European Journal of Operational Research 235 (2014) 252–264

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Innovative Applications of O.R.

# A column generation approach for solving the patient admission scheduling problem

Troels Martin Range [a],*, Richard Martin Lusby [b], Jesper Larsen [b]

[a] Department of Business & Economics, COHERE, University of Southern Denmark, Campusvej 55, 5230 Odense M, Denmark
[b] Department of Management Engineering, Technical University of Denmark, Produktionstorvet, Building 426, 2800 Kgs. Lyngby, Denmark

ARTICLE INFO

*Article history:*
Received 10 January 2013
Accepted 19 October 2013
Available online 30 October 2013

*Keywords:*
OR in health services
Scheduling
Column generation
Dynamic constraint aggregation
Dual disaggregation
Branch and bound

ABSTRACT

This paper addresses the Patient Admission Scheduling (PAS) problem. The PAS problem entails assigning elective patients to beds, while satisfying a number of hard constraints and as many soft constraints as is possible, and arises at all planning levels for hospital management. There exist a few, different variants of this problem. In this paper we consider one such variant and propose an optimization-based heuristic building on branch-and-bound, column generation, and dynamic constraint aggregation to solve it. We achieve tighter lower bounds than previously reported in the literature and, in addition, we are able to produce new best known solutions for five out of twelve instances from a publicly available repository.

## 1. Introduction

In this paper we consider a new approach for solving the Patient Admission Scheduling (PAS) problem. In many countries, the allocation of patients to wards (more specifically, to beds) is carried out by a central planning unit in the hospital management. The algorithm presented in this paper can assist hospital management at the operational level in assigning elective patients to beds, while attempting to satisfy as many individual patient preferences as is possible and making sure critical medical equipment is available. Hospitals usually have slack in order to deal with emergency patients, which are, naturally, difficult to plan for. Emergency patients are not considered in this study.

The area of patient admission scheduling in hospitals is important for a high-quality health care system. To provide health care services is perhaps one of the most complex industries worldwide. Planning and managing the operations of a hospital in an efficient manner requires a sound knowledge of the hospital system, and understanding the patient flow is critical for this. Most literature within patient admission scheduling focuses on the strategic (and tactical) level and describes tools and cases for assisting hospital management with tactical and strategic planning problems. Some examples are: Jittamai and Kangwansura (2011), Harper (2002) and Kusters and Groot (1996). The aim here is to have the right number of beds available in order to increase the efficiency of the hospital.

The problem that we investigate in this paper is the operational version of the PAS problem. This problem has only recently started to receive attention. It is indirectly part of scheduling operating rooms, which has gained some attention. The survey in Guerriero and Guido (2011) lists some 129 publications that focus on the optimization of hospital operating theaters.

In the few, existing articles available on the operational level PAS problem slightly different versions of the problem are considered. In Hutzschenreuter, Bosman, Blonk-Altena, Aarle, and Poutré (2008) an agent-based model for the problem is presented. A model of the patient flow is generated in order to admit an optimal mix of patients from different departments with a special focus on the common usage of (central) hospital resources. The work is based on a detailed description of the patient flow and its probabilities. This system can be used to develop policies for achieving a good patient mix and to facilitate efficient hospital operations.

In contrast to this, Chen et al. (2010) develop a genetic algorithm for an admission scheduling problem with only one department. The algorithm uses historical data and optimizes a long-term admission strategy instead of suggesting a specific schedule.

In the remaining work in this area, Demeester, Souffriau, Causmaecker, and Vanden Berghe (2010), Bilgin, Demeester, Misir, Vancroonenburg, and Vanden Berghe (2012) and Ceschia and Schaerf (2011) all consider the same version of the PAS problem. This definition of the problem was originally presented in Demeester et al. (2010). In addition, the authors have published

---

* Corresponding author. Tel.: +45 6550 3685; fax: +45 6550 3237.
*E-mail addresses:* tra@sam.sdu.dk (T.M. Range), rmlu@dtu.dk (R.M. Lusby), jesla@dtu.dk (J. Larsen).

0377-2217/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ejor.2013.10.050

a website containing instances,[1] current best known solutions, and a solution validator.

The problem has a fixed number of elective patients. In practice, elective patients often wait for an admission date; however, in this particular context, each patient has already been assigned both an admission date and a discharge date, and these are known with certainty. In addition, each patient states their medical requirements. The aim is then to determine which beds to assign to which patients in order to maximize the patient comfort and efficiency of the medical operations. This planning has to be done in compliance with a few hard constraints; for example, admission and discharge dates must be respected, mandatory medical equipment needs to be in the room assigned, and male and female patients cannot share the same room simultaneously. In this paper we present a new mathematical formulation of the problem and devise an efficient column generation-based heuristic for solving it. Our methodology utilizes a dynamic constraint aggregation procedure to overcome some of the problems associated with solving large integer programming formulations. The methodology is tested on the benchmark instances provided by Demeester et al. (2010), where we report tighter linear programming (LP) bounds, significantly faster running times than previous mathematical programming approaches in achieving these, and five new best known integer solutions.

It should finally be noted that in an effort to get closer to the real-life problem Ceschia and Schaerf (2012) introduce the *dynamic PAS* problem or *PAS under uncertainty (PASU)*. This problem is based on the PAS problem as defined in Demeester et al. (2010) with the addition that patients do not necessarily get admitted to the hospital on their planned admission date, and may also stay an extra day longer.

In this paper, our aim is to focus on the PAS problem, introduce a new method for solving the problem, and benchmark its performance against existing approaches available. The structure of this paper is, therefore, as follows. In Section 2, we introduce the PAS problem in more detail and present a binary integer programming formulation to model it. Section 3 describes the proposed solution methodology as well as a detailed discussion on the dynamic constraint aggregation procedure. Here we also introduce a new dual disaggregation strategy. Computational results are described in Section 4, and conclusions from this study are drawn in Section 5.

## 2. Problem definition

In this section we consider the PAS problem in more detail. In particular, we provide a detailed overview of the constraints of the problem and introduce the required terminology and notation that is used throughout the paper. The description of the problem is consistent with that of Demeester et al. (2010).

The PAS problem requires that a set of patients requiring medical attention, $\mathcal{P}$, is assigned to a set of hospital beds over a prespecified daily time horizon. We denote the set of consecutive days that comprise this planning horizon as $\mathcal{T}$. Each patient, $p \in \mathcal{P}$, is assumed to have a known admission date, $a_p \in \mathcal{T}$, and a known discharge date, $d_p \in \mathcal{T}$, which together define the duration of the patient's stay in the hospital.[2] We let $\mathcal{W} = \{(p,t) \in \mathcal{P} \times \mathcal{T} | a_p \leqslant t < d_p\}$ be the set of all *patient-time combinations*; in other words, the set of all patient days which must be assigned at the hospital.

For treatment each patient requires one or more so-called *specialisms* (e.g. cardiology, oncology, gerontology, etc.). Typically, patients need just one specialism. The case where patients need multiple specialisms is excluded from our study. Each bed is located in a particular *room*, and each room belongs to a particular *department* at the hospital. Every department has the ability to cater for the treatment of a variety of specialisms, but with varying degrees of expertise. Some departments (e.g. pediatrics, gerontology) also have an *age policy*, which means that the department can only admit patients of a certain age. As is the case for departments, rooms also have a ranked list of specialisms for which they are suitable. Furthermore, every room has a certain number of identical beds, termed the *room capacity* (typically this is one, two, or four beds) and is equipped with a set of *properties* (e.g. oxygen and telemetry) that can be used for treatment. The presence of certain room properties for a patient may either be *required* (mandatory) or simply *preferred*. Each room also has a *gender policy*. This stipulates that all patients in the room must be male, female, segregated (the same gender on any given day), or unrestricted. The rooms at the hospital can be classified as a particular *room type*, where all rooms of a particular type are identical. We denote the set of all room types at the hospital as $\mathcal{R}$, and for each $r \in \mathcal{R}$ we let $N_r$ be the number of available rooms of this type. Since all beds in a given room are identical and all the rooms of a particular type are also identical, one can equivalently view the PAS problem as finding an assignment of patients to room types over the planning horizon.

In the process of assigning patients to room types, several constraints are enforced. These are classified as either *hard* or *soft* constraints. The former must be respected, while the latter can be violated; however, a penalty is incurred for each violation. In the version of the PAS problem that we consider there are five different types of hard constraints. Firstly, there is the obvious requirement that no two patients can simultaneously occupy the same bed. Viewed from a room type assignment perspective, this is equivalent to stating that the capacity of the room type must not be violated on any day. Secondly, on each day of the planning horizon each room type's gender policy must be respected. Thirdly, for each patient's admission the period of stay is contiguous. Finally, each patient should be assigned a room type consistent with their age and the required mandatory equipment.

All other constraints are considered soft and will be described in turn. In this way we will remain consistent with the overall approach to the problem posed by Demeester et al. (2010). We begin by discussing all penalties that are incurred when assigning a given patient, $p \in \mathcal{P}$, to a particular room type, $r \in \mathcal{R}$. All such individual penalties can be accumulated to obtain one patient room type assignment penalty, $d_{pr}$. Firstly, each patient has a preference for the capacity of the room in which they will stay. A patient who is assigned a room type with a capacity larger than their preference incurs a penalty. Secondly, a patient should be placed in a department whose most competent specialism matches the patient's required specialism. Lower levels of expertise are penalized. Similarly, a patient should receive a room type that best matches the patient's specialism. Imperfect matches are again penalized. Finally, in terms of room properties, patients should receive a room type that is equipped with their preferred properties. Missing preferred properties are penalized. The only penalty not related to room type assignment concerns the transfer of patients between two different rooms on consecutive days. Since it is undesirable to transfer patients during their stay, every time a patient is transferred, a transfer-out penalty, $t^o$, as well as a transfer-in penalty, $t^i$, are incurred. Table 1 gives the weights associated with violating each constraint. These are identical to Demeester et al. (2010), again for consistency.

Given the problem description outlined above, one can formally state the PAS problem as follows: find the minimum penalty allocation of patients to room types so that each patient is allocated to

---

[1] http://allserv.kahosl.be/~peter/pas.

[2] In the benchmark instances proposed by Demeester et al. (2010) some patients have a discharge dates beyond the planning horizon. For these patients we provide a virtual discharge date corresponding to the end of the planning horizon.