Approximate dynamic programming for patient admission scheduling

Thiago Augusto de Oliveira Silva ·
Wim Vancroonenburg ·
Greet Vanden Berghe

Keywords Patient admission scheduling · Meta-heuristic · Approximate Dynamic Programming

Planning and scheduling problems that result from operational level decision-making often involve repeatedly making decisions over time. From an operations research perspective, such problems are frequently addressed by formulating a deterministic optimization model that represents the decision-making's long-term goal. This model is then employed to solve a short-term view on the problem which includes the data available at that moment in time, which is subsequently re-optimized by updating and solving the model again when new data becomes available. One drawback of this approach is its *reactive* nature, which essentially means that it deals with new information only as it becomes available. Its decisions are therefore always trailing and mitigating the current situation, which may be a source of inefficiency. By contrast, an *anticipative* approach may be more appropriate in such a setting to better handle the problem's dynamic nature, by making decisions that are more *robust* with regard to information yet to arrive.

This study focuses on the dynamic patient admission scheduling problem (DPASP), as defined by Ceschia and Schaerf [1]. Patient admission scheduling is an important process in the operation of every hospital, enabling the hospital to manage the inflow of elective *inpatients* (patients requiring at least one overnight stay) and to match demand for hospital services (surgery, ima-

T. A. O. Silva
Departamento de Engenharia de Produção
Instituto de Ciências Exatas e Aplicadas
Universidade Federal de Ouro Preto

W. Vancroonenburg, G. Vanden Berghe
Department of Computer Science, CODeS & imec
KU Leuven
Ghent Technology Campus, Gebroeders De Smetstraat 1, 9000 Gent, Belgium

Proceedings of the 12th International Conference on the Practice and Theory of Automated Timetabling (PATAT-2018), Vienna, Austria, August 28–31, 2018

ging, treatments) and resources (beds, personnel) with their available capacity. The DPASP was defined to develop solution approaches that support this decision-making problem. The problem involves the assignment of admission dates to elective inpatients, and the assignment of rooms to both elective and emergency inpatients for the duration of their stay. Capacity for performing (surgical) treatments, available rooms and equipment for patient stays are however limited. Operational policies regarding room assignments (gender separation, age segregation) further complicate assignment and scheduling decision-making. More importantly, the DPASP correctly addresses the admission scheduling process in a *dynamic* context, where decisions are made daily when new information concerning admission requests becomes available. In the DPASP, the set of inpatients is gradually revealed over a pre-defined planning horizon, requiring scheduling and assignment decisions to be repeatedly made as new inpatient admissions are registered and inpatients arrive at the hospital. The overall quality of a solution for an instance of the DPASP is measured by a (long-term) objective function that calculates the *cost* (a weighted summation of penalties, resulting from not matching operational policies, delaying admissions, unused capacity) at the end of the planning horizon. Ceschia and Schaerf [1] developed a Simulated Annealing (SA) based meta-heuristic that solves the daily decision-making problem, given the current state of the hospital: the currently admitted inpatients, the assigned rooms for their stay and the queue of forthcoming inpatients. This solution method is subsequently applied in a re-optimization fashion as time passes with regard to the planning horizon.

The present contribution proposes the application of a Simulation Based Approximate Dynamic Programming (SB-ADP) approach [2, 3] to deal with the dynamic and uncertain nature of this problem. Using this methodology, the DPASP is viewed as a control problem where the system, during the planning horizon (D), is represented by the daily state (x_d, $d \in D$) of the hospital (admitted patients, room occupation, inpatient queue), and evolves from x_d to the next state x_{d+1} in accordance with actions u_d (new admissions, assigning rooms) derived from some *policy* $\pi(x_d)$ and by next day uncertainties (ω_{d+1}) realization (new arrivals, delays). The aim of this ADP-modeling is to find a policy (π^*) that minimizes the so-called Value Function over the horizon D:

$$J_d^\pi(x_d) = C_d^\pi(x_d, u_d^\pi(x_d)) + \mathbb{E}_{w_{d+1}}\left\{J_{d+1}^\pi(t(x_d, u_d^\pi(x_d), w_{d+1}))|x_d, u_d^\pi(x_d)\right\} \quad (1)$$

In the context of the DPASP, the Value Function corresponds to the remaining cost (*Cost-To-Go*) of the long-term objective function ($C_d^\pi(\cdot)$) obtained at the end of the planning horizon, given the current state of the system at a point in time before the end of D. Due to the inability to enumerate the entire state, action and outcome space, the SB-ADP approach employs a forward-stepping approximation method for approximating the optimal policy by using estimations of Value Functions. Before the optimization phase, the approach simulates the process to collect data for fitting linear regression models on state *basis* functions in order to estimate the value associated with the next day state. Internally, the SA meta-heuristic developed by Ceschia and

Schaerf [1] is employed to evaluate the Cost-To-Go (by solving the remaining days of the planning horizon using the SA method).

The present study discusses the development of the basis functions and evaluates the SB-ADP approach in a computational experiment. This experiment employs the instance generator for the DPASP developed by Ceschia and Schaerf [1], to generate off-line training data for th SB-ADP approach, as well as to sample random paths used by the SB-ADP method. The study and its experiments are currently ongoing. The first results will be presented at the conference.

Acknowledgements Wim Vancroonenburg is a post-doctoral researcher funded by Research Foundation Flanders - FWO Vlaanderen. Thiago Silva was funded by Coimbra Group during the development of this study. We kindly thank Sara Ceschia and Andrea Schaerf for providing the source code of their method. Editorial consultation provided by Luke Connolly (KU Leuven).

References

1. Ceschia, S., Schaerf, A.: Dynamic patient admission scheduling with operating room constraints, flexible horizons, and patient delays. Journal of Scheduling **19**(4), 377–389 (2016). DOI 10.1007/s10951-014-0407-8. URL `https://doi.org/10.1007/s10951-014-0407-8`
2. Powell, W.B.: Approximate Dynamic Programming: Solving the curses of dimensionality, vol. 703. John Wiley & Sons (2007)
3. Powell, W.B.: What you should know about approximate dynamic programming. Naval Research Logistics (NRL) **56**(3), 239–249 (2009)