Constructive Approaches to Radiotherapy Scheduling

Sanja Petrovic Pedro Leite-Rocha *

Abstract— **With the increase in the number of cancer cases and new government regulations for cancer treatment, radiotherapy scheduling has gained a lot of importance recently. In this paper, four constructive approaches to radiotherapy scheduling are introduced, as well as a GRASP based algorithm for improving the solution obtained by the constructive approaches. These algorithms are investigated on real-world data obtained from the Nottingham University Hospitals NHS Trust, UK.**

Keywords: Grasp, patient scheduling, radiotherapy

1 Introduction

The number of cancer cases in the United Kingdom has continually increased in the last decades. Every year 200,000 people are diagnosed with cancer in England and 120,000 people lose their lives to the disease [1].

The radiotherapy patient scheduling has been recognised as a key factor for increasing the quality of treatments, as it is of paramount importance to deliver the treatment by the imposed waiting time target and to enable consecutive treatment sessions without interruptions. Before starting a radiotherapy treatment, a patient needs to go through several phases, including localisation of treatment fields using a CT scanner or simulator, radiotherapy planning in which the dosage and the best way to deliver radiation is determined, and verification of a plan using a simulator. This paper addresses the problem of patient scheduling on linacs (linear accelerators), which are used to deliver radiation once the pre-radiotherapy treatment phases are completed.

To our knowledge, only few papers deal with the problem of scheduling radiotherapy treatments. Kapamara *et al.* give a review of radiotherapy patient scheduling problem (both pre-treatment and treatment), concluding that this problem is similar to a dynamic stochastic jobshop problem [2]. Conforti *et al.* define mathematical models for scheduling of radiotherapy treatments [3]. The objective in the proposed model is to schedule as many patients as possible in a short period of time (e.g. one week), considering treatment slots of equal length and one type of linac. However, this model does not consider all the constraints imposed on real-world radiotherapy scheduling. Lev and Caltagirone developed a discrete event simulation model of the radiotherapy scheduling problem [4]. Larsson offered an automated scheduling system, but relying on simple formulae rather than scheduling algorithms [5].

Some research work carried out on patient scheduling in other types of treatments are relevant to this work. Proctor *et al.* propose a simulation model for a radiotherapy centre in the UK [6]. They analyse two strategies to improve cancer waiting times: 1) acquiring more equipment, such as simulator and/or linac, and 2) changing working policy, such as not requiring radiographers treat the same patients for all sessions, increasing working hours of the linacs, etc. Thomas presents a model based on Monte-Carlo distribution to calculate the percentage of spare capacity required to keep waiting times to treatment short [7]. He analyses the outcome of the model if some parameters are changed, such as no treatment on bank holidays, the number of machines, etc. Holmberg and McClean give a workload indicator for treatment planning [8]. They developed a method to measure the capacity of radiotherapy treatments in a hospital/clinic, in order to predict the workload.

This paper deals with a real-world radiotherapy scheduling problem present in the Nottingham University Hospitals NHS Trust (UK). Our initial research results were concerned with two constructive algorithms for radiotherapy scheduling which first order the patients and then try to book the required number of treatment sessions either forward, starting from the first available date, or backward, starting from the latest possible date, i.e. the waiting time target [9]. In the research presented in this paper, we generalise these constructive approaches and compare the obtained results. In addition, we investigate the impact of a number of parameters on radiotherapy scheduling, namely the allocation of free room on linacs to be used for higher priority patients, the accumulation of patients before scheduling instead of scheduling on daily basis, and delay in scheduling patients closer to their waiting time targets. Finally, an algorithm based on the meta-heuristic GRASP is developed to investigate if

*Automated Scheduling Optimisation and Planning Research Group, School of Computer Science, University of Nottingham, Nottingham, UK. {sxp,plr}@cs.nott.ac.uk The authors would like to thank the Engineering and Physics Sciences Research Council (EPSRC), UK, for supporting this research (Ref No. EP/C549511/1).

it can improve the solutions obtained by the constructive approaches.

The paper is organised as follows. Section 2 defines the radiotherapy treatment scheduling problem considered in the research. In Section 3, the four constructive approaches are described and experiments are carried with each of them using real-world data obtained by the hospital. In Section 4, a GRASP based algorithm and its results are presented. Section 5 gives a summary of the paper and points directions for future work.

2 Radiotherapy Treatment Scheduling

A radiotherapy treatment consists of exposing the patient to radiation, with the intent of destroying the cancerous cells while minimising damage to the surrounding organs. To allow enough time for the healthy organs to recover, the radiation is divided in sessions. To each patient, a number of sessions is assigned required to achieve the prescribed dose of radiation. Each patient is allocated the energy type of linac to be treated with (low i.e. 6MV, high i.e. 10MV or electron). Patients are grouped into different categories based on the treatment priority (emergency, urgent or routine) and treatment intent (radical, palliative or adjuvant), each of which is given a waiting time target (due date) that is measured from the time when a course of radiotherapy is first recommended by a clinical oncologist and it actually starting.

Each day, a number of patients enter at the booking system. Arriving at the system means that the details of the patient's treatment have been decided, and a schedule can be created for that patient. The schedule is partially booked with patients that entered on the previous days.

Radiotherapy treatment scheduling is defined as the problem of scheduling a number of sessions for each patient that enter the booking system on the given day n, $n = 1, \ldots, N$, where N is the number of days considered in the scheduling horizon subject to the following constraints. The utilisation of each linac does not exceed its capacity measured by the amount of hours that the working shifts consist of. Each session lasts a pre-defined time depending on the type of energy required. Patients are available for scheduling after their release time, when all pre-radiotherapy phases are completed. Only emergency patients can have sessions on weekends, while the others must be scheduled on week days. If a patient has 5 or less sessions, they must all be scheduled on the same week. Palliative patients must have at least 2 sessions before a weekend. The sessions for a patient must be scheduled on consecutive days (weekends are not counted as interruptions for non-emergency patients).

The quality of the constructed schedule is measured by the average weighted tardiness of patients, where tardiness occurs if a patient breaches the waiting time target. Weights of the patients are determined by their priority, i.e. they are set as 10, 3, 1 for emergency, urgent and routine patients, respectively. The objective is:

$$\text{minimise} \left(\frac{1}{P} \cdot \sum_{p=1}^{P} w_p \cdot \rho_p \right), \tag{1}$$

where P is the number of patients that entered the booking system ($P = P_1 + P_2 + \ldots + P_N$), P_n, $n = 1, \ldots, N$ is the number of patients on day N, and w_p and ρ_p are the weight and tardiness of patient p, respectively.

3 Constructive Approaches

Four constructive approaches to radiotherapy scheduling are developed. On each day, they schedule patients available for scheduling, while the performance of the algorithm is measured at the end of the scheduling horizon. They consist of two phases. In the first phase, the patients are ordered lexicographically by their due date, priority and the required number of sessions, in the sense that the first criterion for ranking is due date, while the ties are broken by the priority and then by the number of sessions. The algorithms differ from each other in second phase, in which the patients are scheduled from the ordered list produced in the first phase. The remainder of the section gives the description of the approaches and the experimental results.

3.1 Target Approach

The target approach is a generalisation of the two algorithms presented in [9]. These algorithms operate in a forward (backward) manner from the release date r_p (due date d_p) of each patient, trying to schedule the required number of sessions subject to the given constraints. If it is not possible to accommodate all the required sessions, the algorithms move the start day forward (backward) and try again. This is repeated until the patient is scheduled or all days between r_p and d_p have been tried. Finally, if this fails, then the start day is moved to the first available day after d_p.

The developed generalised approach tries to schedule the treatment as close as possible to a pre-defined date within the $[r_p, d_p]$ window. The parameter *target* determines the first day to try to schedule the patient on in the following way:

$$\text{first\_day\_to\_try} = r_p + (d_p - r_p) \cdot target. \tag{2}$$

If this day fails, the algorithm tries the days in the vicinity of the *target* date within the $[r_p, d_p]$ window. If it is not possible to schedule the patient in the specified window, the start day is moved to the first available day after d_p This algorithm enables investigation of the effects of scheduling of patients of different priorities closer to their release or due date.

Experimental Data. Real-world data about radiotherapy patients from the hospital in the period of 5 years were available for experiments. They were used to create 33 different sets of experimental data with 18 months of data each. However, the results obtained with the first 6 months of data will not be presented, since they are used only to partially fill the booking table in a realistic way. The results obtained with the remaining 12 months will be used to evaluate the performance of the algorithms ($N = 365$).

The oncology ward contains 4 linacs, two of which emit high energy, one emits low and one emits electron radiation. Linacs are available from 8:45 to 18:00, Monday to Friday, while weekend sessions 9:00 to 13:00 on Saturdays and Sundays are for emergency patients only. High energy radiation sessions have duration of 15 minutes, while the low energy and electron sessions last for 12 minutes, with the exception of the first session of all radiation types which is 5 minutes longer. In our experiments, we did not use the available data about patients directly, but generated the data with the same characteristics. It enables us to conduct a greater number of experiments and to vary some characteristics and investigate their effect on the result (such as the mean of patients per day). The number of patients per week day and weekend day follow normal and exponential distributions, respectively. Using the available data, the probabilities of the following parameters were calculated: the priority of each patient, the intent of treatment for each different priority, the required type of radiation for each priority, and the required number of sessions for each required type of radiation. The waiting times targets are set as 2 days for emergency patients, 14 for palliative and 28 for radical [10]. Due to lack of data from the hospital, the waiting times targets are calculated from the release date instead of decision to treat date.

Experiments and Results. Experiments are run on the sets of generated data with different values of *target* parameter (0.0, 0.33, 0.66 and 1.0) for urgent and routine patients, while for emergency patients the *target* is set as 0.

The inter quartile mean (IQM) [11] of the average weighted tardiness of patients and standard deviations are given in Figure 1 to present the performance of the algorithm obtained with different values of *target* parameter. The IQM is equal to the mean discarding the 25% smaller values and the 25% largest values. It was chosen for being insensitive to outliers (like the median), but also being based on a larger number of observations (like the mean).

It can be observed that the target approach performs better when it has a *target* value close to 0 for both urgent and routine patients, i.e. patients are scheduled as close to their release date as possible. Also, the standard deviation of the objective function value is smaller for these *target* values. This is contradictory with the results found in [9], in which the backward algorithm starting from the due dates of patients performed better. We believe the results differ because of the way the two experiments were conducted. In this paper, each experiment starts from an empty schedule, while the first 6 months are discarded to avoid start-up effects. In the previous experiments, one month of data were used starting from a schedule previously filled with a pre-defined utilisation level. Three levels were introduced, light, normal and heavy, within which the load of the first day is 90, 95 and 98% respectively of the available capacity, on the second day, it is 75, 80 and 85 and then it decreases by 5 on every following day [9]. We conclude that the way and time period used to partially fill in the booking system affects the results.

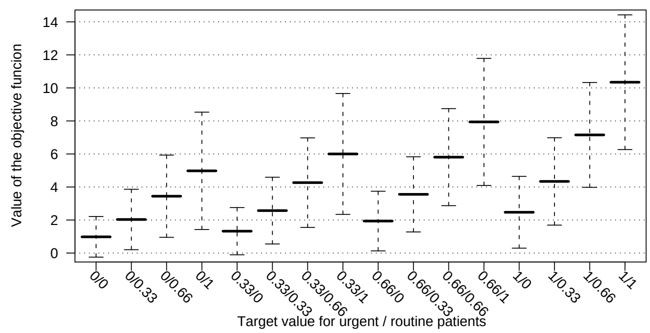

Figure 1: IQM of the objective function with different target values for urgent and routine patients

3.2 Utilisation Threshold Approach

In this approach, a *threshold* of machine utilisation is defined for each priority of patients. It means that when the utilisation of a machine reaches the specified *threshold* for a given patient priority, then no more patients of that priority can be scheduled on that machine on that day, thus leaving the machine available for patients of different priority. Using this algorithm with varied *thresholds* for different patient priorities we can investigate whether allocating more time on the machines for patients of higher priority can lead to better schedules.

Experiments and Results. In these experiments, the *target* for urgent and routine patients is defined to be 0, as it was the best configuration found in the previous experiments.

The values used as *threshold* are 100%, 98% and 96% for urgent patients and 100%, 98%, ..., 90% (with a discretisation step of 2%) for routine patients, whilst the *threshold* for emergency patients was set to be 100%. The other lower values were tried for urgent patients, but did not lead to good results. The results are presented in Figure 2.

ISBN: 978-988-98671-0-2

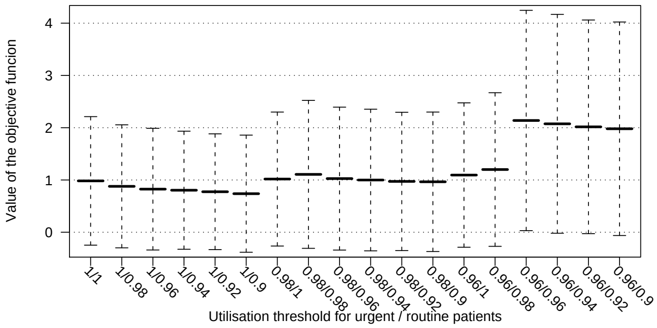

Figure 2: IQM of the objective function with different *threshold* values for urgent and routine patients

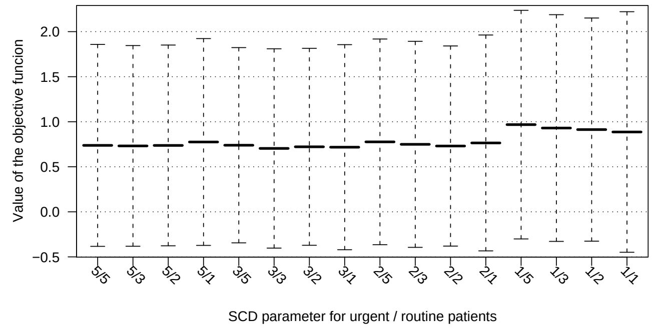

Figure 3: IQM of the objective function with different *SCD* values for urgent and routine patients

The best results occur when the *threshold* is 100% and 90% for urgent and routine patients respectively, i.e., when 10% of the time on the machine is reserved only for emergency and urgent patients. When the *threshold* value for both categories is lower than 96% the solution is greatly worsen. We believe this is due to the smaller number of emergency patients, and therefore it does not lead to better schedule to reserve more time exclusively for emergency patients.

3.3 Schedule Creation Day (SCD) Approach

In the previous approaches, schedules were created for patients on the same day they enter. In the third constructive approach, which we named Schedule Creation Day (SCD) approach, we select specific days of the week for each patient priority when a schedule can be created. If a patient arrives on a day when it is not specified to create a schedule for that priority, the schedule will be created on the first following allowed day.

The motivation behind this approach is to investigate whether the accumulation of patients to be scheduled will lead to better schedules. Obviously, the search space becomes larger and it may lead to solutions of higher quality.

A parameter introduced in this approach denotes the day when the patients of a given priority can be scheduled. Possible values of the parameter are: 7 - applicable only to emergency patients, any day of the week, 5 - any nonweekend day, 3 - Mondays, Wednesdays and Fridays, 2 - Tuesdays and Fridays, 1 - only Fridays.

Experiments and Results. In these experiments, we set the values of the *target* for urgent and routine patients to be 0, the *threshold* for emergency and urgent patients to be 100% and for routine 90%, as this was the best configuration found in the previous experiments. The values of the *SCD* parameter were varied for urgent and routine patients, whilst it was kept at 7 for emergency patients. The results are presented in Figure 3.

The lowest value of the objective function occurs when the *SCD* parameter for both priorities is set to 3, i.e. when urgent and routine patients can have their schedule created only 3 times a week, on Mondays, Wednesdays and Fridays, and not on every day as it is the current practice in the hospital.

3.4 Maximum Number of Days in Advance (MNDA) Approach

The fourth constructive approach introduces waiting for some days before creating a schedule for the patient after he/she has arrived. The maximum number of days in advance (*MNDA*) parameter is introduced, which specifies the maximum number of days before a patient's due date as the day when a schedule can be created. In this way, it is possible to reserve more space for the patients of higher priorities in the earlier dates.

Experiments and Results. The *MNDA* parameter of urgent and routine patients assumes values of 28, 21, 14 and -1 in the experiments, where a *MNDA* of -1 means that there is no limit and a schedule can be created for the patient as soon as he/she arrives. Experiments are conducted with the values of the *target* for urgent and routine patients set to be 0, the *threshold* for emergency and urgent patients to be 100% and for routine 90%, and the *SCD* parameter to be 7 for emergency patients and 3 for urgent and routine patients, as this was the best configuration found in the previous experiments. Results are presented in Figure 4.

The best results were obtained with a *MNDA* parameter set to be 28 for urgent and routine patients. The only patients arriving more than 28 days before their due date are the radical urgent and routine patients, who arrive 29 days before. These represent around 53% of the patients. This suggests that it might be a good idea to not create a schedule for the radical non-emergency patients as soon as they arrive, but to wait one day after they arrive to do it. This way, it is possible to create a better schedule for the other patients first.

ISBN: 978-988-98671-0-2

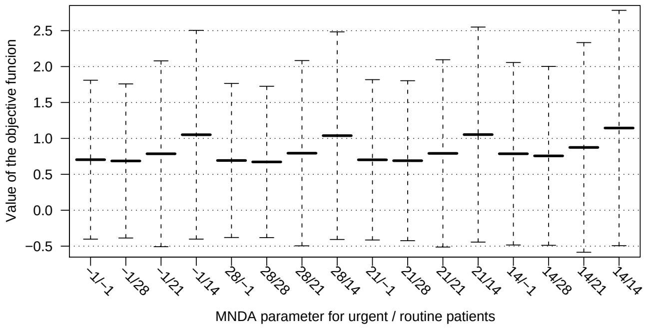

Figure 4: IQM of the objective function with different *MNDA* values for urgent and routine patients

4 GRASP Based Approach

While investigating the behaviour of different constructive approaches, an interesting question arises, whether the constructed schedules can be improved by a metaheuristic approach. A GRASP (Greedy Random Adaptive Search Procedure) [12] algorithm is developed which consists of two phases that are repeated for a number of iterations. The pseudo-code for GRASP is presented in Figure 5.

```
1  orderedList ← OrderPatients(listOfPatients);
2  bestSol ← ConstrApproach(orderedList);
3  P ← LengthOf(orderedList);
4  for i ← 1 to iterations do
       /* First phase of GRASP                          */
5      for p ← 1 to P do
6          q ← RandomExpDistBetween(1, P − p + 1);
7          randomList[p] ← orderedList[q];
8          RemoveFromList(q, orderedList);
9      endfor
10     randSol ← ConstrApproach(randomList);
       /* Second phase of GRASP                         */
11     for j ← 1 to localSearchIterations do
12         p1 ← RandomLatePatient();
13         p2 ← RandomPatientScheduledBefore(p1);
14         if SwapIsFeasible(p1, p2) and
           SwapIsImprovement(p1, p2) then
15             Swap(p1, p2);
16         endif
17     endfor
       /* Update best solution                          */
18     if ValueOf(randSol) < ValueOf(bestSol) then
19         bestSol ← randSol;
20     endif
21 endfor
```

Figure 5: GRASP

In the first phase, patients are ordered lexicographically in the same way as in the constructive approaches (line 1). The list of patients is then randomly re-ordered according to an exponential distribution with rate 1, i.e. each position in $randomList$ is filled with a random patient from $orderedList$, where the probability of the q^{th} patient from $orderedList$ to be placed in the p^{th} position in $randomList$ is proportional to e^{-q} (lines 5-9). Then, the constructive approach is called to create an initial solution using the new re-ordered list (line 10).

In the second phase, after the initial solution has been constructed, a local search is applied. It randomly chooses a patient $p1$ who breached the waiting time target (line 12) and tries to swap it with a randomly chosen patient $p2$ among the patients who require the same type of radiation and starts his/her treatment on an earlier date than $p1$ (line 13). If the swap is feasible, i.e. all sessions of both patients can be scheduled without violating the constraints and the resulting solution is better, the new solution is kept (lines 14-16). Otherwise, the swap is discarded. This is repeated for a number of local search iterations.

If the resulting solution at the end of an iteration is better than the best solution found, the best solution found is updated (lines 18-20).

Experiments and Results. To evaluate GRASP as a method for improving the solution obtained by the constructive approach, the experimental environment is used as with the constructive approaches. The constructive approach is used within the GRASP method with the best configuration of parameters found in the previous experiments: *target* as 0 for all priorities, *threshold* as 100% for emergency and urgent patients and 90% for routine patients, *SCD* parameter as 7 for emergency and 3 for urgent and routine patients, and *MNDA* parameter as -1 for emergency and 28 for urgent and routine patients. GRASP is run 33 times with different random seeds for each one of the 33 generated instances, which gives in total 1089 executions. Both the number of iterations within each run of GRASP and the number of local search iterations are set as 100. Figure 6 presents the histogram of improvement of the objective function achieved by GRASP over the one found by the constructive approach alone. The improvement is defined as

$$\text{improvement} = 1 - \frac{\text{GRASPSol}}{\text{ConsSol}}, \qquad (3)$$

where GRASPSol and ConsSol are the values of the objective function achieved by GRASP and the constructive approach, respectively. The cases where this improvement is equal to 0 (39% of the cases) have been excluded from the histogram in Figure 6.

In 38% of the experiments GRASP achieved an improvement greater than 0%, and in 23% the improvement was negative, meaning that GRASP caused the schedule to

ISBN: 978-988-98671-0-2

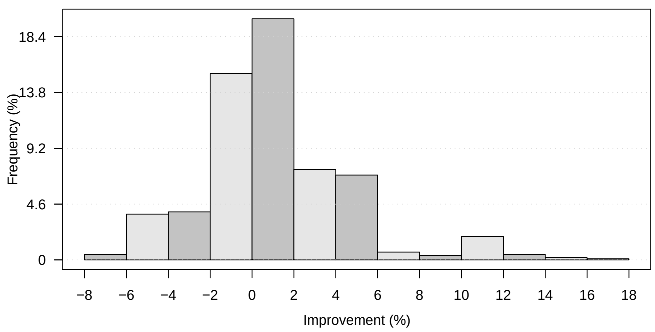

Figure 6: Histogram of improvement achieved by GRASP

be worse than the one found by the constructive approach. This may happen because, even though GRASP will takes the best solution for a given day, that solution may not be the best starting schedule in the succeeding days.

5 Summary

In this paper, the problem of scheduling the treatments of radiotherapy patients in the Nottingham University Hospitals NHS Trust is investigated. Four constructive approaches were presented, which consider the start day of the treatment sessions, the reservation of a machine capacity for patients of a given priority, the allocation of certain days of the week when schedules can be created, and the definition of a maximum time window before the patient's breach date when the schedule can be created. An algorithm based on GRASP was also developed in order to try to improve the schedules created by the constructive approaches.

Future work includes the investigation of look ahead techniques to try to anticipate how many number of patients of each category might arrive in the succeeding days and implementing a dynamic way of calculating the parameters used in the constructive approaches depending on the input data. Additional experiments with different mean of patients per day are also included in future work.

So far, we have considered allocation of treatment days to patients. Next step will be to consider actual time of treatments to include more constraints to realistically capture the real-world radiotherapy scheduling, such as patient preference for being treated in morning or afternoon sessions, requirements for transport to and from the hospital, etc.

6 Acknowledgement

The authors would like to acknowledge the support of the Nottingham University Hospitals NHS Trust (UK), for giving the data-sets used in the experiments.

References

[1] UK Department of Health. The NHS cancer plan: a plan for investment, a plan for reform. Technical report, NHS (National Health Services), United Kingdom, September 2000.

[2] T. Kapamara, K. Sheibani, O. C. L. Haas, D. Petrovic, and C. R. Reeves. A review of scheduling problems in radiotherapy. In *Proceedings of the International Control Systems Engineering Conference (ICSE 2006)*, Coventry, U. K., 2006. Coventry University Publishing.

[3] D. Conforti, F. Guerriero, and R. Guido. Optimization models for radiotherapy patient scheduling. to appear in 4OR: A Quarterly Journal of Operations Research, 2007.

[4] B. Lev and R. J. Caltagirone. Evaluation of various scheduling disciplines by computer systems. In *WSC '74: Proceedings of the 7th conference on Winter simulation*, volume 1, pages 365–370, New York, NY, USA, 1974. ACM.

[5] S. N. Larsson. Radiotherapy patient scheduling using a desktop personal computer. *Clinical Oncology*, 5:98–101, 1993.

[6] S. Proctor, B. lehaney, C. Reeves, and Z. Khan. Modelling patient flow in a radiotherapy department. *OR Insight*, 20(3):6–14, September 2007.

[7] S. J. Thomas. Capacity and demand models for radiotherapy treatment machines. *Clinical Oncology*, 15(6):353–358, 2003.

[8] O. Holmberg and B. McClean. A method of predicting workload and staffing level for radiotherapy treatment planning as plan complexity changes. *Clinical Oncology*, 15(6):359–363, 2003.

[9] S. Petrovic, W. Leung, X. Song, and S. Sundar. Algorithms for radiotherapy treatment booking. In *Proceedings of the 25th Workshop of the UK Planning and Scheduling Special Interest Group (PlanSIG2006)*, pages 105–112, Nottingham, UK, 14-15th December 2006.

[10] P. Barrett-Lee. Reducing delays in cancer treatment: some targets. Technical report, Royal College of Radiologists, 1993. from the Joint Council for Clinical Oncology.

[11] Wikimedia Foundation, Inc. *Wikipedia, the Free Encyclopedia*, chapter Interquartile mean, page http://en.wikipedia.org/wiki/Interquartile_mean. Wikimedia Foundation, 20 April 2008.

[12] T. A. Feo and M. G. C. Resende. Greedy randomized adaptive search procedures. *Journal of Global Optimization*, 6(2):109–133, March 1995.