

Particle Therapy Patient Scheduling: Time Estimation for Scheduling 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

Abstract. In the particle therapy patient scheduling problem (PTPSP) cancer therapies consisting of sequences of treatments have to be planned within a planning horizon of several months. In our previous works we approached PTPSP by decomposing it into a day assignment part and a sequencing part. The decomposition makes the problem more manageable, however, both levels are dependent on a large degree. The aim of this work is to provide and a surrogate objective function that quickly predicts the behavior of the sequencing part with reasonable precision, allowing an improved day assignment w.r.t. the original problem.

Keywords: particle therapy patient scheduling, time estimation, bilevel optimization, surrogate objective function, iterated greedy metaheuristic

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, produced by either cyclotrons or synchrotrons, that can serve up to five treatment rooms in an interleaved way. Several sequential tasks that do not require the beam, like the positioning of the patients, have to be performed in the treatment room before and after each irradiation. Switching the beam between treatment rooms allows an effective utilization of the particle accelerator and increases the throughput of the facility.

In a typical midterm planning scenario a schedule over the next few months for performing therapies, consisting of a sequence of treatments, has to be determined. Midterm planning for classical radiotherapy has already attracted some research starting with the works from Kapamara et al. [4] and Petrovic et al. [7]. In the following a variety of methods has been applied ranging from GRASP [8] and steepest hill climbing methods [3] to MILP approaches [2, 1]. Due to the one-to-one correspondence of treatment rooms and accelerators it is sufficient 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 days.

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

In a recent work [6] we studied the midterm planning for 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 for each day, which are, however, not independent. We provided a construction heuristic, a GRASP, and an iterated greedy (IG) metaheuristic. Our computational experiments showed that the IG performed the best. In the subsequent work [5] we improved our IG by exchanging the construction operator with one that is able to preserve more sequencing related information from the incumbent solution. This and the application of a new local search method allowed to outperform the previous approaches.

The current work focuses on further improving the day assignment level by using a fast to compute, more accurate surrogate model for estimating the optimal objective values for the sequencing subproblems. In particular, during the day assignment phase the total duration of the schedule produced by the sequencing part has to be estimated for each day. We previously applied a lower bound that, however, yielded avoidable overfull days. Hence, we study here surrogate functions that consider more aspects of the problem at hand. The main goal is to predict the overall objective function contribution of the sequencing part with reasonable precision while being computationally fast enough to be used in our existing approaches.

2 Particle Therapy Patient Scheduling Problem

In the particle therapy patient scheduling problem (PTPSP) a set of therapies T has to be scheduled on consecutive days D considering a set of resources R .

Each therapy $t \in T$ consists of a set of daily treatments (DTs) $U_t = \{1, \dots, \tau_t\}$. A DT comprises all tasks required to provide an irradiation and has a window of days at which it can be applied. 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 DTs, and a required break of at least two consecutive days each week. Each DT $u \in U_t$ has a processing time $p_{t,u} \geq 0$ and a set of required resources $Q_{t,u} \subseteq R$. In the execution of a DT each resource $r \in Q_{t,u}$ is required during part of the whole processing time specified by the time interval $[P_{t,u,r}^{\text{start}}, P_{t,u,r}^{\text{end}}) \subseteq [0, p_{t,u})$.

The days in the planning horizon D are partitioned into weeks. A subset $D' \subseteq D$ are working days (usually Mondays till Fridays) on which DTs can be performed. For each day $d \in D'$ we have a fundamental opening time $[\widetilde{W}_d^{\text{start}}, \widetilde{W}_d^{\text{end}})$ that limits the availability of all resources on the considered day.

Resources $r \in R$ have on each working day a regular availability period $[W_{r,d}^{\text{start}}, W_{r,d}^{\text{end}})$ and an extended availability period $[W_{r,d}^{\text{end}}, \widetilde{W}_d^{\text{end}})$ at which they can be used, where the usage of the latter one results in additional costs. Moreover, the availability of resource r may be interrupted on day d by a set of unavailability intervals $\bigcup_{w=1, \dots, \omega_{r,d}} [\overline{W}_{r,d,w}^{\text{start}}, \overline{W}_{r,d,w}^{\text{end}})$.

A solution for the PTPSP is a tuple (Z, S) , where $Z = \{Z_{t,u} \in D' : t \in T, u \in U_t\}$ are the days and $S = \{S_{t,u} \geq 0 : t \in T, u \in U_t\}$ are the times at

which the DTs are planned. A solution is feasible if all resource availabilities, and all operational constraints are respected. We aim at minimizing the use of extended availability periods while finishing each therapy as early as possible. More formally, the objective is to minimize

$$\gamma^{\text{ext}} \sum_{r \in R} \sum_{d \in D'} \eta_{r,d} + \gamma^{\text{finish}} \sum_{t \in T} (Z_{t,\tau_t} - Z_{t,\tau_t}^{\text{earliest}}), \quad (1)$$

where γ^{ext} and γ^{finish} are scalar weights, $\eta_{r,d} = \max(\{S_{t,u} + P_{t,u,r}^{\text{end}} - W_{r,d}^{\text{end}} \mid t \in T, u \in U_t, r \in Q_{t,u}, Z_{t,u} = d\} \cup \{0\})$ is the used time of the extended availability period of resource r on day d , and $Z_{t,\tau_t}^{\text{earliest}}$ is a lower bound on the earliest possible finishing day for the last DT of therapy t (see [6]).

3 Solution Approach and Time Estimation

The PTPSP naturally decomposes into the day assignment (DA) level in which DTs are assigned to days and the time assignment (TA) level that consists of finding starting times for the DTs. In other words, Z are the first level and S are the second level decision variables. Clearly, those two levels are dependent on a large degree. Nevertheless, this problem decomposition is beneficial because we can separate the detailed resource model from the remaining operational constraints. Thus, in the TA level each day becomes independent and can be solved separately.

A central aspect of the DA is to find a well-paired allocation of DTs to days that causes as little use of extended availability periods as possible. Since determining $\eta_{r,d}$ requires the exact starting times, the usage of the resources' availability periods for a given candidate set of DTs has to be estimated. Thus, the DA uses a modified version of (1) that replaces $\eta_{r,d}$ with the surrogate $\hat{\eta}_{r,d} = \max(0, \hat{\lambda}_{r,d} - h_{r,d})$, where $\hat{\lambda}_{r,d}$ estimates the required time and $h_{r,d}$ denotes the aggregated regular availability of resource r on day d . In our previous works [6, 5] we used for $\hat{\lambda}_{r,d}$ the trivial lower bound given by aggregated resource demands $\sum_{(t,u): t \in T, u \in U_t, r \in Q_{t,u}, Z_{t,u}=d} (P_{t,u,r}^{\text{end}} - P_{t,u,r}^{\text{start}})$. Consequently, the DA frequently underestimated the resource consumption, which resulted in avoidable use of extended availability periods in the TA. In this work, we aim at more accurately estimating $\hat{\lambda}_{r,d}$ for the main bottleneck resources, the beam and the rooms.

4 Estimating the Makespan under Complete Resource Availability

In the following we first concentrate on estimating the makespan required for a given non-empty set G of DTs under the assumption that all required resources are available without any further restrictions. We start by determining estimations of the makespan for three special cases. Afterwards, an estimation for the general case is derived that is based on the estimation for these special cases. Let $n_r = |\{(t,u) \in G \mid r \in Q_{t,u}\}|$ be the number of DTs requiring resource

$r \in \{1, 2, 3, B\}$, where 1, 2, 3 represent the rooms and B the beam, respectively. Furthermore, let

$$\bar{P}_r = \begin{cases} \frac{\sum_{(t,u) \in G} (P_{t,u,r}^{\text{end}} - P_{t,u,r}^{\text{start}})}{n_r} & \text{if } n_r > 0 \\ 0 & \text{else} \end{cases} \quad (2)$$

be the average time resource $r \in \{1, 2, 3, B\}$ is required by DTs in G . Moreover, let P^{irb} and P^{ira} be the minimum durations a room is required before and after the beam resource, respectively, and let P^{orb} and P^{ora} be the minimum times required by any DT before and after the usage of the room resource, respectively.

In the first special case we assume that all DTs in G require the same room. Hence, w.r.t. to the room resource all DTs have to be scheduled in a strictly sequential way. The beam will have substantial breaks. In this case the makespan can be estimated using the total time the respective room resource is required and some constant offset for the tasks outside of the room, i.e., by

$$\max\{\bar{P}_1 n_1, \bar{P}_2 n_2, \bar{P}_3 n_3\} + P^{\text{orb}} + P^{\text{ora}}. \quad (3)$$

Observe that because only one room is used exactly one term of the maximum function in (3) is greater than zero.

The second special case supposes that the DTs are provided in two rooms. DTs will be scheduled alternately between the two rooms. It can be assumed that the tasks apart the irradiation take in general longer than the irradiation itself. Consequently, there will be frequently breaks on the beam resource. In most cases, the makespan will be determined by the utilization of the room that is required the most. An estimation of the makespan for this special case is given by (3) again. In contrast to the previous scenario two terms of the maximum function are greater than zero.

The third special case assumes that the DTs are distributed evenly among the three treatment rooms. In such situations the rooms will be used in an interleaved way s.t. the beam cycles between all three rooms. In this way, the beam will typically be used most efficiently and it can be expected that the beam is used without idle time. The makespan can be estimated by the total time the beam resource is used plus a constant offset for the first and last scheduled DTs:

$$\bar{P}_B n_B + P^{\text{irb}} + P^{\text{ira}} + P^{\text{orb}} + P^{\text{ora}}. \quad (4)$$

In practice we will mostly have a mixture of the three discussed cases. A lower bound for the makespan can be derived by combining (3) and (4):

$$MS^{\text{LB}} = \max\{\bar{P}_1 n_1, \bar{P}_2 n_2, \bar{P}_3 n_3, \bar{P}_B n_B + P^{\text{irb}} + P^{\text{ira}}\} + P^{\text{orb}} + P^{\text{ora}}. \quad (5)$$

Equation (5) is a lower bound for the makespan since P^{orb} , P^{ora} , P^{irb} , and P^{ira} are the minimum durations that have to precede and follow the first and last use of the respective resources and the fact that the total resource requirement is a trivial lower bound. Basically, MS^{LB} assumes that there is a schedule without idle time on the resource that is used the most. Let $n_{\max} = \max_{r \in \{1,2,3\}} n_r$ and

$n_{\min} = \min_{r \in \{1,2,3\}} n_r$. We can expect MS^{LB} to be a tight estimate if either $n_{\max} \geq n_B - n_{\max} - 1$, i.e., one room clearly dominates, or $n_{\max} \leq n_{\min} + 1$, i.e., the DTs are evenly distributed among the three rooms.

To strengthen the estimation also for cases in-between, we consider the simplified scenario in which all DTs have exactly the same timing and resource requirements, except that they are distributed among the three rooms. A good schedule would certainly cycle between all three rooms, but not to an extent that remaining DTs have to be scheduled sequentially in a single room.

Let N_{123} be the maximal number of cycles between the three treatment rooms, such that all remaining DTs can be scheduled alternatingly between two rooms. In such a scenario, the following condition must hold:

$$n_{\max} - N_{123} - 1 = (n_{\min} - N_{123}) + (|G| - n_{\max} - n_{\min} - N_{123}). \quad (6)$$

The intuition of the formula above is to compare the number of DTs that remain in each room after cycling between all three rooms for N_{123} times. Note that the minus one represents the fact that the schedule might start and end with the room that is required the most. Equation (6) yields $N_{123} = |G| - 2n_{\max} + 1$. After excluding the corner cases where N_{123} becomes negative or larger than n_{\min} we obtain

$$N_{123} = \min(n_{\min}, \max(0, |G| - 2n_{\max} + 1)). \quad (7)$$

We can now strengthen the estimation of the makespan by using for N_{123} cycles between the rooms the estimation for the third special case and for the remaining DTs the estimation for the second special case as follows:

$$\begin{aligned} MS^{\text{ES}} = \max\{ & \bar{P}_B n_B + P^{\text{irb}} + P^{\text{ira}}, \\ & \bar{P}_1 n_1, \bar{P}_2 n_2, \bar{P}_3 n_3, \\ & 3\bar{P}_B N_{123} + \bar{P}_1 (n_1 - N_{123}), \\ & 3\bar{P}_B N_{123} + \bar{P}_2 (n_2 - N_{123}), \\ & 3\bar{P}_B N_{123} + \bar{P}_3 (n_3 - N_{123})\} + P^{\text{orb}} + P^{\text{ora}}. \end{aligned} \quad (8)$$

Notice that MS^{ES} is in contrast to MS^{LB} not a lower bound anymore.

5 Application of the Time Estimation in PTPSP

In this section the ideas developed in Section 4 will be used to obtain enhanced estimations for the total times the beam and each room is required. To this end, for a considered day $d \in D'$ let G be the set of all DTs assigned to day d . Since the beam and the rooms are normally available the whole day, we can assume that they have in general the same regular availability periods.

The total time the beam resource is required can be estimated almost analogously to (8) with the only difference that we have to disregard the time after

the last DT has stopped using the beam. Thus, in the estimation, given by

$$\begin{aligned} \widehat{\lambda}_{B,d} = \max\{ & \bar{P}_B n_B + P^{\text{irb}}, \\ & \bar{P}_1 n_1 - P^{\text{ira}}, \bar{P}_2 n_2 - P^{\text{ira}}, \bar{P}_3 n_3 - P^{\text{ira}}, \\ & 3\bar{P}_B N_{123} + \max(P^{\text{irb}}, \bar{P}_1 \cdot (n_1 - N_{123}) - P^{\text{ira}}), \\ & 3\bar{P}_B N_{123} + \max(P^{\text{irb}}, \bar{P}_2 \cdot (n_2 - N_{123}) - P^{\text{ira}}), \\ & 3\bar{P}_B N_{123} + \max(P^{\text{irb}}, \bar{P}_3 \cdot (n_3 - N_{123}) - P^{\text{ira}})\} + P^{\text{orb}}, \end{aligned} \quad (9)$$

we have to subtract P^{ira} whenever the room resources are used for the estimation.

The total time the rooms are needed is estimated by

$$\widehat{\lambda}_{r \in \{1,2,3\},d} = \max\{T_r n_r, 3\bar{P}_B N_{123} + \max(T_r(n_r - N_{123}), P^{\text{irb}} + P^{\text{ira}})\} + P^{\text{orb}}. \quad (10)$$

In contrast to the beam resource we can only use the considered room for the prediction. We can strengthen the estimation for the room resource that is used the most, i.e., for $r_{\max} = \arg \max_{r \in \{1,2,3\}} n_r$. This room is most likely the last one used and, hence, it is used at least as much as the beam resource. The strengthened estimation for room resource r_{\max} is then given by

$$\widehat{\lambda}_{r_{\max},d}^* = \max\{\widehat{\lambda}_{r_{\max},d}, \bar{P}_B n_B + P^{\text{irb}} + P^{\text{ira}}\}. \quad (11)$$

6 Computational Study

In this section we study the performance impact of applying the presented time estimation within our so far best performing approach, the enhanced iterated greedy (EIG) from [5]. Moreover, we determine the accuracy of the surrogate functions on final solutions.

All experiments are applied on the benchmark instances from [5] which resemble expected situations at MedAustron. The instances consider 50, 70, 100, 150, 200, and 300 therapies, which have to start within windows of 14 days. The beam and the three rooms are regularly available from $\widetilde{W}_d^{\text{start}}$ for 14 hours and have an extended availability period of 10 hours. Besides the beam and the rooms, there are further resources, such as the personnel, which are, however, sufficiently dimensioned to be not the primary reasons of substantial use of extended service time. A characteristic of the instances is that there is a ramp-up phase until the facility is used at full capacity followed by a wind-down phase until the last therapy is finished. At full capacity and strongly depending on the specific DTs there can be planned around 60 DTs.

Table 1 compares the performance between the EIG, as presented in [5], with the variant of the EIG where in the DA the time required from the beam and room resources phase is estimated by (9), (10), and (11). Both algorithms use as termination criterion a time limit of 20 CPU-minutes and are executed on each of the benchmark instances for 30 times. Table 1 shows the mean objective values obj and the median use of extended service periods $ext[h]$ in hours

Table 1. Average objective values obj and average use of extended service periods in hours $ext[h]$ of 30 runs with a time limit of 20 CPU-minutes and corresponding standard deviations $\sigma(obj)$ and $\sigma(ext)$ for EIG and EIG with time estimation.

Instance	EIG				EIG+TE			
	obj	$\sigma(obj)$	$ext[h]$	$\sigma(ext)$	obj	$\sigma(obj)$	$ext[h]$	$\sigma(ext)$
100-01	11.558	0.976	5.483	0.976	9.046	0.468	2.633	0.468
100-02	15.508	2.372	8.158	2.372	9.737	0.774	1.933	0.774
100-03	8.488	0.497	3.192	0.497	6.435	0.161	0.883	0.161
100-04	14.257	1.566	7.117	1.566	8.884	0.473	1.125	0.473
100-05	13.826	1.755	7.817	1.755	6.823	0.235	0.000	0.235
150-01	18.916	1.760	7.417	1.760	14.068	0.387	1.733	0.387
150-02	52.166	4.722	39.500	4.722	43.950	3.475	30.092	3.475
150-03	32.886	3.757	20.233	3.757	32.740	3.390	20.142	3.390
150-04	18.620	1.659	6.875	1.659	12.395	0.466	0.033	0.466
150-05	24.286	3.833	15.483	3.833	10.628	0.631	0.917	0.631
200-01	48.102	3.935	34.000	3.935	35.945	5.275	17.225	5.275
200-02	38.085	2.824	21.533	2.824	35.206	3.855	16.442	3.855
200-03	31.158	3.574	13.075	3.574	20.454	0.895	1.108	0.895
200-04	30.913	2.576	16.800	2.576	18.860	1.324	2.075	1.324
200-05	29.846	2.384	17.092	2.384	19.876	2.550	5.600	2.550
300-01	23.654	2.739	12.067	2.739	16.429	1.459	4.000	1.459
300-02	61.320	5.063	41.200	5.063	52.510	5.979	27.608	5.979
300-03	41.415	4.254	25.392	4.254	23.707	2.900	6.350	2.900
300-04	108.118	7.221	85.608	7.221	77.244	4.501	50.367	4.501
300-05	18.684	1.634	6.800	1.634	13.219	0.586	0.833	0.586

with the corresponding standard deviations $\sigma(obj)$ and $\sigma(ext)$ of finally obtained solutions. The results indicate that the application of the presented estimation considerably reduces the used extended service periods over all benchmark instances. The surrogate functions are, however, not necessarily a lower bound. Thus, we might occasionally overestimate the required time for the bottleneck resources yielding underutilized days. This has in general the consequence that the finishing day of therapies are delayed, which is penalized with the second term of our objective function. This raises the question whether this trade-off is indeed beneficial w.r.t. the objective function (1) using the same weights as in [5, 6]. This is indeed the case, since the EIG with the presented time estimation performs on all benchmark instances significantly better than the one without according to a Wilcoxon rank sum test with a significance level of 95%. The performance improvement can be explained by the increased accuracy of $\hat{\lambda}_{r,d}$. While using the trivial lower bound given by aggregated resource demands, the EIG's DA level underestimates on average the required time from the beam and the most used room by 27.7 and 101.6 minutes (i.e., by 6.5% and 15.3%) with a standard deviation of 20.2 and 57.2, respectively. With the presented estimation the DA level is on average off by 9.17 minutes for the beam and by 10.9 minutes

for the most used room (i.e., by 2% and 2.4%) with a standard deviation of 8.6 and 9.6, respectively.

7 Conclusion

In this work, we presented a surrogate model for estimating the total times the bottleneck resources required to optimally schedule sets of DTs. This surrogate model is applied to quickly estimate the use of extended service times at the upper DA level of the PTPSP. We evaluated the effects of the presented surrogate model in the so far best performing algorithm, the EIG from [5]. Results show that on all considered benchmark instances the use of extended service periods as well as the whole objective value can be significantly decreased. This can be explained by the substantial gain in accuracy of the new surrogate model and with it the better adjustment of the two levels.

The focus of this work is on resources that are shared by all DTs and are tightly coupled with the throughput of the facility. There are, however, certain resources, like the anesthetist, that are required by some DTs and are available only for the first half of the working day. The interaction of DTs requiring those resources are not considered so far and sometimes result in the use of extended service times that might be avoidable by further improvements.

References

1. Burke, E.K., Leite-Rocha, P., Petrovic, S.: An Integer Linear Programming Model for the Radiotherapy Treatment Scheduling Problem. CoRR abs/1103.3391 (2011)
2. Conforti, D., Guerriero, F., Guido, R.: Optimization models for radiotherapy patient scheduling. *4OR* 6(3), 263–278 (2008)
3. Kapamara, T., Petrovic, D.: A heuristics and steepest hill climbing method to scheduling radiotherapy patients. In: Proceedings of the International Conference on Operational Research Applied to Health Services (ORAHS). Catholic University of Leuven, Leuven, Belgium (2009)
4. 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)
5. Maschler, J., Hackl, T., Riedler, M., Raidl, G.R.: An Enhanced Iterated Greedy Metaheuristic for the Particle Therapy Patient Scheduling Problem. In: Proceedings of the 12th Metaheuristics International Conference. Barcelona, Spain (2017)
6. 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. pp. 223–244. Udine, Italy (2016)
7. 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)
8. Petrovic, S., Leite-Rocha, P.: Constructive and GRASP approaches to radiotherapy treatment scheduling. In: Advances in Electrical and Electronics Engineering - IAENG Special Edition of the World Congress on Engineering and Computer Science 2008. pp. 192–200. IEEE (2008)