Dynamic Multi-Appointment Patient Scheduling for Radiation Therapy

Antoine Sauré[a,*], Jonathan Patrick[b], Scott Tyldesley[c], Martin L. Puterman[a]

[a] *Sauder School of Business, University of British Columbia, , 2053 Main Mall, Vancouver, BC V6T 1Z2, Canada*
[b] *Telfer School of Management, University of Ottawa , 55 Laurier Avenue East, Ottawa, ON K1N 6N5, Canada*
[c] *British Columbia Cancer Agency , 600 West 10th Avenue, Vancouver, BC V5Z 4E6, Canada*

Abstract

Seeking to reduce the potential impact of delays on radiation therapy cancer patients such as psychological distress, deterioration in quality of life and decreased cancer control and survival, and motivated by inefficiencies in the use of expensive resources, we undertook a study of scheduling practices at the British Columbia Cancer Agency (BCCA). As a result, we formulated and solved a discounted infinite-horizon Markov decision process for scheduling cancer treatments in radiation therapy units. The main purpose of this model is to identify good policies for allocating available treatment capacity to incoming demand, while reducing wait times in a cost-effective manner. We use an affine architecture to approximate the value function in our formulation and solve an equivalent linear programming model through column generation to obtain an approximate optimal policy for this problem. The benefits from the proposed method are evaluated by simulating its performance for a practical example based on data provided by the BCCA.

Keywords: Scheduling, OR in health services, Markov decision processes, Linear programming, Approximate dynamic programming

---

*Corresponding author. Tel.: +1 604 632 4241; fax: +1 604 822 9574.
Email addresses: antoine.saure@sauder.ubc.ca (Antoine Sauré), patrick@telfer.uottawa.ca (Jonathan Patrick), styldesl@bccancer.bc.ca (Scott Tyldesley), martin.puterman@sauder.ubc.ca (Martin L. Puterman)

1. Introduction

External beam radiation (hereafter referred to as radiation therapy) is a cancer treatment that uses high-energy rays to kill or shrink tumor cells. It is the principal therapy for most types of cancer, but it is also used in combination with other treatments and therapies (e.g. chemotherapy, hormonal therapy and surgery). When a cure is not possible, radiation therapy can be used for palliative purposes. In British Columbia, approximately 52% of cancer patients require radiation therapy some time during the course of their illness and 40% receive radiation therapy within five years of diagnosis (Source: BCCA registry and treatment databases).

Long standing evidence suggests that delays in radiation therapy are associated with tumor progression, persistence of cancer symptoms, psychological distress and decreased cancer control and survival rates (O'Rourke and Edwards, 2000; Fortin et al., 2002; Waaijer et al., 2003; Coles et al., 2003; Mackillop, 2007; Chen et al., 2008). For this reason, many cancer institutions around the world have adopted wait time benchmarks for the time from when the patient is ready to begin treatment to the start of it. In Canada, the maximum acceptable wait suggested by the Canadian Association of Radiation Oncologists for all non-emergency and non-urgent cases is 14 days (Norris, 2009). Unfortunately, fewer than 75% of the radiation therapy treatments in British Columbia are initiated within this time frame (see Table 1).

Table 1: Service levels for patients who received radiation therapy in British Columbia between 2004 and 2008. Source: BCCA registry and treatment databases.

year	patients	% initiated within		
		14 days	28 days	56 days
2004	9,834	76.4%	97.5%	100.0%
2005	10,144	71.3%	96.2%	99.9%
2006	10,168	73.8%	97.5%	100.0%
2007	10,487	77.0%	96.3%	99.9%
2008	10,318	70.0%	96.2%	99.9%

Delays in radiation therapy are a direct consequence not only of an imbalance between capacity and demand but also a result of inefficient patient scheduling. Three relevant aspects make scheduling radiation therapy treatments especially challenging. First, radiation therapy treatments can be classified into multiple types. The classification is usually made on the basis of cancer site, treatment intent and urgency level. Second, radiation therapy treatments are spread out over time. For most types of cancer, radiation therapy is delivered in daily consecutive sessions for a time period that may vary between 1 day and 8 weeks, with breaks on weekends. Third, radiation therapy sessions do not necessarily have the same duration. Each session is scheduled for a time period ranging from 12 to 60 minutes. The combination of these three aspects of radiation therapy treatment means that a simple first-come-first-served policy will inevitably perform very poorly. This is either because later arriving higher priority demand would be forced to wait longer or else because the uneven session lengths would create unusable gaps in the system resulting in wasted capacity (not unlike in the game Tetris). Thus, due to the importance of timely access to care (as mentioned above) as well as the difficulty of determining intelligent schedules through simpler means, more sophisticated mathematical models are necessary than have hitherto been used for this problem.

To that end, we formulate the radiation therapy appointment scheduling problem as a discounted infinite-horizon Markov decision process (MDP). To deal with an intractable number of variables and constraints, we first approximate the value function in the equivalent linear programming formulation of the MDP using an affine architecture and then solve the dual of the resulting approximate linear programming model through column generation. From the solution we derive an approximate optimal booking policy which we test via simulation. We assume that the machines used for radiation therapy – the treatment units – do not differ significantly and thus that treatments can be delivered by any unit. The total treatment capacity is determined by aggregating individual capacities from multiple machines. This assumption is realistic considering the characteristics of the facilities analyzed for this study, where either the treatment units are identical or they will become operationally identical through scheduled replacements.

From a methodological point of view, our work represents a significant extension of the dynamic multi-priority patient scheduling model and solution approach developed by Patrick et al. (2008). In addition to multiple priority types, we consider patients who receive treatment across multiple days and for irregular lengths of time (see Figure 1). Furthermore, we allow the possibility of using overtime on different days of the planning horizon, and not necessarily for entire treatments. These additional complications are essential for any realistic attempt to model the scheduling of radiation therapy treatments. The new dimensions to the problem and the possibility of enlarging the system capacity through the use of overtime, together

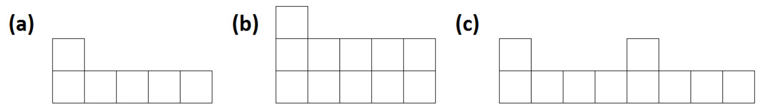

Figure 1: Graphical representation of some possible treatments patterns. Pattern (b), for example, represents a treatment consisting of a first session of three appointment slots and four additional sessions of two appointment slots each.

with the non-convex nature of the overtime cost, make this problem much more difficult to model and solve. To the best of our knowledge, ours is the first paper to incorporate multi-priority and multi-appointment requirements when optimizing advance scheduling policies in a dynamic setting. Our work also constitutes a novel application of approximate dynamic programming to a problem that has received very limited attention in the operations research literature.

The paper is organized as follows. Section 2 summarizes the literature relevant to our work. Section 3 provides a detailed description of the radiation therapy appointment scheduling problem. Sections 4 and 5 describe the proposed MDP model and solution approach, respectively. In Section 6 we provide some managerial insights based on applying our methodology in a small sample problem. We also evaluate the benefits of the proposed method by simulating its performance in a more practical example based on BCCA patient data. Finally, Section 7 states our main conclusions and suggests possible extensions.

2. Literature review

The literature on patient scheduling is usually classified as either *allocation* scheduling or *advance* scheduling. Allocation scheduling refers to methodologies for assigning specific resources and starting times to patients, but only once all patients for a given service day have been identified. Advance scheduling, on the other hand, refers to methodologies for scheduling patient appointments in advance of the service date, when future demand is still unknown. Most studies in the patient scheduling literature address allocation scheduling problems. Magerlein and Martin (1978), Cayirli and Veral (2003), Mondschein and Weintraub (2003), Gupta and Denton (2008), Cardoen et al. (2010) and Begen and Queyranne (2011) provide comprehensive reviews. Our work, however, falls into the second class of problems.

Even though patient scheduling problems have been studied extensively over the last decade, the allocation of medical capacity in advance of the service date and in the presence of multiple types of patients has received limited attention. Rising et al. (1973) describe a case study where simulation models are used to evaluate alternative decision rules for scheduling appointments in an outpatient clinic. The authors consider two types of patients, walk-ins and patients with advance appointments, and focus on the effects of different booking rules on patient throughput and physician utilization. Gerchak et al. (1996) study the advance scheduling of elective surgeries at an operating room when a random portion of the daily capacity is devoted to emergency surgeries. They analyze the tradeoffs between capacity utilization and delays and prove that the optimal capacity allocation policy is not necessarily a control limit policy. Lamiri et al. (2008) extend the model developed by Gerchak et al. by specifying the set of elective surgeries to be performed in each day of a planning horizon and propose a solution method combining Monte Carlo simulation with mixed-integer programming. Vermeulen et al. (2009) develop a dynamic although case-specific method for scheduling CT-scan appointments at a radiology department. They show through the use of simulation that simple rules based on comparing available and needed capacity for different patient groups allow a significant improvement in the number of patients scheduled on time. Schütz and Kolisch (2012) adopt a revenue management approach to address the problem of determining whether or not to accept requests for examinations on an MRI scanner. Although they consider a single booking period, requests coming from different patient types are allowed in advance and during the examination day.

Most related to our work are the papers of Patrick et al. (2008) and Erdelyi and Topaloglu (2009). They study a problem that involves allocating a fixed amount of daily capacity among entities with different priorities. Erdelyi and Topaloglu consider a finite planning horizon and focus on a class of policies that are characterized by a set of protection levels, whereas Patrick et al. consider an infinite planning horizon and seek optimal scheduling policies using approximate dynamic programming. In both cases, the authors assume that each entity to be served consumes only one unit of capacity. None of the studies mentioned above considers multi-appointment requirements.

The majority of the contributions from the operations research community to radiation therapy are associated with treatment planning (Mišić et al., 2010; Kim et al., 2012; Lim and Cao, 2012). The radiation therapy appointment scheduling problem, in particular, has attracted only limited attention.

Conforti et al. (2010) describe an integer programming model intended to support radiation therapy scheduling decisions. Their model includes both multiple priority classes and multiple capacity requirements but differs from ours in that their model is static and thus provides an operational and myopic solution. By myopic we mean a solution that considers only what is known at the present, ignoring the impact of today's decisions on the future. In contrast, we provide a dynamic policy that integrates future cost over an infinite horizon. We choose an infinite planning horizon because we are interested not only in generating specific patient schedules but also in identifying and understanding the properties of good booking policies.

3. Appointment scheduling

The appointment scheduling process for radiation therapy is the process by which available treatment capacity is assigned to incoming demand. It relies on the expertise of one or two booking agents. Each day the booking agent receives requests for appointments for treatments classified into I different pre-specified types. Treatments are classified on the basis of cancer site, treatment intent, urgency level and capacity requirements – that is the number of daily consecutive sessions into which the requested treatment is divided and the number of appointment slots required to deliver each session. Appointment slots are typically 12 or 15 minutes long and quite often the initial session is one appointment slot longer than the rest of the sessions in order to provide additional time for patient and system setup. Each treatment type is associated with a vector $\vec{r_i} = \{r_{ij}\}_{j=1}^{l_i}$, where l_i represents the number of sessions of a treatment of type i. This vector describes the duration, in appointment slots, of each of its sessions. For example, $\vec{r} = (2, 1, 1, 1, 1)$ represents a treatment consisting of an initial session of two appointment slots followed by four additional sessions of one appointment slot each (see Figure 1a). Daily demand for each treatment type i follows a discrete probability distribution with mean m_i requests per day. The demand distribution is independent for each type and does not change over time. For modeling convenience, we assume scheduling decisions are made once a day and the booking agent may schedule treatments at most N days in advance. Assuming that decisions are made once a day is not too far from reality. In practice, booking decisions usually take place in batches a couple of times a day. In this manner, the booking agent – typically a radiation therapist – is able to assess the future workload of the system, take into account clinical considerations and prioritize urgent cases.

Once the booking requests for a given day are known, the booking agent observes the number of appointment slots scheduled from today to the end of the planning horizon and determines how to allocate the available treatment capacity to the waiting demand. As an alternative action, he/she may decide to postpone some of the booking decisions to the next day. The number of appointment slots booked on any given day cannot exceed $C_r + C_o$ units. Parameters C_r and C_o are fixed and correspond to the daily regular-hour and overtime capacities respectively. Since the booking agent may schedule treatments at most N days in advance, the number of days in the planning horizon, M, must be large enough to allow the completion of any treatment initiated on day N. That is $M = N + \max_i\{l_i\} - 1$.

Under current practice, these booking decisions are made without explicitly considering that some treatments can be initiated further into the future allowing the appointments for more urgent treatments to be scheduled earlier. The presence of highly variable demand, limited treatment capacity, multiple urgency levels and multiple appointment requirements make it impossible for the scheduler to adequately assess the real impact of today's decisions on the future performance of the system. This lack of foresight generates several inefficiencies that usually translate into unnecessary delays, an unsystematic prioritization of patients, a larger number of isolated appointment slots that cannot be used and a higher overtime utilization. The model developed in this paper seeks to provide the scheduler with a means of adequately assessing the future impact of today's decisions in order to more intelligently allocate capacity.

4. A Markov decision process model

In this section, we formulate the radiation therapy appointment scheduling problem as a discounted infinite-horizon MDP model. We expand the dynamic multi-priority patient scheduling model developed by Patrick et al. (2008) by introducing multiple appointment requests, multiple session durations and allowing parts of the appointments to be delivered using overtime.

4.1. Decision epochs

Scheduling decisions are made at the end of each day.

4.2. State space

At the end of each day, the booking agent has access to the current schedule from today to the end of the planning horizon as well as the number of treatments of each type waiting to be booked. The current schedule provides the booking agent with the number of regular-hour and overtime slots available on each day of the planning horizon. Thus, a state of the system, denoted by $\vec{s} \in S$, takes the following form:

$$\vec{s} = (\vec{u}, \vec{v}, \vec{w}) = (u_1, \ldots, u_M, v_1, \ldots, v_M, w_1, \ldots, w_I)$$

Here u_m represents the number of regular-hour appointment slots already booked on day m, v_m the number of overtime slots already booked on day m and w_i the number of treatments of type i waiting to be booked. The definition of $\vec{w}$ considers not only the new demand for treatment but also all treatments not scheduled previously. As a consequence of using a rolling planning horizon, at the beginning of each decision epoch $u_M = v_M = 0$.

In order to use the mathematical programming models described later, we require a finite state space. For this reason, we implicitly assume upper bounds to the number of treatments of each type waiting to be booked. The bounds are set sufficiently high so they are of little practical significance.

4.3. Action sets

At the end of each day, the booking agent must decide on which day to start each of the treatments waiting to be scheduled. In some cases, this implies delivering parts of some treatments using overtime. Alternatively, the agent may postpone to the next day the scheduling decisions for some treatments. Overtime and postponements are intended to relieve the stress on the system. Any action available to the booking agent can be represented by:

$$\vec{a} = (\vec{x}, \vec{y}) = (x_{11}, x_{12}, \ldots, x_{IN}, y_1, \ldots, y_M)$$

Here x_{in} represents the number of treatments of type i booked today to start on day n and y_m the number of overtime slots booked today on day m.

Three important observations follow. First, we have chosen to model the capacity utilization component of the state space using two state variables instead of one because y_m is a non-convex function of the total number of appointment slots already booked on day m and the capacity allocation decisions. In this way, we do not need to define any additional variables and constraints to deal with the non-convexity of the overtime cost. Second, the number of treatments of type i that remain unscheduled at the end of the day can be written as $w_i - \sum_{n=1}^{N} x_{in}$. Third, it is important to note that we are not assigning patients to specific appointment slots, we are just allocating capacity to each treatment request. Once capacity is allocated, a second level of scheduling is needed which assigns patients to specific appointment times and specific machines.

The set of feasible actions compatible with state $(\vec{u}, \vec{v}, \vec{w}) \in S$, denoted by $A_{(\vec{u},\vec{v},\vec{w})}$, must satisfy the following constraints:

$$\sum_{n=1}^{N} x_{in} \leq w_i \qquad \forall i \tag{1}$$

$$u_m + \sum_{i=1}^{I} \sum_{\substack{k=\max\{m \\ -l_i+1,1\}}}^{\min\{m,N\}} r_{i(m-k+1)} x_{ik} \leq C_r + y_m \qquad \forall m \tag{2}$$

$$v_m + y_m \leq C_o \qquad \forall m \tag{3}$$

$$x_{in} \in \mathbb{Z}^+ \quad \forall i, n \qquad y_m \in \mathbb{Z}^+ \quad \forall m$$

Constraint (1) limits the number of bookings for each treatment type to be less than or equal to the number of treatments waiting to be booked. Constraint (2) restricts the total number of appointment slots booked today for day m to be less than or equal to the available treatment capacity that day. This is equivalent to ensuring that the number y_m of overtime slots booked today for day m is sufficient to cover the new bookings made for that day. Constraint (3) limits the total overtime utilization on day m to be less than the overtime capacity. All action variables are positive and integer.

4.4. Transition probabilities

Once all scheduling actions are taken, the only source of uncertainty in the transition to the next state of the system is the number of new requests for each type of treatment. Thus, as a result of choosing booking action $\vec{a} = (\vec{x}, \vec{y})$ in state $\vec{s} = (\vec{u}, \vec{v}, \vec{w})$, $\vec{a} \in A_{\vec{s}}$ and $\vec{s} \in S$, and having q_i new requests of type i, the state of the system the next day, denoted by $\vec{s'} = (\vec{u'}, \vec{v'}, \vec{w'})$, will be determined by the following probability distribution:

$$
p(\vec{s'}|\vec{s},\vec{a}) = \begin{cases} \prod_{i=1}^{I} \Pr(q_i) & \text{if } \vec{s'} \text{ satifies equations (4) to (7),} \\ 0 & \text{otherwise.} \end{cases}
$$

$$
u'_m = u_{m+1} + \sum_{i=1}^{I} \sum_{\substack{k=\max\{(m+1)\\ -l_i+1,1\}}}^{\min\{(m+1),N\}} r_{i[(m+1)-k+1]} x_{ik} - y_{m+1} \qquad m < M, \tag{4}
$$

$$
v'_m = v_{m+1} + y_{m+1} \qquad m < M, \tag{5}
$$

$$
w'_i = w_i - \sum_{n=1}^{N} x_{in} + q_i \qquad \forall i, \tag{6}
$$

$$
u'_M = v'_M = 0 \tag{7}
$$

Equations (4), (5) and (7) define the new number of regular-hour and overtime appointment slots booked on day m as a function of the number of previous slots booked on day $m+1$ plus all new bookings that affected day $m+1$. Equation (6) determines the new number of treatments waiting to be booked as the number of treatment requests that have not yet been booked plus new demand. The term $\Pr(q_i)$ corresponds to the probability of having q_i new requests for treatments of type i.

4.5. Costs

The total cost associated with choosing booking action $\vec{a} = (\vec{x}, \vec{y})$ in state $\vec{s} = (\vec{u}, \vec{v}, \vec{w})$, $\vec{a} \in A_{\vec{s}}$ and $\vec{s} \in S$, comes from three sources: the penalties associated with the resulting patient wait times, the cost associated with the use of overtime, and the penalties associated with postponing some of the booking decisions. We represent the total cost as follows:

$$
c(\vec{s},\vec{a}) = \sum_{i=1}^{I} \sum_{n=1}^{N} c_{in} x_{in} + \sum_{m=1}^{M} h_m y_m + \sum_{i=1}^{I} g_i \left(w_i - \sum_{n=1}^{N} x_{in} \right) \tag{8}
$$

where

$$
c_{in} = \sum_{k=1}^{n} \lambda^{k-1} f_{ik} \quad \forall i, n \qquad h_m = \lambda^{m-1} h \quad \forall m
$$

In Equation (8), c_{in} represents the penalty (if any) for starting a treatment of type i on day n, h_m is the discounted overtime cost associated with an overtime booking on day m and g_i corresponds to the penalty for postponing to the next day the booking of a treatment of type i. The values of c_{in} are obtained by discounting the penalties f_{ik} associated with each additional day of wait before the start of a treatment. The relative orders of the wait time penalties are determined by expert opinion and investigated through sensitivity analysis. They are defined in relation to existing guidelines for acceptable waits and taking into consideration estimates of the importance of radiation therapy for different disease sites. The overtime cost is denoted by h and the discount factor by $\lambda < 1$.

To avoid keeping track of the wait times associated with postponed requests, our model assumes that the portion of the demand that is not booked today is handled the same as the new demand tomorrow. The postponement penalties, however, are set orders of magnitude higher than the wait time penalties and the overtime cost. Postponements thus constitute a last resort to relieve stress on the system and are intended to ensure problem feasibility.

4.6. Optimality equations

The value function in our formulation, denoted by $v(\vec{s})$, corresponds to the total expected discounted cost over the infinite horizon. Of course, we are not so much interested in determining the value function for a given policy as in finding the optimal stationary policy. To identify such a policy we need to solve the following optimality equations:

$$v(\vec{s}) = \min_{\vec{a} \in A_{\vec{s}}} \left\{ c(\vec{s}, \vec{a}) + \lambda \sum_{\vec{s'} \in S} p(\vec{s'}|\vec{s}, \vec{a}) v(\vec{s'}) \right\} \quad \forall \vec{s} \in S \tag{9}$$

The challenge is that even for very small instances the size of the state space and the size of the corresponding action sets make a direct solution to (9) impossible. The state variable $\vec{s} = (\vec{u}, \vec{v}, \vec{w})$ and the action variable $\vec{a} = (\vec{x}, \vec{y})$ have $(M + I + M)$ and $(I \times N + M)$ dimensions, respectively. Assuming that w_i can take D_i possible values, this means that we might have up to $(C_r+1)^M \times \prod_{i=1}^{I} D_i \times (C_o+1)^M$ different states and $\prod_{i=1}^{I} D_i^N \times (C_o+1)^M$ different (not necessarily feasible) actions. As an example, the formulation for the practical example presented in Section 6.2 involves more than 10^{300} possible states and 10^{500} potential actions.

5. Solution approach

In order to deal with an intractable number of states and actions, we first transform our MDP model into its equivalent linear programming form. The linear programming approach to discounted infinite-horizon MDPs, initially presented by d'Epenoux (1963), is based on writing the optimality equations in (9) as follows:

$$\max \sum_{\vec{s} \in S} \alpha(\vec{s}) v(\vec{s}) \tag{10}$$

subject to

$$c(\vec{s}, \vec{a}) + \lambda \sum_{\vec{s'} \in S} p(\vec{s'} | \vec{s}, \vec{a}) v(\vec{s'}) \geq v(\vec{s}) \qquad \forall \vec{s} \in S,\ \vec{a} \in A_{\vec{s}}$$

The value of $\alpha(\vec{s})$ represents the weight of state $\vec{s} \in S$ in the objective function. The solution to the equivalent linear programming model in (10) is the same as the solution to the optimality equations in (9) when $\vec{\alpha}$ is strictly positive (Derman, 1970; Kallenberg, 1983). We normalize $\vec{\alpha}$ to $\sum_{\vec{s} \in S} \alpha(\vec{s}) = 1$ and consider it as an exogenous probability distribution over the initial state of the system. The equivalent linear programming model, however, does not avoid the curse of dimensionality. The model in (10) has one variable for every state $\vec{s} \in S$ and one constraint for every feasible state-action pair $(\vec{s}, \vec{a})$, $\vec{s} \in S$ and $\vec{a} \in A_{\vec{s}}$, making its solution impossible.

Fortunately, a whole field of potential methods for dealing with the curse of dimensionality, called *Approximate Dynamic Programming* (ADP), has been developed in the last decade (Bertsekas and Tsitsiklis, 1996; Powell, 2007). A common approach within this field consists of using an approximation architecture to represent the value function in the MDP formulation or, equivalently, the variables in the equivalent linear programming model. The *approximate linear programming approach* to ADP was initially introduced by Schweitzer and Seidmann (1985) and has recently been reconsidered by de Farias and Roy (2003), Adelman and Mersereau (2008), Patrick et al. (2008), Desai et al. (2009) and Adelman and Klabjan (2011).

To solve the radiation therapy appointment scheduling problem we chose the following affine approximation to $v(\vec{u}, \vec{v}, \vec{w})$:

$$v(\vec{u}, \vec{v}, \vec{w}) = W_0 + \sum_{m=1}^{M} U_m u_m + \sum_{m=1}^{M} V_m v_m + \sum_{i=1}^{I} W_i w_i \tag{11}$$

$$\vec{U}, \vec{V}, \vec{W} \geq 0, W_0 \in \mathbb{R}$$

In this way the approximate mathematical programming model is still linear and the new variables are directly interpretable. Any other more sophisticated approximation such as a linear combination of a set of more general basis functions would make the subproblem in our column generation algorithm a non-linear integer programming problem. The values of $\{U_m\}_{m=1}^M$, $\{V_m\}_{m=1}^M$ and $\{W_i\}_{i=1}^I$ represent the marginal cost of having an additional regular-hour appointment slot occupied on day m, the marginal cost of having an additional overtime slot used on day m and the marginal cost of having one more treatment of type i waiting to be booked, respectively. If we substitute (11) into (10) and restrict $\vec{\alpha}$ to be a probability distribution, we obtain:

$$\max_{\vec{U},\vec{V},\vec{W},W_0} \left\{ W_0 + \sum_{m=1}^{M} \mathbb{E}_\alpha[u_m] U_m + \sum_{m=1}^{M} \mathbb{E}_\alpha[v_m] V_m + \sum_{i=1}^{I} \mathbb{E}_\alpha[w_i] W_i \right\} \tag{12}$$

subject to

$$\begin{aligned} (1-\lambda)W_0 + \sum_{m=1}^{M} \mu_m(\vec{s},\vec{a}) U_m + \sum_{m=1}^{M} \nu_m(\vec{s},\vec{a}) V_m + \sum_{i=1}^{I} \omega_i(\vec{s},\vec{a}) W_i \leq c(\vec{s},\vec{a}) \\ \forall \vec{s} \in S,\, \vec{a} \in A_{\vec{s}} \\ \vec{U}, \vec{V}, \vec{W} \geq 0, W_0 \in \mathbb{R} \end{aligned}$$

where

$$\begin{aligned} &\mathbb{E}_\alpha[u_m] = \sum_{\vec{s} \in S} \alpha(\vec{s}) u_m(\vec{s}) \quad \forall m && \mu_m(\vec{s},\vec{a}) = u_m(\vec{s}) - \lambda u'_m(\vec{s},\vec{a}) && \forall m \\ &\mathbb{E}_\alpha[v_m] = \sum_{\vec{s} \in S} \alpha(\vec{s}) v_m(\vec{s}) \quad \forall m && \nu_m(\vec{s},\vec{a}) = v_m(\vec{s}) - \lambda v'_m(\vec{s},\vec{a}) && \forall m \\ &\mathbb{E}_\alpha[w_i] = \sum_{\vec{s} \in S} \alpha(\vec{s}) w_i(\vec{s}) \quad \forall i && \omega_i(\vec{s},\vec{a}) = w_i(\vec{s}) - \lambda \left(w_i(\vec{s}) - \sum_{n=1}^{N} x_{in} + m_i \right) && \forall i \end{aligned}$$

The *approximate* equivalent linear programming model in (12) has a tractable number of variables, $(2M+I+1)$, but still an intractable number of constraints. For this reason, we solve its dual (13) using column generation.

$$\min_{\vec{X}} \sum_{\vec{s}\in S,\ \vec{a}\in A_{\vec{s}}} c(\vec{s},\vec{a})X(\vec{s},\vec{a}) \tag{13}$$

$$\begin{array}{rcll}
(1-\lambda)\displaystyle\sum_{\vec{s}\in S,\ \vec{a}\in A_{\vec{s}}} X(\vec{s},\vec{a}) & = & 1 & \\
\displaystyle\sum_{\vec{s}\in S,\ \vec{a}\in A_{\vec{s}}} \mu_m(\vec{s},\vec{a})X(\vec{s},\vec{a}) & \geq & \mathbb{E}_\alpha[u_m] & \forall m \\
\displaystyle\sum_{\vec{s}\in S,\ \vec{a}\in A_{\vec{s}}} \nu_m(\vec{s},\vec{a})X(\vec{s},\vec{a}) & \geq & \mathbb{E}_\alpha[v_m] & \forall m \\
\displaystyle\sum_{\vec{s}\in S,\ \vec{a}\in A_{\vec{s}}} \omega_i(\vec{s},\vec{a})X(\vec{s},\vec{a}) & \geq & \mathbb{E}_\alpha[w_i] & \forall i \\
\vec{X} & \geq & 0 &
\end{array}$$

Column generation finds the optimal solution to (13), the *master problem*, starting with a small set of feasible state-action pairs and iteratively adding the state-action pair associated with the most violated primal constraint. Unfortunately, there is no easy way to identify an initial set of feasible state-action pairs to (13). For this reason, our methodological approach involves two additional mathematical programming models. The first model finds an initial set of state-action pairs by combining the action-set constraints and the definition of $\mu_m(\vec{s},\vec{a})$ in (12). It is solved for each treatment type and focuses on dual feasibility rather than optimality. The second model corresponds to a Phase I method. It starts from the state-action pairs provided by the first model and incorporates new state-action pairs into the formulation until a feasible solution to (13) is found (see Appendix A for more details).

The model used to identify the state-action pair associated with the most violated primal constraint, the *pricing problem*, is itself an optimization problem. Given the dual values associated with the current solution to (13), $\{U_m\}_{m=1}^M$, $\{V_m\}_{m=1}^M$ and $\{W_i\}_{i=1}^I$, the next state-action pair to enter the basis is given by (14).

$$\arg\min_{\vec{s}\in S,\vec{a}\in A_{\vec{s}}} \left\{ c(\vec{s},\vec{a}) - (1-\lambda)W_0 - \sum_{m=1}^{M} \mu_m(\vec{s},\vec{a})U_m - \sum_{m=1}^{M} \nu_m(\vec{s},\vec{a})V_m - \sum_{i=1}^{I} \omega_i(\vec{s},\vec{a})W_i \right\} \tag{14}$$

$$\sum_{i=1}^{I} \sum_{k=\max\{m-l_i+1,1\}}^{\min\{m,N\}} r_{i(m-k+1)}x_{ik} \geq y_m \qquad \forall m \tag{15}$$

$$u_M = 0 \qquad v_M = 0 \tag{16}$$

Three additional constraints are added to guarantee that only valid state-action pairs are generated. Constraint (15) limits the number of overtime bookings on day m to be less than or equal to the total number of new bookings for that day. Constraints in (16) ensure that the M-th day in the planning horizon of any valid state-action pair has no appointments in it.

The column generation algorithm iterates until no primal constraint is violated or until we are close enough to quit (stopping criterion of -0.0001), giving us $\{U_m^*\}_{m=1}^M$, $\{V_m^*\}_{m=1}^M$ and $\{W_i^*\}_{i=1}^I$. These values are then used to identify the *approximate* optimal policy. In practice, rather than computing and storing the approximate optimal actions for each state, a resource-intensive task, we only compute them as needed. To this end, we solve:

$$d^*(\vec{s}) \in \underset{(\vec{x},\vec{y}) \in A_{\vec{s}}}{\arg\min} \left\{ \sum_{i=1}^{I} \sum_{n=1}^{N} C_{in} x_{in} + \sum_{m=1}^{M} H_m y_m \right\} \tag{17}$$

where

$$C_{in} = c_{in} + \lambda \sum_{k=(n-1)}^{(n-1)+l_i-1} r_{i(k+1-n+1)} U_k^* - (g_i + \lambda W_i^*) \qquad \forall i, n$$

$$H_m = \begin{cases} h & m = 1 \\ h_m + \lambda V_{m-1}^* - \lambda U_{m-1}^* & m > 1 \end{cases}$$

The integer programming model in (17) is obtained by inserting the approximate value function defined by $\{U_m^*\}_{m=1}^M$, $\{V_m^*\}_{m=1}^M$ and $\{W_i^*\}_{i=1}^I$ into the right hand side of the optimality equations in (9) and ignoring the constant terms. The coefficients accompanying the booking actions in (17) have direct interpretations and their values provide a good description of the approximate optimal policy. C_{in} balances the penalty associated with a wait time of n days and the cost due to the loss of available treatment capacity in the future, $c_{in} + \lambda \sum_{k=(n-1)}^{(n-1)+l_i-1} r_{i(k+1-n+1)} U_k^*$, against the benefit due to the fact that the booking decision is not postponed and the treatment request does not re-appear in tomorrow's demand, $g_i + \lambda W_i^*$. H_m balances the cost associated with a one-unit increase in the system total capacity on day m and the cost due to the loss of available overtime capacity on day $m-1$ tomorrow, $h_m + \lambda V_{m-1}^*$, against the benefit from freeing regular-hour capacity on day $m-1$ tomorrow, λU_{m-1}^*. It is important to note that when the expected daily demand for appointment slots does not exceed the regular-hour capacity (i.e.

when capacity is not a binding constraint), $\{U_m^*\}_{m=1}^M$, $\{V_m^*\}_{m=1}^M$ and $\{W_i^*\}_{i=1}^I$ are equal to zero. In this case, the coefficients of the booking actions do not need to be adjusted to take into account the impact of today's decisions on the future and the approximate optimal policy becomes a myopic policy (C_{in} and H_m are simply defined by $c_{in} - g_i$ and h_m, respectively). We refer the reader to Appendix B for a partial description of the analytical solution to (12) for the sample problems presented in Section 6.

One of the difficulties with column generation is the large number of iterations needed to find the optimal solution to the master problem (often referred to as the *tailing-off effect* of column generation, Lübbecke and Desrosiers 2005). In order to reduce execution times, we experimented with alternative ways of finding an initial set of feasible state-action pairs, different stopping criteria and several methods for solving the master problem in each iteration. We found the best execution times are achieved when the master problem is solved using barrier methods starting from the immediate previous solution. We also implemented the in-out approach proposed by Ben-Ameur and Neto (2007) to further reduce computation time.

The implementation of the column generation algorithm and the integer programming model used to identify the approximate optimal policy was performed in GAMS 23.5 with CPLEX 12.2 as the solver.

6. Results

This section provides some insights regarding the properties of the approximate optimal policy and illustrates the potential benefits that can be obtained from its use. We first consider a small example of the problem and analyze the impact of different treatment specifications on the values of C_{in} and H_m. Then, we evaluate the approximate optimal policy by simulating its performance for a more practical example based on British Columbia Cancer Agency (BCCA) operations.

6.1. Policy insights

In this section, we analyze six scenarios to illustrate some relevant properties of the coefficients defining the approximate optimal policy. Each scenario assumes the values of c_{in} are determined by discounting a constant wait time penalty f_i for each day of wait beyond a recommended wait time target (wait time penalties are incurred for each appointment slot). The regular-hour capacity is set equal to the average demand, request arrivals are assumed

Poisson and postponements are not allowed. In theory, assuming that postponements are not allowed can result in infeasibility. However, in practice, this is not an issue as there are sufficient regular-hour and overtime appointment slots to guarantee feasibility. Demand distributions are truncated at three times their mean values.

(a) *Base case* – we consider a hypothetical radiation therapy center with a regular-hour capacity of 50 appointment slots per day and an overtime capacity of 6 appointment slots per day. The facility initially delivers only one type of treatment consisting of five consecutive one-appointment-slot sessions. The average demand for this type of treatment is 10 requests per day. The wait time target is 10 days and the wait time penalty is \$50 per day per appointment slot. The overtime cost is \$100, the length of the booking horizon is 25 days and the discount factor is 0.99. A discount factor of 0.99 reflects the medium-term planning horizon that is often applicable in health-care settings. In this way, costs are relatively similar over the short-term and less valued far in the future.

(b) *Different number of sessions* – we consider five treatment types consisting of three, four, five, six and seven one-appointment-slot sessions, respectively. The mean demand for each type is set at 2 requests per day. All other parameters stay the same.

(c) *Different session durations* – we have five treatment types, all with only one treatment session, but the duration of the session depends on the type, requiring between three and seven appointment slots. All other parameters remain the same.

(d) *Different wait time penalties* – we study five treatment types, all of them consisting of five one-appointment-slot sessions. What varies is the wait time penalty, taking values of \$10, \$25, \$50, \$75 and \$100 per day per appointment slot, respectively. All other parameters stay the same.

(e) *Different wait time targets* – we still consider five treatments types consisting of five one-appointment-slot sessions each. The wait time penalties are set back at \$50 per day per appointment slot, but the wait time targets are 5, 8, 10, 12 and 15 days, respectively.

(f) *Different demand rates* – we have five treatment types, all of them consisting of two one-appointment-slot sessions. The demand rate varies depending on the type, taking values from three to seven requests per day, respectively. All other parameters remain the same as the base case.

Note that even for the small instances presented in this section the size of the state space and the size of the corresponding action sets make a direct solution to (9) impossible. For example, the formulation for the base case scenario described above involves more than 10^{75} possible states.

Figure 2 provides a graphical representation of C_{in}, the adjusted cost associated with starting a treatment of type i on day n, for each scenario. We focus our analysis on this coefficient since it is the most relevant when identifying the approximate optimal booking actions. The values of H_m are, with the exception of the value for day 1, very close to zero, if not zero. This, together with the fact that the values of C_{in} on day 1 are negative and lower than the overtime cost, allow us to conclude that the approximate optimal policy will allocate available treatment capacity according to the values of C_{in} and use overtime whenever demand is present and regular-hour capacity is not available.

Figure 2 suggests the following properties for the approximate optimal policy: (1) to book as much demand as possible on workday 1 and then within the wait time targets (see Fig. 2a); (2) the larger the number of sessions, the earlier the booking (see Fig. 2b); (3) the shorter the duration of each session, the later the booking (see Fig. 2c); and (4) the more urgent the treatment or the shorter the wait time target, the earlier the booking (see Fig. 2d and 2e). The first property follows from the negative values of C_{in}, with the values associated with workday 1 being the most negative. The other three properties are due to the fact that longer, more intense and more urgent treatments have lower negative C_{in} values. Additionally, Figure 2f suggests that the approximate optimal policy does not depend on the demand rate associated with each treatment type. The drops in the values of C_{in} in Figure 2c are due to an abrupt decrease in the value of U_m^* in the last days of the booking horizon. U_m^* becomes less relevant late in the booking horizon since most treatments will be booked within the wait time target, leaving the system empty further into the future.

The values of H_m, the adjusted cost associated with a one-unit increase in the system total capacity on day m, are almost identical across all scenarios. The value of H_m on workday 1 is equal to the overtime cost. Then, as a consequence of the opportunity cost due to the loss of available overtime capacity and the benefit from freeing regular-hour capacity, H_m quickly decreases until it reaches zero. Late in the planning horizon, the opportunity costs vanish and H_m becomes equal to the discounted overtime cost h_m.

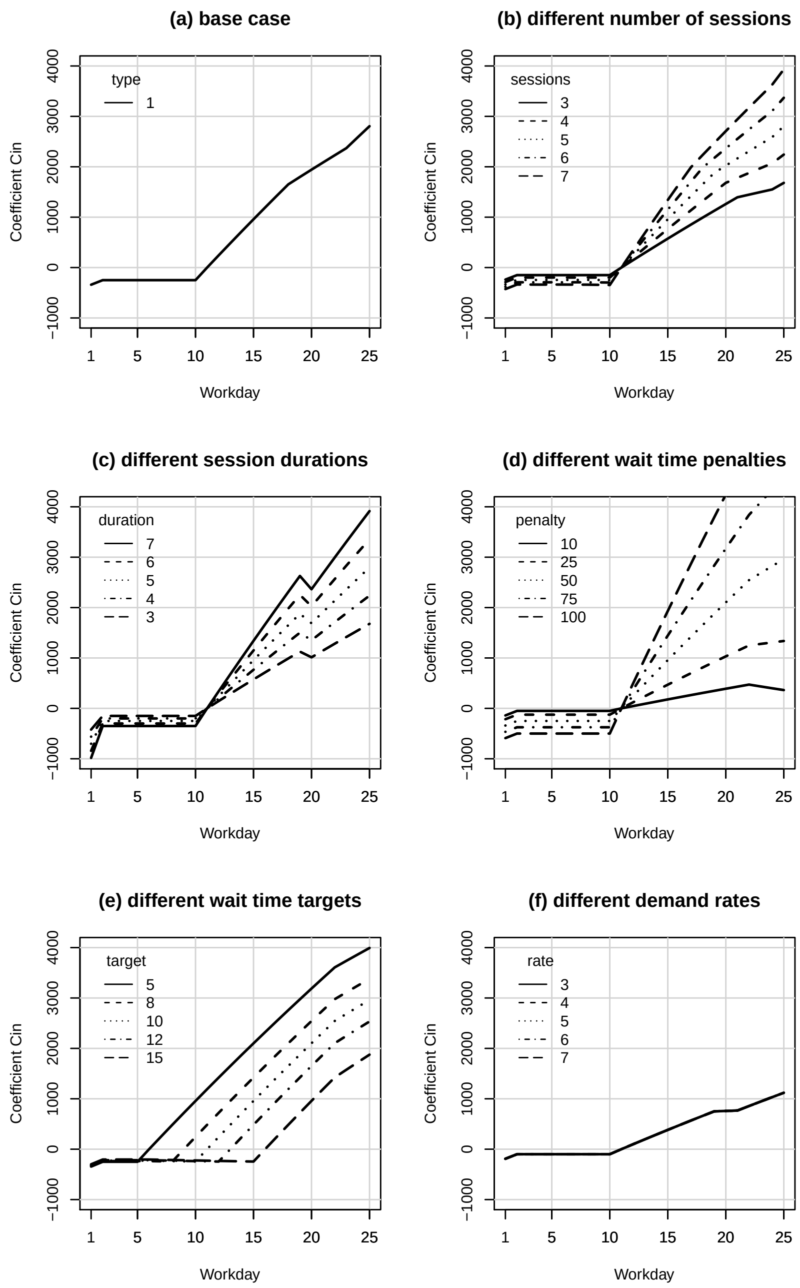

Figure 2: Graphical display of C_{in}, the adjusted cost associated with starting a treatment of type i on day n, for the six scenarios defined in Section 6.1. Scenarios are labeled as: (a) base case, (b) different number of sessions, (c) different session durations, (d) different wait time penalties, (e) different wait time targets, and (f) different demand rates.

6.2. A practical example

This section describes the potential benefits from using the proposed solution approach by evaluating its performance for a practical-sized example based on an approximation to BCCA operations using historical data.

The BCCA is a publicly funded organization with a mandate to provide a cancer control program for the people of British Columbia, Canada. This program includes prevention, early detection, treatment and various forms of care. Services provided by the BCCA are delivered through five regional cancer centers. The instance of the problem analyzed in this section is based on all breast, head and neck, lung and prostate treatments registered and completed at the Vancouver Cancer Center (VCC) between April 1, 2009 and March 31, 2010 (2,061 cases in total, equivalent to 60% of the total number of treatments delivered within these dates).

An analysis of the data identified 461 different capacity requirements out of 2,061 treatments delivered. For this reason, we chose to classify treatments based on urgency, cancer site and treatment intent and then to represent each class using the most frequent capacity requirements within that group. Table 2 describes the 18 treatment types obtained through this exercise. They account for 62% of the treatments in our data. Treatments of type 1, for example, represent urgent lung, prostate and breast palliative cases and consist of an initial session of two appointment slots plus four additional sessions of one appointment slot each. Appointments slots are 12-minutes long. Note that, in practice, the values defining the optimal value function approximation can be used to estimate the impact of today's decisions on the future performance of the system not only for the treatment types considered in their computation but also for any other possible treatment type since they represent the marginal cost of resource consumption irrespective of the treatment type.

Between April 2009 and March 2010, the booking agents at the VCC faced an average demand of 8.25 treatment requests per day, for the above cancer sites. A Poisson distribution with rate 8.25 provided the best fit for the total number of requests per day (p-value = 0.724). We assumed the demand rate for each treatment type is proportional to the number of cases of each type and set a maximum number of arrivals, obtained from historical data, in order to maintain a finite state space. The demand process is described by the arrival rates given in Table 2.

Regular-hour capacity is set at 120 appointment slots, which is equivalent to three identical treatment units operating eight hours a day. This number

Table 2: Characteristics of the treatment types used to evaluate the performance of the proposed methodology. Types are grouped according to common urgency (see Table 3).

type	capacity requirements	sessions/ slots	arrival rate [reqs./day]	most frequent case (site and intent)
1	1×2 + 4×1	5/ 6	0.19	Lung/Prostate/Breast
2	1×2	1/ 2	0.11	Palliative
3	1×2 + 3×1	4/ 5	0.11	(Urgent)
4	1×2 + 15×1	16/17	1.43	Breast Adjuvant
5	1×2 + 15×1 + 1×2 + 3×1	20/22	0.59	
6	1×3 + 15×2	16/33	0.45	
7	1×2	1/ 2	1.42	Lung/Breast/Prostate
8	1×2 + 4×1	5/ 6	1.36	Palliative
9	1×2 + 9×1	10/11	0.57	(Non-urgent)
10	1×2 + 3×1	4/ 5	0.38	
11	1×2 + 14×1	15/16	0.18	
12	1×1	1/ 1	0.18	
13	1×2 + 19×1	20/21	0.29	Head and Neck Radical
14	1×3 + 34×2	35/71	0.21	
15	1×2 + 32×1	33/34	0.30	Prostate Radical
16	1×2 + 36×1	37/38	0.29	
17	1×2 + 21×1 + 1×2 + 14×1	37/39	0.15	
18	1×2 + 32×1	33/34	0.04	Prostate Adjuvant

is almost six appointment slots lower than the average daily demand. The overtime capacity is 15 appointment slots or one extra hour per treatment unit. The overtime cost is set at 100 per appointment slot, the discount factor remains at 0.99 and no postponements are allowed. The booking horizon and the planning horizon are set at 100 and 136 workdays, respectively. In practice, the approximate optimal policy has proved to be independent of the length of the booking horizon provided that it exceeds the wait time target of the least urgent treatment type. The wait time penalties are defined in Table 3. They were specified by co-author Dr. Tyldesley, an experienced radiation oncologist at BCCA, on the basis of existing Joint Collegiate Council for Oncology guidelines for acceptable waits, which provide good practice and maximal acceptable wait times for urgent, radical, palliative and adjuvant cases (Ash et al., 2004), and estimates of the importance of radiation therapy in relation to disease sites.

Figure 3 shows the values of C_{in} for treatment types 1, 4, 7, 14, 17 and 18 (one type per wait time penalty scheme) for the first 25 workdays of the booking horizon. These six treatment types were chosen because they

Table 3: Wait time penalties associated with each treatment type.

types	daily penalty within workday interval						
	[0,1]	(1,5]	(5,10]	(10,20]	(20,30]	(30,40]	(40,100]
1– 3	0	100	150	150	150	150	150
4– 6	0	0	0	50	100	100	150
7–12	0	0	65	100	100	100	150
13–14	0	0	80	150	150	150	150
15–17	0	0	0	40	80	100	150
18	0	0	0	50	90	100	150

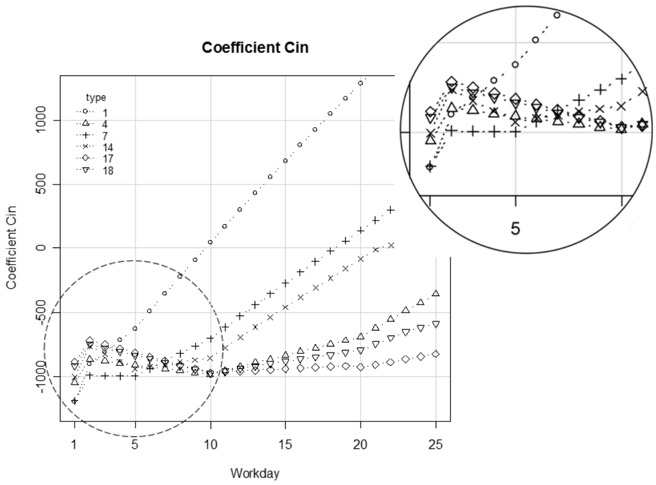

Figure 3: Partial display of C_{in}, the adjusted cost associated with starting a treatment of type i on workday n, for treatment types 1, 4, 7, 14, 17 and 18.

provide a good description of the overall booking preferences associated with each wait time penalty scheme. As in the previous section, we focus our attention on coefficient C_{in} since it is the most relevant when identifying the approximate optimal booking actions. The value of H_m in this case is equal to zero from workday 2 to far into the future. This, together with the fact that the values of C_{in} on workday 1 are negative and orders of magnitude lower than the overtime cost, demonstrates that the approximate optimal policy for the evaluation instance will use overtime whenever demand is present and regular-hour capacity is not available.

Figure 3 demonstrates that the values of C_{in} are strictly increasing for treatment type 1 and non-monotonic for all other types. This is a direct consequence of the differences in capacity requirements and the wait time

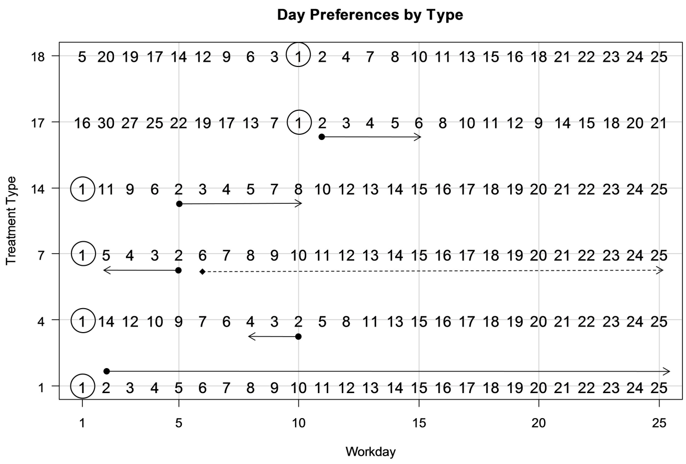

Figure 4: Partial display of the booking day preferences for treatment types 1, 4, 7, 14, 17 and 18. For each treatment type, a circle indicates the most preferable booking day and the arrows a recommended booking order.

penalties between types. Taking each type independently, the values of C_{in} identify the booking day preferences associated with each treatment class (see Figure 4). Treatments should be initiated according to these preferences and the available treatment capacity. Figure 4 shows that treatments of type 1 should be initiated as soon as possible. Treatments of type 7, however, should be scheduled starting with workday 1, then workday 5, 4, 3 and 2. If a treatment of type 7 has not been scheduled within the first five workdays, then it should be booked on the first workday with available capacity after workday 5. The preferences for treatment types 17 and 18 make evident the lack of capacity in the system. It is necessary to defer these two types of treatments in order to deliver more urgent types earlier.

In order to demonstrate the impact of today's decisions on the future performance of the system, which is one of the main contributions of our model, Figure 5 shows the difference between the values of C_{in} associated with the approximate optimal policy and the myopic policy for the six treatment types shown in Figure 3. The differences reflect the importance of the opportunity cost that is not considered by the myopic policy. Under the approximate optimal policy less urgent treatment types such as 17 and 18 have higher C_{in} values early on in the booking horizon and therefore are less likely to be booked early. This observation supports the conclusions obtained from

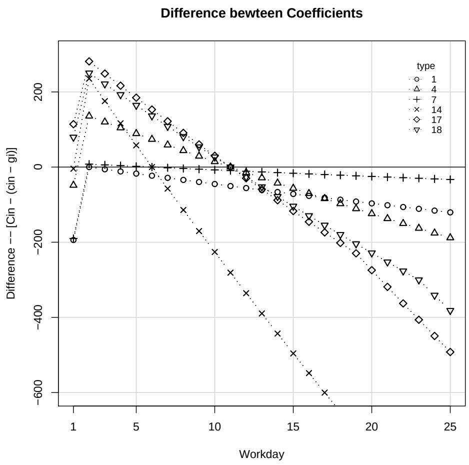

Figure 5: Difference between the capacity allocation coefficients defining the approximate optimal policy and the myopic policy for treatment types 1, 4, 7, 14, 17 and 18.

Figure 4.

The benefits from the proposed method are evaluated by simulating the performance of the scheduling system under the approximate optimal policy. Results are compared to the myopic policy. The simulation length was set at 1,500 days and statistics were collected for each of 10 runs after a warm-up period of 750 days.

Table 4 demonstrates that the approximate optimal policy outperforms the myopic policy with respect to the total number of cases initiated within 1, 5 and 10 workdays. The same improvement is observed for most treatment types individually. The average percentages of treatments initiated within 1, 5 and 10 workdays increase from 5% to 26%, 29% to 53% and 73% to 96%, respectively. This is achieved at the expense of a negligible but statistically significant increase in the average overtime utilization of 3 minutes a day. The approximate optimal policy achieves better service levels by scheduling less important treatments in terms of wait time penalties and capacity requirements later in the planning horizon, allowing other more important treatments to be scheduled earlier. Treatments of types 15 to 18 are booked starting on workday 10 allowing most of the treatments of types 7 to 14 to be scheduled within the first five workdays. The main drawback with regard to the proposed policy is the increase in the wait times for treatments of types 1 to 3. However, the fact that urgent treatments wait slightly longer demon-

Table 4: Simulation results. The bold font indicates the policy that provides the highest service level for each treatment type and wait time target. Only the results for which a significance test shows the mean service level has improved are highlighted ($\alpha = 0.05$).

	% of the cases initiated within				
type	1 workday	5 workdays	10 workdays	15 workdays	20 workdays
			myopic policy		
1	70 ± 10	**94 ± 4**	**100 ± 0**	100 ± 0	100 ± 0
2	**82 ± 8**	**95 ± 5**	**100 ± 0**	100 ± 0	100 ± 0
3	72 ± 12	**95 ± 5**	**100 ± 0**	100 ± 0	100 ± 0
4	1 ± 2	4 ± 3	57 ± 9	98 ± 2	100 ± 0
5	1 ± 1	4 ± 3	55 ± 10	98 ± 2	100 ± 0
6	1 ± 1	4 ± 3	53 ± 10	97 ± 3	100 ± 0
7	2 ± 2	75 ± 4	**100 ± 0**	100 ± 0	100 ± 0
8	2 ± 2	25 ± 9	76 ± 5	100 ± 1	100 ± 0
9	2 ± 2	16 ± 10	62 ± 9	99 ± 2	100 ± 0
10	2 ± 2	37 ± 8	89 ± 2	100 ± 0	100 ± 0
11	3 ± 2	15 ± 10	61 ± 10	99 ± 1	100 ± 0
12	4 ± 2	92 ± 2	100 ± 0	100 ± 0	100 ± 0
13	1 ± 1	19 ± 12	90 ± 3	100 ± 0	100 ± 0
14	1 ± 1	16 ± 10	61 ± 9	99 ± 2	100 ± 0
15	1 ± 1	**2 ± 2**	54 ± 9	97 ± 3	100 ± 0
16	0 ± 0	**2 ± 1**	52 ± 10	97 ± 3	100 ± 0
17	1 ± 1	2 ± 2	52 ± 8	97 ± 3	100 ± 0
18	1 ± 1	**3 ± 3**	55 ± 11	98 ± 2	100 ± 0
total	5 ± 2	29 ± 4	73 ± 6	99 ± 1	100 ± 0
			approximate optimal policy		
1	66 ± 7	81 ± 8	97 ± 2	100 ± 0	100 ± 0
2	75 ± 9	84 ± 8	97 ± 2	100 ± 0	100 ± 0
3	66 ± 11	80 ± 10	97 ± 3	100 ± 0	100 ± 0
4	**17 ± 6**	**17 ± 6**	**95 ± 4**	100 ± 0	100 ± 0
5	**11 ± 5**	**11 ± 5**	**94 ± 5**	100 ± 0	100 ± 0
6	**12 ± 5**	**12 ± 5**	**94 ± 5**	100 ± 0	100 ± 0
7	43 ± 9	82 ± 8	97 ± 3	100 ± 0	100 ± 0
8	**31 ± 9**	**79 ± 9**	**96 ± 3**	100 ± 0	100 ± 0
9	**27 ± 8**	**78 ± 9**	**96 ± 3**	100 ± 0	100 ± 0
10	**38 ± 8**	**81 ± 8**	**96 ± 3**	100 ± 0	100 ± 0
11	**25 ± 9**	**77 ± 9**	**96 ± 3**	100 ± 0	100 ± 0
12	**36 ± 7**	**97 ± 2**	100 ± 0	100 ± 0	100 ± 0
13	**23 ± 9**	**80 ± 8**	**97 ± 3**	100 ± 0	100 ± 0
14	**11 ± 4**	**77 ± 9**	**95 ± 3**	100 ± 0	100 ± 0
15	0 ± 0	0 ± 0	**94 ± 5**	100 ± 0	100 ± 0
16	0 ± 0	0 ± 0	**94 ± 4**	100 ± 0	100 ± 0
17	0 ± 0	0 ± 0	**94 ± 5**	100 ± 0	100 ± 0
18	0 ± 0	0 ± 0	**93 ± 6**	100 ± 0	100 ± 0
total	**26 ± 7**	**53 ± 6**	**96 ± 3**	100 ± 0	100 ± 0

strates a willingness to trade-off a small increase in wait time for urgent treatments for a larger gain for less urgent treatments.

The results show no statistically significant difference between the two policies in terms of regular-hour capacity utilization. Both policies use on average about 99.5% of the available regular-hour capacity. However, the approximate optimal policy outperforms the myopic policy with respect to the average discounted cost (p-value = 0.000). The average discounted cost associated with the approximate optimal policy is $121,973.57 and the average discounted cost associated with the myopic policy is $185,843.06.

The total time required to identify the approximate optimal policy was 1 hour and 45 minutes and the total time needed to simulate each policy was 3 hours and 20 minutes. The computer used to run our algorithm was a 3.00GHz Quad Core PC with 16GB of RAM. The simulation model was also implemented using GAMS 23.5 and CPLEX 12.2.

7. Conclusions

This paper describes the use of approximate dynamic programming as a means of solving a radiation therapy appointment scheduling problem, a problem computationally intractable using standard methods. Starting from the linear programming approach to discounted infinite-horizon MDPs and employing an affine structure to represent the value function in the equivalent linear programming model, our methodological approach provides a systematic way of identifying effective booking guidelines for radiation therapy. These guidelines could be used in practice to significantly reduce wait times for radiation therapy and thus to decrease the impact of delays on cancer patients' health.

The results presented in Section 6.2 show how the approximate optimal policy outperforms the myopic policy which is an approximation of the current practice. The percentage of treatments initiated within the clinical benchmark (10 workdays) increases, on average, from 73% to 96% under the proposed policy. This increase brings with it a negligible but statistically significant increase in the average overtime utilization of 3 minutes a day.

Possible extensions to this paper involve elements of the original problem that were excluded from our formulation. The inclusion of multiple treatment units and unit compatibility restrictions, in addition to wait lists, would significantly increase the complexity of the problem. We are also interested in proving the analytical solution for the coefficients defining the optimal value function approximation (see Appendix B), computing better bounds on the optimality gap and applying other approximate dynamic programming techniques to the radiation therapy appointment scheduling problem. Our models could also be extended to consider cancellations and no-shows.

Acknowledgments

This work was partially supported by NSERC Discovery Grant 5527-07 and by a research grant from the Canadian Institutes of Health Research (CIHR) in support of *The CIHR Team in Operations Research for Improved Cancer Care.* We thank all the administrators, physicians, therapists and clerks at BCCA's Vancouver Centre for their contribution in this project. In particular, we are grateful of the valuable insights provided by Mike Darud, Pablo Santibáñez and John French. We also wish to thank Maurice Queyranne for many helpful comments through the course of this project.

Appendix A. Finding an initial set of feasible state-action pairs

We first find an initial set of state-action pairs, one pair per treatment type, by combining the action-set constraints and the definition of $\mu_m(\vec{s},\vec{a})$ in (12).

$$(\vec{s},\vec{a})_i \in \underset{\vec{s}\in S,\ \vec{a}\in A_{\vec{s}}}{\arg\max} \{\eta_i : \eta_i \leq \mu_m(\vec{s},\vec{a})\ \forall m,\ v_m = y_m = 0\ \forall m,\ w_i = 3m_i,\ w_j = 0\ \forall j \neq i\} \quad \forall i$$

Then, we iteratively add new state-action pairs to (A.1) based on its current dual solution and the pricing problem in (14) until the value s is equal to zero.

$$\begin{aligned}
&\min s && && \quad\quad (A.1)\\
&\text{subject to} \\
&(1-\lambda)\sum_{\vec{s}\in S,\ \vec{a}\in A_{\vec{s}}} X(\vec{s},\vec{a}) &=&\ 1 \\
&\sum_{\vec{s}\in S,\ \vec{a}\in A_{\vec{s}}} \mu_m(\vec{s},\vec{a})X(\vec{s},\vec{a}) &\geq&\ \mathbb{E}_\alpha[u_m] - s && \forall m \\
&\sum_{\vec{s}\in S,\ \vec{a}\in A_{\vec{s}}} \nu_m(\vec{s},\vec{a})X(\vec{s},\vec{a}) &\geq&\ \mathbb{E}_\alpha[v_m] - s && \forall m \\
&\sum_{\vec{s}\in S,\ \vec{a}\in A_{\vec{s}}} \omega_i(\vec{s},\vec{a})X(\vec{s},\vec{a}) &\geq&\ \mathbb{E}_\alpha[w_i] - s && \forall i \\
&\vec{X}, s &\geq&\ 0
\end{aligned}$$

Appendix B. Optimal value function approximation

$$U_m^* = \begin{cases} U_{m+l_1}^* & m = 1, \ldots, (T_1 - 1) \\ \lambda^{(m-1)} h & m = T_1, \ldots, (M-1) \\ 0 & m = M \end{cases}$$

$$V_m^* = 0 \qquad \forall m$$

$$W_i^* = \sum_{k=T_i}^{T_i+l_i-1} r_{i(k+1-T_i)} U_m^* \qquad \forall i$$

l_1 and T_1 represent the length and the wait time target associated with the most urgent treatment type, respectively. The optimal value of W_0 is not listed since it is not part of the definition of the coefficients characterizing the approximate optimal policy.

References

Adelman, D., Klabjan, D., 2011. Computing near-optimal policies in generalized joint replenishment. INFORMS J. Comput. In press.

Adelman, D., Mersereau, A., 2008. Relaxations of weakly coupled stochastic dynamic programs. Oper. Res. 56 (3), 712–727.

Ash, D., Barrett, A., Hinks, A., Squire, C., 2004. Re-audit of radiotherapy waiting times 2003. Clin. Oncol. 16 (6), 387–394.

Begen, M. A., Queyranne, M., 2011. Appointment scheduling with discrete random durations. Math. Oper. Res. 36 (2), 240–257.

Ben-Ameur, W., Neto, J., 2007. Acceleration of cutting-plane and column generation algorithms: Applications to network design. Networks 49 (1), 3–17.

Bertsekas, D., Tsitsiklis, J., 1996. Neuro-Dynamic Programming. Athena Scientific, Belmont, Mass.

Cardoen, B., Demeulemeester, E., Belien, J., 2010. Operating room planning and scheduling: A literature review. Eur. J. Oper. Res. 201 (3), 921–932.

Cayirli, T., Veral, E., 2003. Outpatient scheduling in health care: A review of literature. Production Oper. Management 12 (4), 519–549.

Chen, Z., King, W., Pearcey, R., Kerba, M., Mackillop, W., 2008. The relationship between waiting time for radiotherapy and clinical outcomes: A systematic review of the literature. Radiother. Oncol. 87 (1), 3–16.

Coles, C., Burgess, L., L.Tan, 2003. An audit of delays before and during radical radiotherapy for cervical cancer-effect on tumour cure probability. Clin. Oncol. 15 (2), 47–54.

Conforti, D., Guerriero, F., Guido, R., 2010. Non-block scheduling with priority for radiotherapy treatments. Eur. J. Oper. Res. 201 (1), 289–296.

de Farias, D., Roy, B. V., 2003. The linear programming approach to approximate dynamic programming. Oper. Res. 51 (6), 850–865.

d'Epenoux, F., 1963. A probabilistic production and inventory problem. Management Sci. 10 (1), 98–108.

Derman, C., 1970. Finite State Markovian Decision Processes. Academic Press, Inc., Orlando, FL, USA.

Desai, V., Farias, V., Moallemi, C., 2009. A smoothed approximate linear program. Advances in Neural Information Processing Systems 22, 459–467.

Erdelyi, A., Topaloglu, H., 2009. Computing protection level policies for dynamic capacity allocation problems by using stochastic approximation methods. IIE Trans. 41, 498–510.

Fortin, A., Bairati, I., Albert, M., Moore, L., Allard, J., Couture, C., 2002. Effect of treatment delay on outcome of patients with early-stage head-and-neck carcinoma receiving radical radiotherapy. Int. J. Radiat. Oncol. Biol. Phys. 52 (4), 929–936.

Gerchak, Y., Gupta, D., Henig, M., 1996. Reservation planning for elective surgery under uncertain demand for emergency surgery. Management Sci. 42 (3), 321–334.

Gupta, D., Denton, B., 2008. Appointment scheduling in health care: Challenges and opportunities. IIE Trans. 40 (9), 800–819.

Kallenberg, L., 1983. Linear Programming and Finite Markovian Control Problems. Mathematisch Centrum, Amsterdam.

Kim, M., Ghate, A., Phillips, M. H., 2012. A stochastic control formalism for dynamic biologically conformal radiation therapy. Eur. J. Oper. Res. 219 (3), 541–556.

Lamiri, M., Xie, X., Dolgui, A., Grimaud, F., 2008. A stochastic model for operating room planning with elective and emergency demand for surgery. Eur. J. Oper. Res. 185 (3), 1026–1037.

Lim, G. J., Cao, W., 2012. A two-phase method for selecting IMRT treatment beam angles: Branch-and-prune and local neighborhood search. Eur. J. Oper. Res. 217 (3), 609–618.

Lübbecke, M., Desrosiers, J., 2005. Selected topics in column generation. Oper. Res. 53 (6), 1007–1023.

Mackillop, W., 2007. Killing time: The consequences of delays in radiotherapy. Radiother. Oncol. 84 (1), 1–4.

Magerlein, J., Martin, J., 1978. Surgical demand scheduling: A review. Health Services Res. 13 (4), 418–433.

Mišić, V., Aleman, D., Sharpe, M., 2010. Neighborhood search approaches to non-coplanar beam orientation optimization for total marrow irradiation using IMRT. Eur. J. Oper. Res. 205 (3), 522–527.

Mondschein, S., Weintraub, G., 2003. Appointment policies in service operations: A critical analysis of the economic framework. Production Oper. Management 12 (2), 266–286.

Norris, S., 2009. The wait time issue and the patient wait times guarantee. Tech. rep., Canada: Library of Parliament.

O'Rourke, N., Edwards, R., 2000. Lung cancer treatment waiting times and tumour growth. Clin. Oncol. 12 (3), 141–144.

Patrick, J., Puterman, M., Queyranne, M., 2008. Dynamic multi-priority patient scheduling for a diagnostic resource. Oper. Res. 56 (6), 1507–1525.

Powell, W., 2007. Approximate Dynamic Programming: Solving the Curses of Dimensionality. Wiley-Interscience, Hoboken, N.J.

Rising, E., Baron, R., Averill, B., 1973. A systems analysis of a university-health-service outpatient clinic. Oper. Res. 21 (5), 1030–1047.

Schütz, H., Kolisch, R., 2012. Approximate dynamic programming for capacity allocation in the service industry. Eur. J. Oper. Res. 218 (1), 239–250.

Schweitzer, P., Seidmann, A., 1985. Generalized polynomial approximations in Markovian decision processes. J. Math. Anal. Appl. 110 (2), 568–582.

Vermeulen, I., Bohte, S., Elkhuizen, S., Lameris, H., Bakker, P., Poutré, H. L., 2009. Adaptive resource allocation for efficient patient scheduling. Artificial Intelligence in Medicine 46 (1), 67–80.

Waaijer, A., Terhaard, C., Dehnad, H., Hordijk, G., Leeuwen, M. V., Raaymakers, C., Lagendijk, J., 2003. Waiting times for radiotherapy: Consequences of volume increase for the TCP in oropharyngeal carcinoma. Radiother. Oncol. 66 (3), 271–276.