description on the practical implementation of SA refer the reader to Aarts et al. (2005).

Instead of making minor changes in moving from one solution to the next, it is possible to make more dramatic changes by destroying a significant part of the solution and then repairing it again by applying a construction heuristic. This is the fundamental idea of so-called *destroy and repair* methods and Large Neighborhood Search (LNS), see Shaw (1998). Destroying a part

of the solution and then repairing it again tremendously enlarges a solution's neighborhood. This prohibits a full search of all neighboring solutions and consequently a SA framework is applied.

Pisinger and Ropke (2006) observed that not all destroy or repair methods remain equally good throughout the course of an SA algorithm. In other words, some tend to improve the solution more in the beginning of the solution process, while others are better suited towards the end. The authors therefore suggest to assign modifiable weights to each of the destroy methods as well as weights to each of the repair methods. These weights are then continually updated over the course of the algorithm and are decided using an exponential smoothing of the performance of the individual methods. This gives rise to the so-called ALNS metaheuristic, which *adapts* its selection of the destroy and repair methods based on the respective performances of each of them.

The set of all such destroy and repair methods is given as  $H = I \cup J$ , where  $I$  is the set of repair methods and  $J$  is the set of destroy methods. For each method  $h \in H$  a weight  $w_h$  is defined and is based on the performance of the method; the better it performs the higher its weight. Well-performing methods should have a higher probability of being chosen than worse-performing methods. The probability of choosing a repair method  $h \in I$  is given as

$$\pi_h := \frac{w_h}{\sum_{i \in I} w_i}, \quad h \in I,$$

while the probability of choosing a destroy method  $h \in J$  is denoted

$$\pi_h := \frac{w_h}{\sum_{j \in J} w_j}, \quad h \in J.$$

The performance of the methods is based on observing the changes in objective value over a given number of iterations,  $n_{iter}$ . For each method  $s_h$  defines the accumulated score for method  $h \in H$ , and  $v_h$  states the number times a method has been used. For an iteration where heuristic  $h$  is chosen to move from solution  $x$  to solution  $x'$  the following is added to  $s_h$

$$\sigma := \begin{cases} \sigma_1, & f(x') < f(x^*) \\ \sigma_2, & f(x^*) \leq f(x') < f(x) \\ \sigma_3, & f(x) \leq f(x') \text{ and } x' \text{ accepted} \\ 0, & \text{otherwise} \end{cases}$$

where  $\sigma_1 \geq \sigma_2 \geq \sigma_3 \geq 0$ . Note that  $f(x)$  gives the objective value of solution  $x$ , and that  $x^*$  is the best found solution so far. In the case where a new best solution is found  $\sigma_1$  is added, while in the case where an improvement over the current solution is found,  $\sigma_2$  is added. If the neighboring solution is not an improvement but is still accepted, then  $\sigma_3$  is added to the score. After  $n_{iter}$  iterations the weights are updated as follows

$$w_h := \begin{cases} (1 - \gamma)w_h + \gamma \frac{s_h}{v_h}, & \text{if } v_h > 0 \\ w_h, & \text{if } v_h = 0 \end{cases}$$

after which  $s_h$  and  $v_h$  are reset to zero. The reaction factor,  $\gamma \in [0, 1]$ , indicates how much emphasis is put on the score of the most recent  $n_{iter}$  iterations. If  $\gamma = 0$ , then the weights are kept constant at their initial levels, while  $\gamma = 1$  only uses the score of the most recent  $n_{iter}$  iterations.

### 3.4 Application of ALNS to the DPAS Problem

In what follows we describe an application of the ALNS methodology to the DPAS problem. We present a framework consisting of three destroy methods and two repair methods, and describe each of these in detail. Note that when we remove a patient from a solution, the patient is removed for all days it occupies the assigned room from the current day and into the future. Consequently, when repairing the solution, by inserting the patient again, the patient will be inserted into a room for the remaining part of his/her admission.

Given a solution to DPAS problem, the destroy methods focus on identifying a number of patients,  $N \in \mathbb{Z}$ , to be removed from the solution and subsequently reinserted. If  $N$  is small (say close to one or two), then the method will resemble a standard SA approach. If, on the other hand,  $N$  is large (close to  $|\mathcal{P}|$ ), then the resulting solution could become quite different, essentially running a reinsertion heuristic for all the patients. We choose the value of  $N$  to be within  $4 \leq N \leq \min\{100, \xi \cdot A\}$  using a uniform distribution, where  $\xi \in ]0; 1]$  and  $A$  is the number of active patients in the problem. The value 100 is chosen to limit the computational time used. The parameter  $\xi$  depends on the instance size.

The first destroy method we use is termed **Random Removal**. As the name suggests, this randomly selects  $N$  patients using a uniform distribution and removes them from the solution, the aim being to simply randomize the selection in order to explore other parts of the solution space.

The second removal heuristic aims at removing patients with the highest cost in the solution, and is termed **Worst Removal**. Such patients have obviously not received their most favorable room, and we are therefore interested in finding them another room. To ensure that we do not always select the worst positioned patients, we apply the following randomization which favors the selection of the worst cost patients. Let  $\bar{p}_1, \dots, \bar{p}_A$  be a list of patients sorted in decreasing order of patient cost, which corresponds to the cost of soft constraint violation directly related to the patient. That is, the penalty for the patient being placed in a given room, the penalty of violation of the gender policy, the penalty from the admission delay, and the penalty of not having a bed when there is an overcrowd risk. Suppose that  $r$  is uniformly distributed on the interval  $[0, 1[$ , and that  $\beta \geq 1$  is the randomization factor, we determine  $1 \leq n \leq A$  to be

$$n := \lfloor r^\beta A \rfloor + 1,$$

and select patient  $\bar{p}_n$  for removal. After removing patient  $\bar{p}_n$  we decrease the size of  $A$  by one and reset the indices of  $\bar{p}_1, \dots, \bar{p}_A$  accordingly. This is repeated until  $N$  patients are removed. Note that if  $\beta = 1$  then this selection corresponds to the random removal; however, when we increase  $\beta$ , the likelihood of selecting patients with lower indices increases. If, in the extreme case,  $\beta \rightarrow \infty$ , we would always select the first element of  $A$ .

As patients have different characteristics, removing two patients at random does not necessarily mean that their respective room assignments can be feasibly swapped. Patient differences can arise in admission timing or in treatment requirements. Consequently, the third removal heuristic – inspired by Shaw (1998) – is based on trying to remove *related* patients. This is termed **Related Removal**. A patient is randomly selected and the remaining  $N - 1$  patients are sorted in decreasing order of relatedness. We define relatedness as a weighted sum of four factors:

1.  $\mathbf{rel}_1(p_1, p_2) :=$  Number of days the stays of patients  $p_1$  and  $p_2$  overlap divided by the number

of days the patient with shorter stay has.

2.  $\text{rel}_2(p_1, p_2) :=$  The proportion of all rooms that are feasible for patients  $p_1$  and  $p_2$
3.  $\text{rel}_3(p_1, p_2) :=$  Equal to one if patients  $p_1$  and  $p_2$  have the same gender, zero otherwise
4.  $\text{rel}_4(p_1, p_2) :=$  A room is related cost-wise for the patients if the room has the exact same cost for these i.e.  $c_{p_1 r} = c_{p_2 r}$ . The relatedness based on a given room is then the cost of the lowest cost patient divided by the cost of the highest cost patient i.e.  $\min\{c_{p_1 r}, c_{p_2 r}\} / \max\{c_{p_1 r}, c_{p_2 r}\}$ , where we use the convention that the relatedness is 1 if  $\max\{c_{p_1 r}, c_{p_2 r}\} = 0$ . Thus the average room-based relatedness over feasible rooms is considered a relatedness measure.

$$\text{rel}_4(p_1, p_2) := \frac{1}{|\mathcal{R}_{p_1} \cap \mathcal{R}_{p_2}|} \sum_{r \in \mathcal{R}_{p_1} \cap \mathcal{R}_{p_2}} \frac{\min\{c_{p_1 r}, c_{p_2 r}\}}{\max\{c_{p_1 r}, c_{p_2 r}\}}$$

Each of these measures is normalized and has a value on the interval  $[0, 1]$ . Their relative importance does, however, differ. For example, two patients who are not admitted in overlapping periods are not related, nor are two patients who have no feasible rooms in common. Our final measure of relatedness is hence

$$\text{rel}(p_1, p_2) = \begin{cases} 0, & \text{if } \text{rel}_1(p_1, p_2) = 0 \text{ or } \text{rel}_2(p_1, p_2) = 0 \\ \sum_{i=0}^4 \varrho_i \text{rel}_i(p_1, p_2), & \text{otherwise} \end{cases}$$

where  $\varrho_i$  reflects the relative importance of relatedness measure  $i$ .

After the solution has been destroyed through the removal of  $N$  patients, the removed patients must be reinserted into the resulting partial solution. The first repair method is very simple and is termed **Greedy Insertion**. In this method patients are greedily inserted based on the cheapest insertion. The second, termed **Regret Insertion**, attempts to minimize the cost that will be incurred from not yet admitted patients when inserting a given patient. Recall that it is possible to postpone the admission of a not yet admitted patient  $p$  by up to  $d_p^{\max} - d_p^{\text{plan}}$  days.

Given a partial solution in which  $N$  of the  $A$  active patients are not yet inserted, for each of the removed patients a set of triples can be generated. Each triple is of the form  $(p, r, k)$  and states an insertion possibility for patient  $p$ . In particular, it states that patient  $p$  is assigned room  $r$  with  $k$  days of delay from the  $d_p^{\text{plan}}$ . Associated with each such triple is an assignment cost,  $c_{prk}$ , reflecting the increase in cost that would be occurred if the corresponding insertion were made. If the assignment indicated by a given triple is not feasible, then its cost is assumed to be infinitely large. The **Greedy Insertion** sorts all such triples in increasing order of cost and iteratively considers them. At any given insertion, the triple resulting in the smallest immediate increase in cost is considered. Due to the overstay risk and gender constraints, the cost of all triples needs to be updated after an insertion has been made. This process is repeated until either no more patients can be inserted or all patients have been feasibly inserted.

Greeditly inserting patients has the potential to spoil the possibility of placing other patients later in the insertion process. The **Regret Insertion** is an attempt to combat this. Intuitively we select the patient to insert in such a way which that it impacts the possibilities of remaining

unassigned patients the least. In other words, at any given step, the **Regret Insertion** focuses on the patient we would regret the most if not inserted immediately. We define  $\bar{c}_{pr} := \min_k \{c_{prk}\}$  to be the cheapest day to insert patient  $p$  in room  $r$ . If we then sort the  $\bar{c}_{pr}$  values (there is at most one for each room) in increasing order and let  $r(j)$  be the room at position  $j$  in the sorted list, we obtain the following general regret cost

$$v_p^k := \sum_{j=2}^k (\bar{c}_{pr(j)} - \bar{c}_{pr(1)}).$$

This is an indication of how much worse it will be to insert patient  $p$  in his/her  $k$ 'th best room compared to his/her best room. Selecting the patient  $p$  to insert with the largest value of  $v_p^k$  will try to minimize the regret by reducing the ill-effect caused by not being able to place patients in good rooms towards the end of the algorithm.

---

**Algorithm 1** ALNS with Simulated Annealing

---

```

Input:  $T_{start}, T_{end}, n, \xi, \alpha$ 
 $x_0 \leftarrow \text{runGreedyConstruction}()$ 
 $x \leftarrow x_0$ 
 $x_{best} \leftarrow x_0$ 
 $T \leftarrow T_{start}$ 
 $w \leftarrow \text{initializeMethodWeights}()$ 
while  $T > T_{end}$  do
  for  $i \leftarrow 1, n$  do
     $s \leftarrow \text{resetMethodScores}()$ 
    for  $j \leftarrow 1, 100$  do
       $N \leftarrow \text{random}([4, \min(100, \xi \cdot A)])$ 
       $\text{selectDestroyMethod}(w)$ 
       $\mathcal{P}_A \leftarrow \text{remove}(x, N)$ 
       $\text{selectRepairMethod}(w)$ 
       $x' \leftarrow \text{repair}(x, \mathcal{P}_A)$ 
       $\delta \leftarrow f(x') - f(x)$ 
       $\rho_{acc} \leftarrow \exp\left(\frac{-\delta}{T}\right)$ 
       $r \leftarrow \text{Random}([0, 1])$ 
      if  $\delta < 0$  or  $r < \rho_{acc}$  then
         $x \leftarrow x'$ 
        if  $f(x) < f(x_{best})$  then
           $x_{best} \leftarrow x'$ 
        end if
       $s \leftarrow \text{updateMethodScores}()$ 
    end if
  end for
   $w \leftarrow \text{updateMethodWeights}(s)$ 
end for
   $T \leftarrow \alpha \cdot T$ 
end while
return  $x_{best}$ 

```

---

Algorithm 1 provides an overview of the full SA based ALNS metaheuristic. For each temperature, we execute  $n$  trials. Each trial itself consists of 100 iterations. In other words, we allow exactly 100 solutions to be considered on any trial. Again the value of 100 has been chosen to

limit the run time of the algorithm. The value of  $n$  on the other hand varies with problem size; we set it to  $e \times |\mathcal{R}|$ , where  $e$  is a parameter we tune.

## 4 Computational Results

In this section we test and compare the performance of the proposed metaheuristic on the set of publicly available instances provided by Ceschia and Schaerf (2012)<sup>1</sup>. These artificially generated instances attempt to reflect reality, total some 450 in number, and are divided into nine families of 50 instances. Each family is characterized by a size (**small**, **medium**, or **large**) and a horizon length (**short**, **mid**, or **long**). The size of the instance stipulates the number of departments, the number of rooms, the number of features, and the number of specialisms, while the horizon length indicates the number of days in the planning horizon. In particular, the **short** instances consider 14 days, the **mid** instances look at 28 days, and the **long** instances include 56 days. The number of patients in an instance depends on the horizon length; it doubles if the horizon length doubles. An overview of the instance families is provided in Table 3. Comparing the performance of our algorithm on the same datasets as those used in Ceschia and Schaerf (2012) allows us to not only conclude something about the efficiency of the proposed algorithm, but also to benchmark the quality of the results obtained using the SA approach of Ceschia and Schaerf (2012).

The ALNS framework is programmed in C++, and all experiments have been performed on a dedicated Intel(R) Xeon(R) CPU X5550 @ 2.67GHz with 24 gigabytes of main memory running Ubuntu Linux version 14.04. The framework has several parameters. We omit an extensive discussion on the parameter tuning of the algorithm here, and simply report the best found values for each of the these. The chosen values are the following:  $T_{start} = 30$ ,  $T_{end} = 0.5$ ,  $e = 1.5$ , and  $\alpha = 0.5$ . Depending on the instance size, the parameter  $\xi$  is equal to 0.3 (**small**), 0.2 (**medium**), or 0.075 (**large**). The reaction factor  $\gamma$  is set to 0.1, and the weights on the relatedness measures are  $\varrho_1 = 4$ ,  $\varrho_2 = 3$ ,  $\varrho_3 = 2$ , and  $\varrho_4 = 1$ , respectively. We reward routines that produce solutions that are accepted with scores of  $\sigma_1 = 25$ ,  $\sigma_2 = 15$ , and  $\sigma_3 = 5$ , respectively. Finally, the randomness factor for worst removal is set to,  $\beta$ , is set to 0.2.

Family	Instances	Depts	Rooms	Features	Patients	Specialisms	Days
<b>small-short</b>	50	4	8	4	50	3	14
<b>small-mid</b>	50	4	8	4	100	3	28
<b>small-long</b>	50	4	8	4	200	3	28
<b>medium-short</b>	50	6	40	5	250	10	14
<b>medium-mid</b>	50	6	40	5	500	10	28
<b>medium-long</b>	50	6	40	5	1000	10	56
<b>large-short</b>	50	8	160	6	1000	15	14
<b>large-mid</b>	50	8	160	6	2000	15	28
<b>large-long</b>	50	8	160	6	4000	15	56

Table 3: Instance overview

The results for each of the instance families are summarized in Table 4. Each instance of each

<sup>1</sup>Available at <http://satt.diegm.uniud.it/projects/pasu/>



family has been run 10 times, and each row of the table provides descriptive statistics for the 50 instances which comprise the family. Comparing to the results in Ceschia and Schaerf (2012), for each family we report the largest improvement of any instance, the worst improvement of any instance, and the average improvement (all three given as a percentage). We also report the number of feasible solutions found by both approaches. For the ALNS approach we also report, in parentheses, the number of instances for which it produced a new best known solution. Finally, we also report the average time taken to solve each family size, as well as the average time taken to solve a day of the planning horizon (both in seconds). The latter is simply the average time divided by the number of days in the planning horizon.

Family	Improvement (%)			Solutions		Time (s)	Time (s)
	Best	Worst	Avg.	CS	ALNS	Avg.	Per Day
<b>small-short</b>	-16,54	5,11	-1,35	50	50 (27)	2,11	0,15
<b>small-mid</b>	-3,57	12,46	1,27	50	50 (13)	8,55	0,31
<b>small-long</b>	-4,81	7,02	1,74	49	50 ( 8)	39,52	0,71
<b>medium-short</b>	-6,99	3,37	-0,51	50	50 (25)	61,84	4,42
<b>medium-mid</b>	-5,37	2,72	-0,48	47	49 (31)	273,69	9,77
<b>medium-long</b>	-4,30	2,95	-1,25	49	50 (41)	1384,60	24,73
<b>large-short</b>	-7,56	-0,24	-3,46	48	48 (48)	2252,28	160,90
<b>large-mid</b>	-16,00	-0,55	-8,06	49	49 (49)	5346,35	190,94
<b>large-long</b>	-21,95	-7,16	-13,90	48	49 (49)	22 407,49	400,13

Table 4: Summary of results on benchmark instances

Three noticeable trends can be observed from the results in Table 4. The first is that the ALNS approach appears to perform better the larger the instance size gets. For all families larger than the **short** instances we report, on average better solutions. In some cases the improvement is dramatic. For example, on the **large-long** family we report solutions that are up to 22% better in some cases and on average 14% better. Furthermore, on the **large-long** instances we find new best known solutions for all instances. The second obvious trend is that the computation time drastically increases as the problem size grows. The solution times range from seconds on the **small-short** instances to hours on the **large-long** instances. A run time of 60 seconds per day of the planning horizon is granted by Ceschia and Schaerf (2012). No artificial time limit is imposed in our algorithm. The results for the **short** and **medium** families adhere to this; however, the same is not true for the **large** families. In the worst case, i.e. the **large-long**, we spend nearly seven minutes per day. It was a conscious decision not to enforce an artificial limit as we wanted to verify the quality of the solutions proposed in Ceschia and Schaerf (2012). It is not trivial how best to enforce the 60 second limit per day. Should all days have the same limit? Or should the days early in the time horizon be allocated more time? Naturally, the daily problems get easier to solve towards the end of the horizon as the number of registered patients decreases. Also, 60 seconds per day, irrespective of problem size, seems rather unrealistic. There are significantly more feasible solutions for the **large-long** instances than the **medium-long** instances. As such, a 60 second limit seems too general. Regardless of time limits, we have shown that solutions obtained on the **large-long** instances by Ceschia and Schaerf (2012) are, in some cases, of poor

quality. Finally, the third observation is that the ALNS approach is always at least as good at finding feasible solutions as the approach by Ceschia and Schaerf (2012). For example, on the **medium-mid** instances the ALNS approach finds a feasible solution on 49 of the 50 instances, while the approach by Ceschia and Schaerf (2012) finds 47.

	ALNS*	CS*	ALNS (Avg.)	CS (Avg.)	$\Delta$ Best	$\Delta$ Avg.
<b>small-short</b>	2694,26	2723,02	2763,82	2761,07	-1,06	0,10
<b>small-mid</b>	6119,52	6058,18	6269,55	6165,21	1,01	1,69
<b>small-long</b>	12 394,14	12 179,73	12 627,70	12 383,17	1,76	1,97
<b>medium-short</b>	12 350,88	12 398,46	12 658,08	12 662,78	-0,38	-0,04
<b>medium-mid</b>	26 617,96	27 670,66	27 107,11	27 203,34	-0,53	-0,35
<b>medium-long</b>	61 816,90	62 616,82	62 530,64	63 483,94	-1,29	-1,43
<b>large-short</b>	38 931,65	40 302,35	39 589,10	41 099,36	-3,40	-3,67
<b>large-mid</b>	100 436,55	108 864,84	101 773,81	111 176,49	-6,51	-8,46
<b>large-long</b>	192 537,25	223 225,35	195 431,73	227 118,07	-13,75	-13,95

Table 5: Summary of best and average solutions on benchmark instances

Similar observations can be seen in Table 5. As a means to directly compare our metaheuristic (ALNS in the table) to that of Ceschia and Schaerf (2012) (CS), we compute two different average objective values for both approaches across the 50 instances of each family. The first gives the average best objective value (\*), while the second refers to the average of the mean objective values (Avg.). Recall that each instance is solved 10 times. The table also reports the percentage change between the best solutions and the average solutions for the two different methods ( $\Delta$  Best and  $\Delta$  Avg.). Table 5 shows that the best objective value is around 0.5% to 14% better, while the same is true for the mean objective values. For the **large** families the fact we are better on average, compared to the best solution, indicates that ALNS approach constantly finds good solutions. For the **small** instances the solutions provided by the two approaches are comparable; however, the ALNS provides them in a fraction of the time. The apparent improvement on the **small-short** is somewhat is, however, somewhat misleading. This is due to the ALNS approach producing a significantly better solution on one of the instances (approximately 16.5% better).

Finally, all solutions have been validated using the official solution validator (can be downloaded from <http://satt.diegm.uniud.it/projects/pasu/>) and are available online at <http://www.ms.man.dtu.dk/Research/Instances>.

## 5 Conclusion

In this paper, we have devised a metaheuristic for the DPAS problem. The method is uses a SA framework and the ALNS procedure. This has resulted in an efficient and flexible method that can solve problems of different sizes and with different time horizons. For small problems solutions a found within 40 seconds, while for medium sized instances solutions are obtained within 25 minutes. Larger instances take somewhat longer, on average a few hours. However, it should be stressed that in reality the time taken to solve the problem on a given day of the planning horizon is more relevant. Since in a dynamic setting the problem only needs to be solved on a

daily basis. For this, the times are much shorter, fractions of a second for the small instances and up to around seven minutes for the larger instances. In most cases our method is superior to the methods suggested by Ceschia and Schaerf (2012). For the large instances we report improvements in the range of 3-14% on average, but as large as 22% for individual instances. Furthermore, we find solutions to more of the instances than Ceschia and Schaerf (2012). Out of the 450 instances used in the experiment we find a feasible solution to 445 of them.

If this work should be brought closer to a real-life implementation it is necessary to reduce running times for the large instances in order to get a tool that could work within the planning cycles seen at hospitals. Furthermore, ideas from the mathematical programming could be integrated into the heuristic framework in order to derive a matheuristic method.

## References

- Aarts, E., Korst, J., Michiels, W., 2005. Simulated annealing. *Search Methodologies: Introductory Tutorials in Optimization and Decision Support Techniques*, Search Methodologies: Introductory Tutorials in Optimization and Decis. Support Techniques, 187–210.
- Bilgin, B., Demeester, P., Misir, M., Vancroonenburg, W., Vanden Berghe, G., 2012. One hyperheuristic approach to two timetabling problems in health care. *Journal of Heuristics* 18 (3), 401–434.
- Ceschia, S., Schaerf, A., 2011. Local search and lower bounds for the patient admission scheduling problem. *Computers and Operations Research* 38, 1452 – 1463.
- Ceschia, S., Schaerf, A., 2012. Modeling and solving the dynamic patient admission scheduling problem under uncertainty. *Artificial Intelligence in Medicine* 56, 199–205.
- Ceschia, S., Schaerf, A., 2014. Dynamic patient admission scheduling problem with operating room constraints, flexible horizons and patient delays. *Journal of Scheduling in press*.
- Chen, N., Zhan, Z., Zhang, J., Zhang, J., Liu, O., Liu, H., 2010. A genetic algorithm for the optimization of admission scheduling strategy in hospitals. In: *Proceedings of the IEEE Congress on Evolutionary Computation (CEC) 2010*. IEEE, pp. 1 – 5.
- Demeester, P., Souffriau, W., Causmaecker, P. D., Vanden Berghe, G., 2010. A hybrid tabu search algorithm for automatically assigning patients to beds. *Artificial Intelligence in Medicine* 48, 61–70.
- Harper, P. R., 2002. A framework for operational modelling of hospital resources. *Health Care Management Science* 5, 165 – 173.
- Hutzschenreuter, A. K., Bosman, P. A. N., Blonk-Altena, I., Aarle, J., Poutré, H. L., May 2008. Agent-based patient admission scheduling in hospitals. In: *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008) – Industry and Applications Track*. pp. 42 – 52.

- Jittamai, P., Kangwansura, T., 2011. A hospital admission planning model for emergency and elective patients under stochastic resource requirements and no-shows. 2011 IEEE International Conference on Industrial Engineering and Engineering Management (IEEM), 166–170.
- Kirkpatrick, S., Gelatt, D., Vecchi, M., 1983. Optimization by simulated annealing. *Science* 220 (4598), 671 – 680.
- Kusters, R. J., Groot, P. M. A., 1996. Modelling resource availability in general hospitals design and implementation of a decision support model. *European Journal of Operational Research* 88, 428 – 445.
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., Teller, E., 1953. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics* 21 (6), 1087–1092.
- Pisinger, D., Ropke, S., 2006. An adaptive large neighbourhood search heuristic for the pickup and delivery problem with time windows. *Transportation Science* 40, 455–472.
- Range, T., Lusby, R., Larsen, J., 2014. A column generation approach for solving the patient admission scheduling problem. *European Journal of Operational Research* 235, 252–264.
- Shaw, P., 1998. Using constraint programming and local search methods to solve vehicle routing problems. *Lecture Notes in Computer Science* 1520, 417 –431.
- Vancroonenburg, W., Causmaecker, P. D., Berghe, G. V., 2012. Patient-to-room assignment planning in a dynamic context. In: *Practice and Theory of Automated Timetabling (PATAT 2012)*.