Particle Therapy Patient Scheduling: Time Estimation to Schedule Sets of Treatments*

Johannes Maschler, Martin Riedler, and Günther R. Raidl

Institute of Computer Graphics and Algorithms, TU Wien, Vienna, Austria
{maschler|riedler|raidl}@ac.tuwien.ac.at

1 Introduction

In classical radiotherapy cancer treatments are provided by linear accelerators that serve a dedicated treatment room exclusively. In contrast, particle therapy uses beams that are produced by either cyclotrons or synchrotrons that can serve up to five treatment rooms in an interleaved way. Several sequential activities like stabilization not requiring the beam have to be performed in the treatment room before and after each actual irradiation. Using several rooms and switching the beam between the rooms thus allows an effective utilization of the expensive particle accelerator and increased throughput of the facility.

In a typical midterm planning scenario a schedule for performing the therapies over the next few months has to be determined. Midterm planning for classical radiotherapy has already attracted some research starting with the works from Kapamara et al. [1] and Petrovic et al. [3]. Due to the one-to-one correspondence of treatment rooms and accelerators it suffices to consider a coarser scheduling scenario in which treatments have to be assigned only to days but do not have to be sequenced within the day. In a recent work [2] we studied a simplified problem formulation addressing the midterm planning of the particle therapy treatment center MedAustron in Wiener Neustadt, Austria, which offers three treatment rooms. Our approach consisted in decomposing the problem into a day assignment and a sequencing part, and we provided a construction heuristic, a GRASP, and an Iterated Greedy (IG) metaheuristic. The aim of the current work is to extend the proposed model and to provide and utilize a mechanism that quickly predicts the behavior of the sequencing part with reasonable precision, allowing in particular an improved day assignment.

2 Particle Therapy Patient Scheduling Problem

In the Particle Therapy Patient Scheduling Problem (PTPSP) therapies consisting of daily treatments (DTs) on 8 to 35 subsequent days have to be planned. Each therapy has a window of days at which it is allowed to start. There is a minimal and maximal number of DTs that have to be provided each week, a lower and upper bound of days that are allowed to pass between two subsequent

* We thank EBG MedAustron GmbH, Wiener Neustadt, Austria, for the collaboration on particle therapy patient scheduling and partially funding this work.

DTs, and a required break of at least two consecutive days each week. Moreover, DTs should be provided roughly at the same time within each week. Each DT consists of either five or seven consecutive activities having associated processing times and requiring individual resources such as the particle beam, room, and anesthetist. Resources have each day a regular and an extended availability period at which they can be used, where the use of the latter one results in additional costs. The availability periods of resources can be further interrupted by so-called unavailability periods.

A schedule assigns all DTs of a given set of therapies to days and determines starting times for the associated activities considering all operational constraints. The considered objective minimizes the use of extended availability periods, the finishing day of the therapies, and the variation of the starting times of the DTs.

3 Solution Approach and Time Estimation

Our solution approach consists of decomposing PTPSP into the Day Assignment (DA) in which DTs are assigned to days and the Time Assignment (TA) in which for each day starting times for the respective DTs activities are determined. Clearly, those two levels are dependent on a large degree. Determining the usage of the resources' availability periods for a given candidate set of DTs at a specific day is of crucial interest in any method determining an optimized DA as well as any constructive heuristic for the TA. Optimally solving the associated subproblem, however, would in general require to completely solve the underlying scheduling problem, which is very time-expensive if practically possible at all. Therefore, we investigate different efficient ways to estimate the use of the resources' availability periods with reasonable accuracy and study the impact on the greedy heuristic, GRASP, and IG.

Furthermore, in [2] we did not yet regard the requirement that each therapy's DTs should be provided at roughly the same time. This allowed a more independent calculation of the TA for each day. We now extend our methods to also address this soft-constraint. The main idea is that in the underlying construction heuristic the days are now considered sequentially and DTs are assigned preferably to starting times that are close to starting times of previous DTs.

References

1. Kapamara, T., Sheibani, K., Haas, O., Petrovic, D., Reeves, C.: A review of scheduling problems in radiotherapy. In: Proceedings of the International Control Systems Engineering Conference (ICSE). pp. 207–211. Coventry University Publishing, Coventry, UK (2006)
2. Maschler, J., Riedler, M., Stock, M., Raidl, G.R.: Particle therapy patient scheduling: First heuristic approaches. In: Proceedings of the 11th Int. Conference on the Practice and Theory of Automated Timetabling. Udine, Italy (2016)
3. Petrovic, S., Leung, W., Song, X., Sundar, S.: Algorithms for radiotherapy treatment booking. In: 25th Workshop of the UK Planning and Scheduling Special Interest Group. pp. 105–112. Nottingham, UK (2006)