

A new column-generation-based algorithm for VMAT treatment plan optimization

This article has been downloaded from IOPscience. Please scroll down to see the full text article.

2012 Phys. Med. Biol. 57 4569

(<http://iopscience.iop.org/0031-9155/57/14/4569>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 141.211.173.82

The article was downloaded on 26/06/2013 at 15:02

Please note that [terms and conditions apply](#).

A new column-generation-based algorithm for VMAT treatment plan optimization

Fei Peng^{1,2}, Xun Jia², Xuejun Gu², Marina A Epelman¹,
H Edwin Romeijn¹ and Steve B Jiang²

¹ Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, MI 48109, USA

² Center for Advanced Radiotherapy Technologies and Department of Radiation Medicine and Applied Sciences, University of California San Diego, La Jolla, CA 92093, USA

E-mail: feipeng@umich.edu and sbjiang@ucsd.edu

Received 20 January 2012, in final form 21 May 2012

Published 22 June 2012

Online at stacks.iop.org/PMB/57/4569

Abstract

We study the treatment plan optimization problem for volumetric modulated arc therapy (VMAT). We propose a new column-generation-based algorithm that takes into account bounds on the gantry speed and dose rate, as well as an upper bound on the rate of change of the gantry speed, in addition to MLC constraints. The algorithm iteratively adds one aperture at each control point along the treatment arc. In each iteration, a restricted problem optimizing intensities at previously selected apertures is solved, and its solution is used to formulate a pricing problem, which selects an aperture at another control point that is compatible with previously selected apertures and leads to the largest rate of improvement in the objective function value of the restricted problem. Once a complete set of apertures is obtained, their intensities are optimized and the gantry speeds and dose rates are adjusted to minimize treatment time while satisfying all machine restrictions. Comparisons of treatment plans obtained by our algorithm to idealized IMRT plans of 177 beams on five clinical prostate cancer cases demonstrate high quality with respect to clinical dose–volume criteria. For all cases, our algorithm yields treatment plans that can be delivered in around 2 min. Implementation on a graphic processing unit enables us to finish the optimization of a VMAT plan in 25–55 s.

(Some figures may appear in colour only in the online journal)

1. Introduction

Rotational delivery of intensity modulated radiation therapy (IMRT) was first proposed by Yu (1995) as intensity modulated arc therapy (IMAT). This radiation therapy treatment modality delivers radiation as the gantry rotates around the patient in one or more arcs while, at

the same time, field shapes (apertures) are formed to modulate fluence using a multi-leaf collimator (MLC) system. Recently, this modality has gained popularity and interest due to the introduction of several commercially available implementations, and is now usually referred to as volumetric modulated arc therapy (VMAT). These advances have also made it possible to dynamically change the dose rate and gantry speed. As compared to more standard IMRT treatments, which use a relatively small number of fixed-beam directions, VMAT has the potential to significantly reduce treatment time. This has the additional benefits of decreasing patient discomfort, mitigating intrafraction motion uncertainties and increasing the utilization rate of the equipment.

Since the conception of VMAT, researchers have investigated inverse planning techniques as well as methods to make VMAT plans deliverable in one single arc. Various studies utilized techniques including simulated annealing (Cameron 2005, Earl *et al* 2003) and transformation of an IMRT treatment plan (Crooks *et al* 2003) to design VMAT treatments. These studies showed that, with constant gantry speed and dose rate, either multiple arcs or an excessively long time for a single arc is required to deliver a high-quality treatment. In a theoretical study, Bortfeld and Webb (2009) compared single-arc VMAT treatment to traditional IMRT on a phantom and argued that single-arc VMAT may not be able to replicate the optimal intensity profile without intensity modulation at individual angles—a claim that was denied in a response by Otto (2009). Today, the potential quality of VMAT treatment plans is generally believed to be comparable to IMRT treatment plans (Palma *et al* 2008, Verbakel *et al* 2009).

The nature of VMAT treatments makes the treatment plan optimization problem much more complex than a corresponding IMRT treatment plan optimization problem due to the continuous nature of the gantry motion, gantry speed constraints, dose rate constraints and MLC leaf speed constraints. Even when the set of angles is discretized, exact methods quickly become intractable. Hence, the literature has exclusively focused on simplifications of the model and/or heuristic solution approaches. Many approaches to the VMAT treatment plan optimization problem rely on converting an IMRT plan in some fashion. Shepard *et al* (2007), Wang *et al* (2008) and Luan *et al* (2008) propose algorithms that rely on (i) first identifying an optimal IMRT treatment plan using 10° -spaced beam directions and then (ii) producing a deliverable VMAT plan by minimizing the difference between the sequenced and the ideal IMRT fluence maps. Cao *et al* (2009) started the treatment planning process by performing IMRT planning with direct machine parameter optimization, which directly optimizes aperture shapes and weights. They then used a simulated annealing-based arc sequencer to convert the resulting fluence maps into deliverable VMAT plans while again minimizing the difference between the IMRT and VMAT fluence maps. Craft *et al* (2011) performed IMRT optimization on 2° -spaced beam angles and, noting that adjacent fluence maps were similar, used a combination of beam merging and sliding window leaf sequences to obtain a VMAT deliverable plan whose quality is close to the 'ideal' IMRT plan. However, their observation that adjacent fluence maps in the high-resolution IMRT treatment plan are similar appears to depend strongly on the particular treatment plan optimization algorithm employed, so the approach may not be robust to changes in treatment plan evaluation criteria that would require using a different algorithm.

Several approaches to VMAT treatment plan optimization use the concept of *control points*, which represent beam directions along the arc where machine parameters such as gantry speed, aperture shape and dose rate are controlled. Between control points, the machine parameters are then either kept constant or interpolated in some fashion. Otto (2008) proposed a simulated annealing algorithm that adds control points sequentially and, in each iteration, randomly samples aperture weights and shapes. Bedford (2009) discretized an IMRT fluence map into multiple intensity levels and assigned each level to a control point close to the

corresponding angle in the IMRT plan. A subsequent local search procedure then attempts to improve leaf positions and segment weights. Bzdusek *et al* (2009) started with an IMRT plan on a coarse set of angles, then generated two apertures to approximate each of the fluence maps and placed these at predetermined control points. Additional control points were assigned interpolations of these apertures, and the remaining machine parameters were calculated in a subsequent optimization step. Gözbasi (2010) formulated the problem as a large integer programming model, but for tractability reasons resorted to grouping beamlets and solved a series of approximate problems. Unfortunately, the required computation time of the algorithm is still impractical and the algorithm also failed to provide clinically acceptable results on their test cases.

Finally, Men *et al* (2010b) proposed a column-generation-based approach which sequentially adds aperture shapes at control points while ensuring that each newly generated aperture is compatible with previously chosen ones. However, no bounds on the dose rate were taken into account, and gantry speed was considered to be constant. In this paper we formalize and extend this approach to explicitly take the following physical restrictions into account during the course of the algorithm:

- apertures at consecutive control points must be compatible, i.e. the leaves in individual rows of the MLC must be able to finish their transition from the previous aperture to the next aperture before the gantry reaches the next control point;
- restrictions on machine parameters, such as upper and lower bounds on the gantry speed and dose rate, as well as the upper bound on the rate of change of the gantry speed must be obeyed.

Although some of the algorithmic approaches proposed in the literature and discussed above account for some of these restrictions by either explicitly incorporating them into the treatment plan optimization or by modifying the optimized plan into one that satisfies them in a post-processing step; to the best of our knowledge, none incorporates all of the restrictions explicitly in the treatment plan optimization process. Our algorithm is implemented using the compute unified device architecture (CUDA) to take advantage of the parallel processing power of the graphic processing unit (GPU).

2. Methods and materials

2.1. The VMAT optimization model

Let K denote the total number of (not necessarily equispaced) control points along one or more treatment arc(s) (in the case of multiple arcs, K represents the total number of control points along all arcs, which will be ordered so that the first control point on the second arc succeeds the last control point on the first arc, etc). We associate with each control point k an aperture A_k , a dose rate r_k (in MU s^{-1}), and a gantry speed s_k (in deg s^{-1}). In addition, we can express the fluence rate (in MU deg^{-1}) at control point k as $y_k = r_k/s_k$. Note that each control point represents a snapshot of the continuous gantry rotation, i.e. r_k , s_k , y_k and A_k represent the state of the treatment machine as the gantry passes through control point k . For convenience, we will add a dummy control point $K + 1$ at the end of the arc. For tractability, we will then calculate the dose delivered to each voxel by making the approximation that the aperture, dose rate and gantry speed (and hence fluence rate) are constant throughout the arc spanning from one control point to the next. Note in particular that this means that we do not need to specify any of the variables at the dummy control point $K + 1$. If the angular distances δ_k between pairs of control points k and $k + 1$ are small, then this approximation will be sufficiently accurate (Otto 2008).

The gantry speeds and apertures specified at consecutive control points need to be compatible with each other. This compatibility requirement derives from the characteristics of the MLC system used to form the apertures during the treatment. As the gantry travels from control point k to control point $k + 1$, the leaves of the MLC system need to shift from positions prescribed by aperture A_k to those prescribed by aperture A_{k+1} . Since the speed of such leaf movement is bounded, the time spent by the gantry moving between the control points needs to be sufficiently large (and hence its speed sufficiently small) to allow for the required leaf movement to take place. We will denote the maximum gantry speed that will allow aperture A' at control point $k + 1$ to be reached from aperture A at control point k by $S_{k,k+1}^U(A, A')$ for $k = 1, \dots, K$, where, for convenience, $S_{K,K+1}^U(A, A') = \infty$ for all pairs of apertures A and A' .

In addition to the discretization parameters δ_k ($k = 1, \dots, K$), there are several other machine parameters that need to be taken into account. First, a VMAT delivery machine may have upper and lower bounds on gantry speed (denoted by S^U and S^L), and an upper bound on dose rate (denoted by R^U). Furthermore, there typically is an upper bound on the rate of change in speed that the gantry can sustain, which may be given per degree or per control point. We will denote the upper bound on the change in speed between control points k and $k + 1$ by ΔS_k . Finally, we let \mathcal{A} denote the set of all apertures that are deliverable by the MLC system. In particular, we will assume that left and right leaves within each pair can be positioned at any non-overlapping continuous points within their range, and that motion is independent between leaf pairs. In particular, we assume that there are no interdigitation constraints. After formulating our base model and developing our algorithm, we will discuss how additional constraints, such as a lower bound on dose rate, independent bounds on the fluence rate and interdigitation constraints, can be incorporated into our solution approach when necessary.

Finally, let the discretized set of voxels be \mathcal{V} , and denote the delivered dose distribution by $\mathbf{z} = (z_j : j \in \mathcal{V})^\top$, where z_j is the total dose delivered to voxel $j \in \mathcal{V}$. The quality of the dose distribution \mathbf{z} is evaluated using an objective function $F : \mathbb{R}^{|\mathcal{V}|} \rightarrow \mathbb{R}$. Under our approximation, z_j is equal to the sum of the doses delivered to the voxel from all control points. Moreover, since the delivered dose is linear in fluence, it is convenient to denote the dose received by voxel $j \in \mathcal{V}$ from aperture $A \in \mathcal{A}$ at control point k at unit fluence by $D_{kj}(A)$ ($k = 1, \dots, K$).

The complete VMAT optimization model, which we refer to as the full problem (FP), then reads:

$$\begin{aligned}
 \text{(FP) minimize} \quad & F(\mathbf{z}) \\
 \text{subject to} \quad & \\
 z_j = \sum_{k=1}^K D_{kj}(A_k) \delta_k y_k \quad & j \in \mathcal{V} \\
 y_k = \frac{r_k}{s_k} \quad & k = 1, \dots, K \\
 |s_k - s_{k+1}| \leq \Delta S_k \quad & k = 1, \dots, K - 1 \\
 s_k \in [S^L, S^U] \quad & k = 1, \dots, K \\
 r_k \in [0, R^U] \quad & k = 1, \dots, K \\
 s_k \leq S_{k,k+1}^U(A_k, A_{k+1}) \quad & k = 1, \dots, K \\
 A_k \in \mathcal{A} \quad & k = 1, \dots, K,
 \end{aligned} \tag{1}$$

where the terminal aperture A_{K+1} in (1) can be chosen arbitrarily and is added for convenience only.

We can simplify this model by first noting that for any feasible solution $\mathbf{s}, \mathbf{r}, \mathbf{y}, \mathbf{z}$ and $(A_k : k = 1, \dots, K)$ of (FP) we can obtain another feasible solution with the same apertures

and the same values of \mathbf{y} and \mathbf{z} (and hence the same objective function value), but with all the gantry speeds at their lower bound value S^L by scaling the dose rates by a factor of $S^L/s_k \leq 1$, $k = 1, \dots, K$, to compensate. As a result, we can eliminate the gantry speed variables from the model. In addition, we can then replace the upper bound constraints on the dose rates by upper bounds on fluence rate and eliminate the dose rate variables as well. The problem thus reduces to the following *master problem* (MP):

$$\begin{aligned} \text{(MP) minimize} \quad & F(\mathbf{z}) \\ \text{subject to} \end{aligned}$$

$$z_j = \sum_{k=1}^K D_{kj}(A_k) \delta_k y_k \quad j \in \mathcal{V} \quad (2)$$

$$\begin{aligned} y_k &\in [0, Y^U] & k = 1, \dots, K \\ S^L &\leq S_{k,k+1}^U(A_k, A_{k+1}) & k = 1, \dots, K \end{aligned} \quad (3)$$

$$A_k \in \mathcal{A} \quad k = 1, \dots, K, \quad (4)$$

where $Y^U \equiv R^U/S^L$. We will refer to (MP) as the master problem, for reasons made clear in the following section. By construction, for any feasible solution to (MP) we can construct an equivalent feasible solution to (FP), in the sense that it has the same apertures and values of \mathbf{y} and \mathbf{z} , by setting $s_k = S^L$ and $r_k = y_k s_k$ for $k = 1, \dots, K$. Moreover, there may be several combinations of gantry speeds and dose rates that lead to equivalent solutions to (FP), and some may be more desirable than others, e.g., due to shorter treatment times. We will discuss this in more detail in section 2.5.

Note that (MP) is not convex because of the nonlinearities associated with the apertures in constraints (2)–(4). Therefore, instead of solving problem (MP) exactly, we propose a column-generation-based algorithm.

2.2. An algorithm for solving (MP)

Our proposed algorithm starts with a dose distribution $\mathbf{z} = \mathbf{0}$ and without an aperture specified at any of the control points. The algorithm then proceeds by, in each iteration, attempting to improve the current treatment plan and dose distribution by selecting an aperture at one of the control points for which none has been specified yet (that is, *generating a column* in constraint (2) of (MP)), and terminates when a completely specified treatment plan is obtained. In any given iteration, which is characterized by a set $\mathcal{C} \subseteq \{1, \dots, K\}$ of control points and the corresponding collection of apertures $\{\bar{A}_k, k \in \mathcal{C}\}$, we optimize fluence rates of these apertures by solving the so-called *restricted master problem* (RMP). The RMP is constructed so that a feasible solution to this problem implicitly provides a deliverable plan that can be fully specified by feasibly completing the sequence of apertures while retaining the current aperture fluence rates (in particular, $y_k = 0$ for $k \notin \mathcal{C}$). The optimal solution to the RMP then defines the so-called *pricing problem* (PP) which is used to select the next control point and aperture to be added to the plan. The intuition behind the PP is that, relative to the current solution, each candidate aperture at each control point $k \notin \mathcal{C}$ has an associated *price* which is defined as the rate of improvement in the objective function value if the fluence rate of that aperture is increased (from its current value of 0). The PP then finds the control point/aperture combination that has the best price. In contrast with direct aperture optimization for IMRT treatment planning (Preciado-Walters *et al* 2004, Romeijn *et al* 2005, Shepard *et al* 2002), only a single aperture is added for each control point. This makes our algorithm heuristic in nature and, in particular, a greedy heuristic.

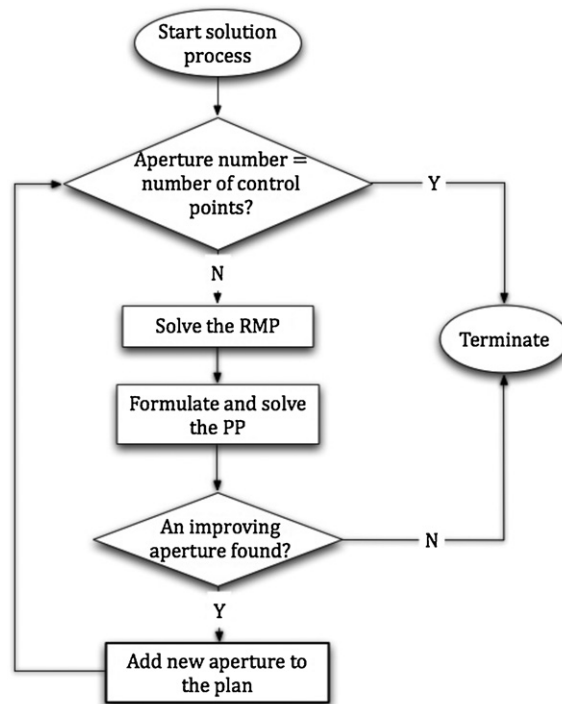


Figure 1. Flow chart for the column-generation-based greedy heuristic for (MP).

More formally, our algorithm can be described as follows:

Column generation based greedy heuristic for (MP)

Step 0 Set $\mathcal{C} = \emptyset$ and $\bar{\mathbf{z}} = \mathbf{0}$.

Step 1 Use the information on the current treatment plan (control points \mathcal{C} , apertures \bar{A}_k for $k \in \mathcal{C}$, and dose distribution $\bar{\mathbf{z}}$) to formulate and solve an instance of the PP.

Step 2 If the optimal value of the PP is nonpositive, go to step 5. Otherwise, denote the optimal solution to the PP by \bar{c} and $\bar{A}_{\bar{c}}$, and replace \mathcal{C} by $\mathcal{C} \cup \{\bar{c}\}$.

Step 3 Solve the instance of the RMP associated with \mathcal{C} and $A_k = \bar{A}_k$, $k \in \mathcal{C}$.

Step 4 If $|\mathcal{C}| < K$, return to step 1.

Step 5 If necessary, complete the treatment plan by identifying feasible apertures (which will have fluence 0) at control points $c \notin \mathcal{C}$, and denote the final set of fluence rates by \bar{y}_k ($k = 1, \dots, K$).

Figure 1 provides a graphical representation of the structure of this algorithm.

In the remainder of this section, we will provide additional details on steps 1–3 of this algorithm in sections 2.3 and 2.4. Section 2.5 discusses obtaining a solution to (FP) based on the solution to (MP) returned by our algorithm. Finally, section 2.6 discusses how additional constraints on the setting of the VMAT delivery machine can be incorporated into (FP) formulation and our solution.

2.3. Restricted master problem

At each iteration of the algorithm, we obtain the RMP from (MP) by fixing the apertures $A_k = \bar{A}_k$ for all control points in the current set $\mathcal{C} \subseteq \{1, \dots, K\}$ and setting the fluence rates $y_k = 0$ for all control points $k \notin \mathcal{C}$. The RMP then reads

$$\begin{aligned}
 \text{(RMP}^{(\mathcal{C})}) \quad & \text{minimize} \quad F(\mathbf{z}) \\
 & \text{subject to} \\
 & z_j = \sum_{k \in \mathcal{C}} D_{kj}(\bar{A}_k) \delta_k y_k \quad j \in \mathcal{V} \\
 & y_k \in [0, Y^U] \quad k \in \mathcal{C}.
 \end{aligned} \tag{5}$$

(RMP^(C)) can be thought of as a restriction of (MP) in that it considers a subset of control points, and considers the apertures \bar{A}_k , $k \in \mathcal{C}$, to be fixed and given. Note that (RMP^(C)) is a nonlinear optimization problem with linear constraints, and it is a convex optimization problem provided that F is a convex function.

Recall that any feasible solution to (RMP^(C)) should correspond to a deliverable plan that can be fully specified by completing the sequence of apertures at control points $k \notin \mathcal{C}$ in a way that is compatible with apertures \bar{A}_k , $k \in \mathcal{C}$, and setting $y_k = 0$, $k \notin \mathcal{C}$. To ensure that this is indeed the case, the set of apertures \bar{A}_k , $k \in \mathcal{C}$ in the restricted problem should satisfy the following condition. For all $k \in \mathcal{C}$, let k^+ be the successor of k in \mathcal{C} , i.e. $k^+ = \min\{k' \in \mathcal{C} : k' > k\}$, where $k^+ \equiv K + 1$ if $k = \max\{k' : k' \in \mathcal{C}\}$. Any feasible solution to (RMP^(C)) for any $\mathcal{C} \subseteq \{1, \dots, K\}$ can be extended to a feasible solution to (MP), provided that

$$S_{kk^+}^U(\bar{A}_k, \bar{A}_{k^+}) \geq S^L, \quad k \in \mathcal{C}. \tag{6}$$

Here we have generalized the notation of (FP) and (MP) by denoting the maximum gantry speed that will allow aperture A' at control point k' to be reached from aperture A at control point $k < k'$ by $S_{kk'}^U(A, A')$ (where again, for convenience, we will let $S_{k, K+1}^U(A, A') = \infty$ for all pairs of apertures A and A' and all $k = 1, \dots, K$). The PP described in the following section is designed to ensure that the apertures added at each iteration of the column generation algorithm satisfy (6).

2.4. Pricing problem

2.4.1. PP derivation. Consider an optimal solution $(\bar{\mathbf{y}}, \bar{\mathbf{z}})$ to (RMP^(C)). Suppose we were to add aperture $A \in \mathcal{A}_c$ at a control point $c \notin \mathcal{C}$, where $\mathcal{A}_c \subseteq \mathcal{A}$ is a set of deliverable apertures that can feasibly be added to the current treatment plan. Based on first-order optimality conditions, the rate of *improvement* (i.e. decrease) in objective function value per unit fluence rate in this aperture can be calculated as

$$\pi_c(A) \equiv \sum_{j \in \mathcal{V}} \pi_j(\bar{\mathbf{z}}) D_{cj}(A_c) \delta_c,$$

where $(\pi_j(\bar{\mathbf{z}}) : j \in \mathcal{V})^\top = \boldsymbol{\pi}(\bar{\mathbf{z}}) \equiv -\nabla F(\bar{\mathbf{z}})$. Our strategy will be to select a control point c and aperture A that solve the following optimization problem, which we will refer to as the *PP*:

$$\text{(PP)} \quad \max_{c \notin \mathcal{C}} \max_{A \in \mathcal{A}_c} \pi_c(A).$$

In order to ensure that solutions to PP satisfy condition (6), we need to specify the set \mathcal{A}_c in such a way that

$$\mathcal{A}_c \subseteq \{A \in \mathcal{A} : S_{c-c}^U(\bar{A}_{c^-}, A) \geq S^L, S_{cc^+}^U(A, \bar{A}_{c^+}) \geq S^L\} \equiv \mathcal{A}_c^U.$$

The choice of $\mathcal{A}_c \subseteq \mathcal{A}_c^U$ in the PP dictates a tradeoff in the behavior of our heuristic algorithm. On the one hand, choosing a smaller set \mathcal{A}_c at the current iteration results in lesser flexibility in the current selection of apertures. On the other hand, this can potentially allow for more flexibility in later iterations of the heuristic. We will consider the following family of potential choices for \mathcal{A}_c in (PP), parameterized by a lower bound s on the speed between two control points:

$$\mathcal{A}_c(s) = \{A \in \mathcal{A} : S_{c^-}^U(\bar{A}_{c^-}, A) \geq s, S_{c^+}^U(A, \bar{A}_{c^+}) \geq s\} \text{ for } s \geq S^L. \tag{7}$$

It is easy to see that $\mathcal{A}_c(s') \subseteq \mathcal{A}_c(s)$ whenever $s' \geq s$, with $\mathcal{A}_c(S^L) = \mathcal{A}_c^U$. Note also that, although the choice of s does not explicitly determine the final gantry speed used at different control points, because it affects the selection of apertures $\mathcal{A}_c(s)$, it implicitly affects the final gantry speed and thus the delivery time. We will explore the impact of the choice of parameter s on the overall performance of the algorithm in our experiments.

If the optimal value of (PP) is non-positive, we conclude that the current solution cannot be improved by adding any aperture at any control point (at least in the framework of our algorithm). In this case, the algorithm will be terminated (see step 2), and fluence rates at control points $k \notin \mathcal{C}$ will remain at zero. However, if an improving control point/aperture is found, the algorithm will proceed to step 3.

2.4.2. Solving the PP. We will now describe a solution approach for the PP. First, it is easy to see that this problem decomposes into a collection of independent problems for the different candidate control points c , so that we can focus on the problem for a given $c \notin \mathcal{C}$:

$$(PP^{(c)}) \max_{A \in \mathcal{A}_c} \sum_{j \in \mathcal{V}} \pi_j(\bar{\mathbf{z}}) D_{cj}(A),$$

where we have also eliminated the scaling constant δ_c .

Next, in the absence of interdigitation constraints, we can decompose $PP^{(c)}$ by MLC row. In particular, if the MLC system consists of M leaf pairs, we can represent any aperture $A \in \mathcal{A}$ as a collection of leaf settings for its individual rows. Let us denote an aperture by $A = (a_1, \dots, a_M)$, where $a_m \in \alpha$ describes the leaf settings of MLC row m of aperture A , and α is the set of deliverable ‘row apertures’. The set \mathcal{A}_c can be written as $\mathcal{A}_c = \times_{m=1}^M \mathcal{A}_{cm}$, with

$$\mathcal{A}_{cm} \subseteq \{a : S_{c^-}^U(\bar{a}_{c^-m}, a) \geq S^L, S_{c^+}^U(a, \bar{a}_{c^+m}) \geq S^L\}, \tag{8}$$

where \bar{a}_{km} is the m th row aperture of \bar{A}_k at control point $k \in \mathcal{C}$ and, with a slight abuse of notation, $S_{kc}^U(a, a')$ denotes the maximum gantry speed that will allow row aperture a' of aperture A' at control point c to be reached from row aperture a at control point $k < c$. Let $D_{cmj}(a)$ denote the dose received by voxel $j \in \mathcal{V}$ from row aperture $a \in \alpha$ in row m and control point c at unit fluence rate, with $D_{cj}(A) = \sum_{m=1}^M D_{cmj}(a_m)$. The pricing problem $(PP^{(c)})$ can then be decomposed by row:

$$(PP^{(cm)}) \max_{a \in \mathcal{A}_{cm}} \sum_{j \in \mathcal{V}} \pi_j(\bar{\mathbf{z}}) D_{cmj}(a).$$

Recall that we would like to consider the family of feasible regions $\mathcal{A}_c = \mathcal{A}_c(s)$ defined by (7) for $s \geq S^L$. The corresponding family $\mathcal{A}_{cm}(s)$ of feasible regions for $(PP^{(cm)})$ is

$$\mathcal{A}_{cm}(s) = \{a : S_{c^-}^U(\bar{a}_{c^-m}, a) \geq s, S_{c^+}^U(a, \bar{a}_{c^+m}) \geq s\} \text{ for } s \geq S^L. \tag{9}$$

Let us analyze the structure of these sets in more detail. Any row aperture a can be characterized as a pair of leaf settings (ℓ, r) , where $0 \leq \ell \leq r \leq N$ are the positions of the left and right leaves, respectively (N , which we assume to be integer, is the range of the MLC row).

Furthermore, let v denote the maximum leaf speed (in distance/second, with the unit of distance depending on the choice of N). Then, for $k < c$,

$$\{a \in \alpha : S_{kc}^U(a, a') \geq s\} = \left\{ (\ell, r) : 0 \leq \ell \leq r \leq N; |\ell - \ell'|, |r - r'| \leq \frac{v\delta_{kc}}{s} \right\},$$

where $\delta_{kc} = \sum_{k'=k}^{c-1} \delta_{k'}$. Therefore,

$$\mathcal{A}_{cm}(s) = \left\{ (\ell, r) : 0 \leq \ell \leq r \leq N; |\ell - \bar{\ell}_{c-m}|, |r - \bar{r}_{c-m}| \leq \frac{v\delta_{c-c}}{s}; |\ell - \bar{\ell}_{c+m}|, |r - \bar{r}_{c+m}| \leq \frac{v\delta_{c+c}}{s} \right\}, \tag{10}$$

with the appropriate and obvious definitions of $(\bar{\ell}_{c-m}, \bar{r}_{c-m})$ and $(\bar{\ell}_{c+m}, \bar{r}_{c+m})$. Thus, the set $\mathcal{A}_{cm}(s)$ is simply the set of all (ℓ, r) such that $0 \leq \ell \leq r \leq N$ and with upper and lower bounds on ℓ and r which get tighter as s increases.

Finally, in order to fully specify $(PP^{(cm)})$, we need to characterize $D_{cmj}(a) = D_{cmj}(\ell, r)$. Specifying $D_{cmj}(\ell, r)$ as a function of continuous variables (ℓ, r) will require an immense amount of data. Instead, we use a common approximation obtained by discretizing each MLC row into N beamlets (thus specifying an appropriate MLC row range), where beamlet n represents the interval $[n-1, n]$. We then precompute traditional beamlet-based dose deposition coefficients D_{cmnj} , i.e. the dose rate to voxel $j \in \mathcal{V}$ from beamlet n in MLC row m at control point c ($n = 1, \dots, N, m = 1, \dots, M, c = 1, \dots, K$) at unit fluence rate and make the following approximation:

$$D_{cmj}(\ell, r) = \int_{\ell}^r \phi_{cmj}(x) dx,$$

where $\phi_{cmj} : [0, N] \rightarrow \mathbb{R}^+$ is the following step function:

$$\phi_{cmj}(x) = D_{cmnj}, \quad n-1 < x \leq n; \quad n = 1, \dots, N.$$

To summarize, the pricing problem (PP^{cm}) with $\mathcal{A}_{cm} = \mathcal{A}_{cm}(s)$ can be written as

$$\max_{(\ell, r) \in \mathcal{A}_{cm}(s)} \int_{\ell}^r \pi_{cm}(x; \bar{\mathbf{z}}) dx, \tag{11}$$

where

$$\pi_{cm}(x; \bar{\mathbf{z}}) = \sum_{j \in \mathcal{V}} \pi_j(\bar{\mathbf{z}}) \phi_{cmj}(x).$$

The problem (11) can be solved efficiently by noting that the only candidate values for ℓ and r that have to be considered are (i) integers, and (ii) the bounds derived from (10).

2.5. From (MP) to (FP): gantry speeds and dose rates

Our column generation algorithm will return a feasible solution $\bar{\mathbf{y}}, \bar{\mathbf{z}}$ and $\bar{A}_k, k = 1, \dots, K$ to the problem (MP). As a final step of our solution procedure, we need to compute gantry speed and dose rate vectors \mathbf{s} and \mathbf{r} consistent with this solution. Substituting the apertures and fluence rates into constraints of (FP), we need to find a feasible solution to the following system:

$$\bar{y}_k = \frac{r_k}{s_k} \quad k = 1, \dots, K \tag{12}$$

$$|s_k - s_{k+1}| \leq \Delta S_k \quad k = 1, \dots, K-1 \tag{13}$$

$$s_k \in [S^L, S^U] \quad k = 1, \dots, K \tag{14}$$

$$r_k \in [0, R^U] \quad k = 1, \dots, K \quad (15)$$

$$s_k \leq S_{k,k+1}^U(\bar{A}_k, \bar{A}_{k+1}) \quad k = 1, \dots, K. \quad (16)$$

This system has a feasible solution since the apertures and fluence rates computed by the algorithm are guaranteed to be compatible with the gantry moving at speed S^L . However, this solution would lead to a treatment with undesirably long delivery time. Instead, we would like to identify a solution to (12)–(16) that can be delivered in the shortest time possible. Adding this objective function and rewriting constraints (12)–(16) in terms of the gantry speeds only, we obtain the following convex optimization problem:

$$\begin{aligned} \text{(SP) minimize} \quad & \sum_{k=1}^K \frac{\delta_k}{s_k} \\ \text{subject to} \quad & s_k \in [S^L, S_k^U] \quad k = 1, \dots, K \\ & |s_k - s_{k+1}| \leq \Delta S_k \quad k = 1, \dots, K-1, \end{aligned} \quad (17)$$

where

$$S_k^U = \min \left\{ S^U, S_{k,k+1}^U(\bar{A}_k, \bar{A}_{k+1}), \frac{R^U}{\bar{y}_k} \right\}, \quad k = 1, \dots, K.$$

If \bar{s}_k is an optimal solution to this problem, then the corresponding dose rates are $\bar{r}_k = \bar{y}_k \bar{s}_k$, $k = 1, \dots, K$.

In order to assess the impact on delivery time of the bound on the change in speed we will also consider a relaxation of (SP) in which constraints (17) are removed. It is easy to see that the corresponding solution will be to choose $s_k = S_k^U$ and $r_k = \bar{y}_k S_k^U$ (for $k = 1, \dots, K$).

2.6. Extensions

We conclude this section by discussing how additional constraints on the setting of the VMAT delivery machine can be incorporated into the model (FP), its reformulation (MP) and our solution heuristic.

2.6.1. Upper bound on fluence rate. Our basic model contained no bounds on the fluence rate, other than those implied by bounds on the gantry speed and dose rate. An independent upper bound on the fluence rate can be easily incorporated by including it among the constraints of (FP), and by calculating Y^U in (MP) and (RMP^(C)) as the minimum of this upper bound and the ratio R^U/S^L .

2.6.2. Lower bounds on fluence and dose rates. Our basic model did not include lower bounds on fluence and dose rates, other than nonnegativity constraints. If such lower bounds Y^L and R^U , respectively, are dictated by the VMAT delivery specification, they can be easily included among the constraints of (FP). However, inclusion of one or both such bounds invalidates the reformulation steps that lead to the equivalent problem (MP). We could then, instead, consider a generalization (MP) in which we add lower bounds on the fluence which are the larger of Y^L and R^L/S^U and apply our heuristic solution procedure without modification (aside from the obvious addition of lower bounds on fluence rates to (RMP^(C))). The resulting mater problem, however, is no longer equivalent to (FP), but rather is a relaxation. The consequences of using this relaxation are:

- (i) An intermediate solution obtained by solving $(RMP^{(C)})$ with $|\mathcal{C}| < K$ may not be extendable to a feasible solution to (FP). This means, in particular, that we have to modify step 2 of the greedy heuristic and add an aperture even if the optimal solution value to the PP is nonpositive.
- (ii) Postprocessing problem (SP) may not have a feasible solution. Therefore, we may have to slightly modify the final treatment plan to be able to satisfy the constraints on the rate of change in gantry speed.

2.6.3. *Interdigitation and other MLC constraints.* Lastly, consider a VMAT delivery system that does not allow interdigitation of MLC leaves in adjacent rows. The main impact of this constraint is a modification of the solution process of the PP discussed in section 2.4.2.

Before discussing the required modification, we will argue that if interdigitation constraints are satisfied in each iteration of the algorithm (i.e. for all apertures $A_k, k \in \mathcal{C}$), and we obtain a VMAT plan by linearly interpolating these apertures between control points in \mathcal{C} , then the interdigitation constraints will also hold for MLC apertures throughout the treatment. To obtain the interpolation, we assume that both the gantry and each of the leaves move at constant speeds between two consecutive control points. Let us examine two adjacent MLC rows m and $m + 1$ at two control points k and k^+ . We assume that

$$\bar{\ell}_{km} \leq \bar{r}_{k,m+1} \text{ and } \bar{\ell}_{k^+m} \leq \bar{r}_{k^+,m+1}, \tag{18}$$

i.e. the left leaf in row m and the right leaf in row $m + 1$ do not overlap at either control point, and denote the time it takes the gantry to rotate from k to k^+ by Δt . At an intermediate time $t \in (0, \Delta t)$, with the gantry at location $c : k < c < k^+$, the leaf positions are computed by linear interpolation as follows:

$$\ell_m = \bar{\ell}_{km} + (\bar{\ell}_{k^+m} - \bar{\ell}_{km}) \cdot \frac{t}{\Delta t} = \bar{\ell}_{km} \cdot \frac{\Delta t - t}{\Delta t} + \bar{\ell}_{k^+m} \cdot \frac{t}{\Delta t}$$

and

$$r_{m+1} = \bar{r}_{k,m+1} + (\bar{r}_{k^+,m+1} - \bar{r}_{k,m+1}) \cdot \frac{t}{\Delta t} = \bar{r}_{k,m+1} \cdot \frac{\Delta t - t}{\Delta t} + \bar{r}_{k^+,m+1} \cdot \frac{t}{\Delta t}.$$

It is easy to see that because of (18), we have $\ell_m \leq r_{m+1}$; using the same technique, we can show that $\ell_{m+1} \leq r_m$, assuming same is true at control points k and $k + 1$. Thus, at any location between the control points, the interdigitation constraints are satisfied.

Now we study the impact of the interdigitation constraints on the PP. Because of the dependence between individual MLC rows created by these constraints, we cannot no longer decompose the problem $PP^{(C)}$ by row. As before, we characterize row aperture $a_m \in \alpha$ at MLC row m of aperture A by (ℓ_m, r_m) where $0 \leq \ell_m \leq r_m \leq N$ and the (integer) N represents the range of the MLC row, and denote by $S_{kc}^U(a'_m, a_m)$ the maximum gantry speed that will allow row m of aperture $A = (a_1, \dots, a_M)$ at a control point c to be reached from row m of aperture $A' = (a'_1, \dots, a'_M)$ at control point $k < c$. To incorporate the interdigitation constraints, we add the inequalities

$$\ell_m \leq \min\{r_{m+1}, r_{m-1}\}, \quad m = 1, \dots, M \tag{19}$$

to the description of the set $A_c(s)$ in (7), where we let $r_0 = r_{M+1} = \infty$. By construction, $A_c(s) \neq \emptyset$ throughout the algorithm.

Table 1. Problem dimensions of different cases (after downsampling).

Case	No. of voxels	No. of beamlets	No. of $D_{cmj} \neq 0$	δ_k ($k = 2, \dots, 175$)	δ_1, δ_{176}
1	13 047	21 417	89 360 831	$\frac{330}{175}$	$\frac{1}{2} \cdot \frac{330}{175}$
2	10 130	17 700	63 359 780	$\frac{330}{175}$	$\frac{1}{1} \cdot \frac{330}{175}$
3	12 102	29 913	106 889 530	$\frac{330}{175}$	$\frac{1}{2} \cdot \frac{330}{175}$
4	9602	21 417	79 889 713	$\frac{330}{175}$	$\frac{1}{1} \cdot \frac{330}{175}$
5	13 769	25 488	114 315 187	$\frac{340}{175}$	$\frac{1}{2} \cdot \frac{340}{175}$

The PP for $c \notin \mathcal{C}$ under interdigitation constraints can then be written as

$$\begin{aligned}
 (\text{PPI}^{(c)}) \text{ maximize } & \sum_{j \in \mathcal{V}} \pi_j(\bar{\mathbf{z}}) \sum_{m=1}^M D_{cmj}(\ell_m, r_m) \\
 \text{subject to} & \\
 & 0 \leq \ell_m \leq r_m \leq N \quad m = 1, \dots, M \\
 & |\ell_m - \bar{\ell}_{c-m}|, |r_m - \bar{r}_{c-m}| \leq \frac{v\delta_{c-c}}{s} \quad m = 1, \dots, M \quad (20)
 \end{aligned}$$

$$\begin{aligned}
 & |\ell_m - \bar{\ell}_{c+m}|, |r_m - \bar{r}_{c+m}| \leq \frac{v\delta_{c+}}{s} \quad m = 1, \dots, M \quad (21) \\
 & \ell_m \leq \min\{r_{m+1}, r_{m-1}\} \quad m = 1, \dots, M,
 \end{aligned}$$

where $D_{cmj}(\ell_m, r_m)$, as before, is approximated by an integral of the beamlet-based step function $\phi_{cmj} : [0, N] \rightarrow \mathbb{R}^+$. Given this discretization of each MLC row into beamlets, this version of the PP can still be solved efficiently using a dynamic programming approach described in Romeijn *et al* (2005), with potential left and right leaf positions including integers and bounds in (20) and (21).

3. Experiments and results

3.1. Data and implementation

Our algorithm was evaluated on a clinical dataset of prostate cancer patients treated at UCSD Moores Cancer Center. The preprocessing steps, including image processing and calculation of the dose deposition coefficients D_{cmj} are performed using our in-house treatment planning system. The dose calculations in this system are based on a finite-sized pencil-beam algorithm with 3D density correction (Gu *et al* 2009a, 2009b). For each case, we used beamlets of size $1 \times 1 \text{ cm}^2$ and voxels of size $4 \times 4 \times 2.5 \text{ mm}^3$. However, in order to reduce its size, the optimization problem used a downsampled voxel grid that selected one of every two voxels along each of the three dimensions in critical structures and one of every four in unspecified tissue. The dimensions (after downsampling) and other characteristics of the five cases are summarized in table 1. Figure 2 illustrates the distribution of control points around the arc. For all cases we used $K = 176$ control points, and the angular distances between control points are shown in table 1. Finally, the machine parameters we used are based on the Varian RapidArc machine (Varian Medical Systems, Inc., Palo Alto, CA, USA) and are shown in table 2 (where $\Delta S_k = \Delta S$ for all $k = 1, \dots, K$).

The prescription dose to the PTV was set to 79.2 Gy, corresponding to a 44-fraction treatment with 1.8 Gy delivered in each fraction. All dose–volume histogram (DVH)-based criteria for the PTV and critical structures, in particular rectum, bladder and femoral heads, that were used to evaluate all treatment plans are based on clinical protocols at UCSD as well




Figure 2. Distribution of 177 control points around the arc for (a) cases 1–4, and (b) case 5.

Table 2. Machine parameters.

R^L (MU s ⁻¹)	R^U (MU s ⁻¹)	S^L (deg s ⁻¹)	S^U (deg s ⁻¹)	ΔS (deg s ⁻¹)	v (cm s ⁻¹)
0	10	0.83	6.0	0.75	2.25

as RTOG protocol Radiation Therapy Oncology Group (2004) (see the tables in the remainder of this section for details). In our optimization model, we chose F to be a voxel-based penalty function that penalizes the dose received by a voxel quadratically but asymmetrically with respect to a threshold dose that depends on the structure containing the voxel; consequently, F is a smooth piece-wise quadratic convex function. The structure-dependent relative weights of these penalty functions were determined using a manual procedure that iteratively modified these weights until satisfactory treatment plans were obtained for all five cases. This led to a common set of weights for all cases, thereby eliminating variation between cases caused by individual parameter tuning. In order to facilitate the comparison between the treatment plans, we normalized each treatment plan by scaling the dose distribution so that 95% of the PTV receives the prescription dose.

Finally, we implemented our algorithm using the CUDA to take advantage of the processing power of the GPU. CUDA is a platform for the implementation of large-scale parallelization of computer code, and has been shown to be very suitable for solving radiation therapy treatment plan optimization problem (Men *et al* 2009, 2010a, 2010b). The RMP is solved using our in-house gradient-based solver. On an Nvidia Tesla C1060 GPU card, our algorithm takes 22–55 s to complete. The high efficiency of our VMAT optimization engine provides a significant step toward a clinical application of VMAT for adaptive radiotherapy, although of course a treatment planning system for adaptive radiotherapy would contain several additional components.

3.2. Benchmark

Before evaluating our VMAT treatment plans, we provide a benchmark for comparison by solving an IMRT treatment plan optimization problem with the same objective function as

Table 3. Performance of 177-beam IMRT treatment plans.

Structure	Threshold Dose (Gy)	Volume Criterion (%)	Case				
			1	2	3	4	5
PTV	73.7	≥99	100	100	100	99	100
	79.2	≥95	95	95	95	95	95
	87.1	≤10	0	1	0	1	0
Rectum	75	≤15	4	16	14	9	9
	70	≤25	6	18	16	11	11
	65	≤35	8	20	19	14	13
	40	≤45	19	30	22	38	25
	65	≤17	7	23	19	42	11
Bladder	40	≤35	14	41	30	65	21
	50	≤10	0/0	0/0	0/0	0/0	0/0
Femoral heads (L/R)	45	≤25	0/0	0/0	0/0	0/0	0/0
	40	≤40	0/0	0/0	0/0	0/1	0/0

(FP) and using the $K + 1 = 177$ beam angles that correspond to control points in the VMAT treatment plan optimization model. These IMRT plans are used to assess the quality of our VMAT plans and are not suitable for clinical use due to the unreasonably long time required to deliver an IMRT plan with 177 beam angles. Since the IMRT treatment plan optimization model allows intensity modulation at each beam angle and contains no constraints on aperture compatibility nor gantry speed, they can be viewed as an ideal, likely clinically unattainable, solution. Therefore, any VMAT treatment plan that is close to this ideal IMRT treatment plan will also be close to the optimal VMAT treatment plan.

Table 3 includes the clinical dose–volume criteria and the 177-beam IMRT plans evaluated with these criteria for all five cases. All plans are able to avoid significant hot and cold spots in the PTV, and are able to satisfy most of the DVH criteria for the critical structures. For case 2, 14% of the bladder volume and 18% of the rectum volume overlap with the PTV, and for case 4 almost 30% of the bladder volume overlaps with the PTV. For these patients, even additional patient-dependent parameter tuning did not allow us to resolve the violations highlighted in the table without introducing more serious ones.

3.3. Performance of VMAT plans

As discussed in section 2.4.1, the column generation heuristic is parameterized by the parameter s , which represents how ‘greedy’ the heuristic is. Moreover, since the choice of s affects the choice of \bar{A}_k , which affects S_k^U in the time-minimization problem (SP), it also impacts the delivery time of the final treatment plans. In this section we test our heuristic under various settings of s , and study the performance of the resulting treatment plans in terms of quality and delivery time.

3.3.1. The most greedy heuristic. Recall that a smaller value of s means that there are more apertures to choose from in each iteration. We will start by choosing the smallest value of s ; i.e. we will start by investigating the most greedy variant of our heuristic. Note that we use the same objective function as for the 177-beam IMRT optimization without any additional tuning, so that the results only reflect the differences between the two modalities and treatment planning techniques.

Table 4. Performance of VMAT treatment plans with $s = S^L = 0.83 \text{ deg s}^{-1}$.

Structure	Threshold Dose (Gy)	Volume Criterion (%)	Case				
			1	2	3	4	5
PTV	73.7	≥ 99	100	100	100	99	99
	79.2	≥ 95	99	95	95	95	95
	87.1	≤ 10	0	1	0	1	0
Rectum	75	≤ 15	4	16	15	9	10
	70	≤ 25	7	19	18	13	13
	65	≤ 35	9	23	21	18	17
	40	≤ 45	26	38	36	36	33
Bladder	65	≤ 17	7	24	14	45	12
	40	≤ 35	14	43	22	68	23
Femoral heads (L/R)	50	≤ 10	1/0	0/0	1/0	1/1	2/0
	45	≤ 25	2/0	1/0	12/2	3/4	4/0
	40	≤ 40	6/0	3/1	16/6	7/7	12/1
Runtime (s)			50.2	55.1	47.2	25.0	35.6
Treatment time (s)			316.0	285.6	318.9	299.8	324.6

The results in table 4 show that our VMAT treatment plans exhibit high quality in general, closely resembling the IMRT treatment plans when evaluated with the same set of clinical criteria. The major differences between the IMRT and VMAT plans are with respect to the dose distribution in the femoral heads, but the VMAT plans still easily meet all clinical criteria.

The second to last row of table 4 contains the plan optimization times, in seconds, using the algorithm for each case, while the last line shows the treatment delivery time for each case. These treatment times range between 4.75 and just under 5.5 min.

3.3.2. Effect of the value of s on treatment quality and time. Unfortunately, the treatment plans in table 4 require relatively long treatment times, which negates some of the benefits of using VMAT treatments. We therefore study next whether we can choose a larger value of s , which can be expected to decrease treatment time, without significant deterioration in treatment plan quality. In these tests we focus on case 2, which is one of the more challenging cases.

Table 5 shows the performance of the VMAT treatment plans obtained with various values of s ranging from S^L to S^U . We can conclude that the treatment plan quality is not sensitive to the choice of s . However, the overall treatment time is reduced dramatically, by a factor of more than 2 from almost 5 min to just over 2 min, when using $s = S^U = 6$.

3.3.3. Least greedy heuristic. Based on the results in the previous section, we focus our further analysis on obtaining shorter treatment times. Table 6 shows the performance of the VMAT treatment plans obtained with $s = S^U = 6$. Similar to case 2 we conclude that there is no significant clinical deterioration in treatment plan quality as compared to the plans in table 4, but the treatment time is greatly reduced for all cases. In particular, the treatment time is consistently just over 2 min for all five cases. Finally, we provide an estimate of the total number of MUs required for treatment given by $\sum_{k=1}^K \delta_k \bar{y}_k$ where \bar{y}_k is the fluence rate at control point k in the final solution (see also the definition of the fluence rates y_k in section 2.1).

3.3.4. VMAT treatment plan quality as compared to benchmark. Figure 3 shows the DVHs of the 177-beam IMRT and VMAT (with $s = S^U = 6$) treatment plans for case 2. As can be

Table 5. Performance of VMAT treatment plans for case 2 obtained with different s values.

Structure	Threshold Dose (Gy)	Volume Criterion (%)	s (deg s ⁻¹)			
			0.83	2	4	6
PTV	73.7	≥99	100	100	100	99
	79.2	≥95	95	95	95	95
	87.1	≤10	1	0	1	1
Rectum	75	≤15	16	16	17	16
	70	≤25	19	19	20	20
	65	≤35	23	22	23	23
	40	≤45	38	35	37	40
Bladder	65	≤17	24	24	25	24
	40	≤35	43	42	43	45
Femoral heads (L/R)	50	≤10	0/0	0/0	0/0	0/4
	45	≤25	1/0	0/1	1/4	0/9
	40	≤40	3/1	3/5	14/10	0/16
Runtime (s)			55.1	55.5	29.9	35.2
Treatment time (s)			285.6	168.6	133.1	128.4

Table 6. Performance of VMAT treatment plans with $s = S^U$.

Structure	Threshold Dose (Gy)	Volume Criterion (%)	Case				
			1	2	3	4	5
PTV	73.7	≥99	99	99	99	99	99
	79.2	≥95	95	95	95	95	95
	87.1	≤10	1	1	3	4	4
Rectum	75	≤15	5	17	15	10	11
	70	≤25	7	20	19	14	16
	65	≤35	10	23	22	20	20
	40	≤45	26	40	36	35	40
Bladder	65	≤17	7	24	14	45	12
	40	≤35	14	45	24	70	25
Femoral heads (L/R)	50	≤10	0/5	0/4	1/1	4/3	7/0
	45	≤25	0/9	0/9	3/2	10/9	13/3
	40	≤40	7/15	0/16	8/5	14/13	24/9
Runtime (s)			24.8	35.2	29.4	22.8	28.5
Treatment time (s)			128.1	128.4	121.7	121.7	122.5
Total MU			503.2	534.7	492.8	483.1	507.6

gleaned from the corresponding tables, the largest differences in these plans are with respect to the dose distribution in the femoral heads, but these are still well within the clinical protocol. Given the restrictions that VMAT delivery imposes on a treatment plan, we feel that the quality of the VMAT plan is not only clinically acceptable, but also high since the degradation from the idealized 177-beam IMRT plan is relatively small.

3.3.5. Disallowing interdigitation. In order to study the effect of using an MLC system that does not allow interdigitation, we implemented the modified pricing problem (PPI) described in section 2.6.3 and applied our resulting algorithm on all five cases. Our experiments revealed no significant change in solution quality, while the solution time increased by less than 7.5 s in all cases (see table 7).




Figure 3. Treatment plans for case 2. Solid: 177-beam IMRT. Dashed: VMAT.

Table 7. Impact of disallowing interdigitation on algorithm solution time.

Solution time (s)	Case				
	1	2	3	4	5
Allowing interdigitation	24.8	35.2	29.4	22.8	28.5
Disallowing interdigitation	30.4	34.8	34.6	25.8	35.8

3.3.6. Rate of change in gantry speed. Finally, we investigate the effect of the constraint on the rate of change in gantry speed on treatment time. As discussed in section 2.5, we can assess the impact of the rate of change restriction by removing the corresponding constraint (17) from problem (SP). As before, we focus this study on case 2.

Figure 4 plots the gantry speeds and dose rates for control points 1–176 for case 2. Graphs (a-1) and (b-1) show gantry speeds, and (a-2) and (b-2) show dose rates, obtained by setting $s = S^U = 6$ in the column generation algorithm, with rate of change constraint (17) omitted for graphs (a-1) and (a-2), but included for graphs (b-1) and (b-2). Graphs (c-1)–(d-2) contain analogous results with $s = 4$. It is clear that without a constraint on the rate of change, the gantry speed changes rapidly during the treatment in order to extend the time spent at apparently favorable control points while limiting the time spent at less favorable ones. Note as well that at most control points, the speed remains equal to the value s used in the column generation algorithm even after solving (SP).

When the rate of change constraint is enforced, as in parts (a-1) and (c-1), the sequence of speeds is smoothed so that the change in speed between control points is no more than ΔS_k . Note that, as discussed in section 2.5, the rate of change constraint does not have any impact on treatment plan quality. However, if it were technologically feasible to increase or eliminate the bound on the rate of change in gantry speed, the treatment time could be shortened by about 30% to less than 1.5 min, as shown in table 8.

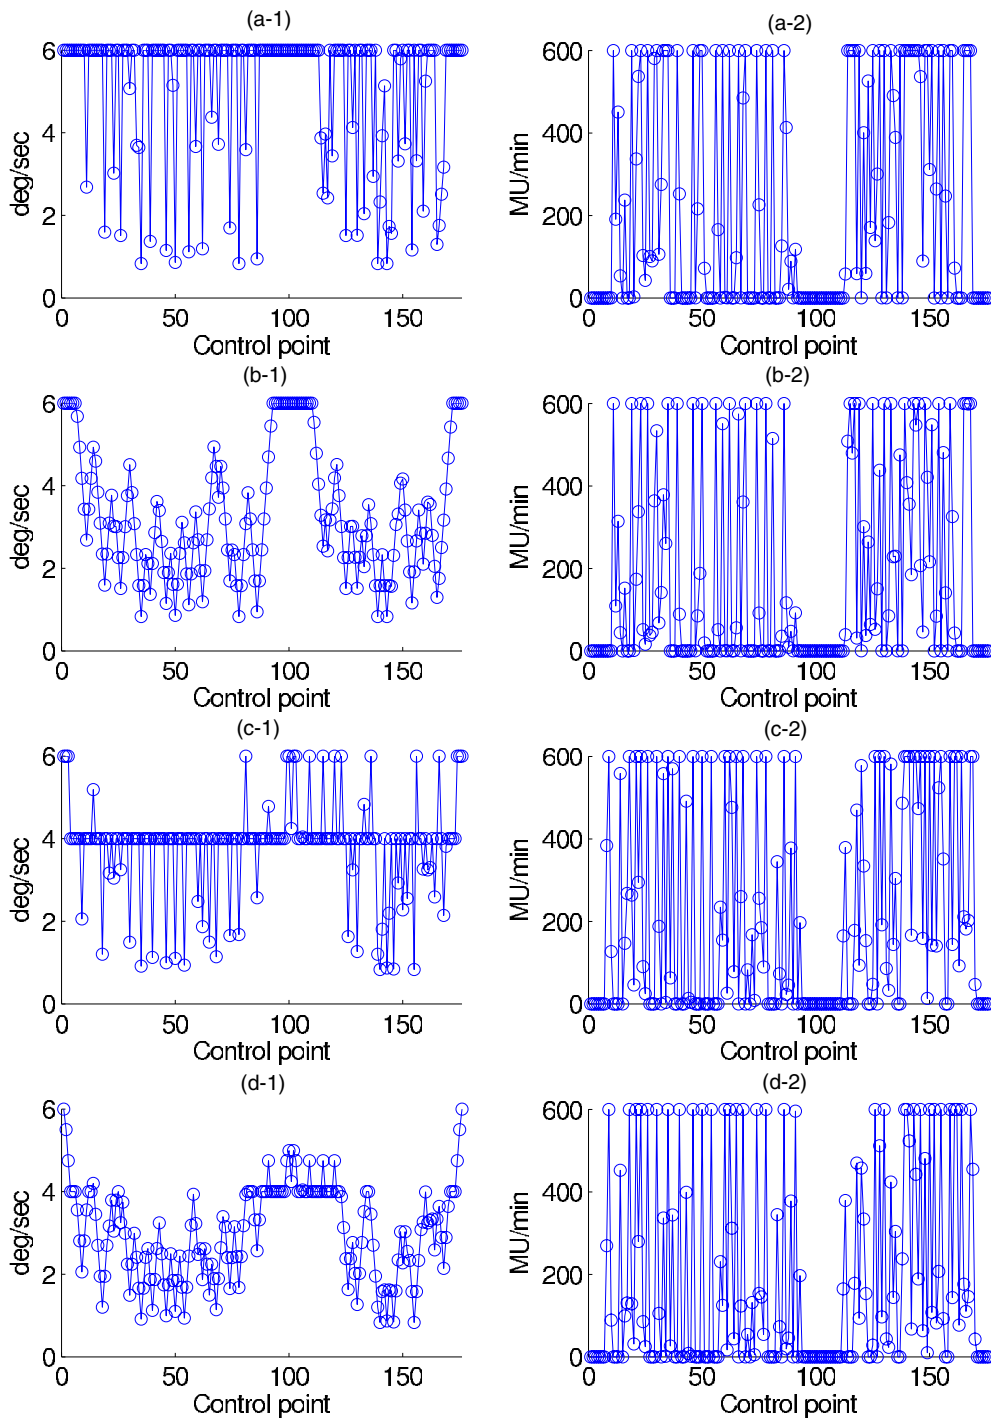


Figure 4. Gantry speeds (left column) and dose rates (right column) versus control points for case 2: (a) $s = 6 \text{ deg s}^{-1}$, ignoring rate of change constraint; (b) $s = 6 \text{ deg s}^{-1}$, considering rate of change constraints; (c) $s = 4 \text{ deg s}^{-1}$, ignoring rate of change constraint; (d) $s = 4 \text{ deg s}^{-1}$, considering rate of change constraints.

Table 8. Impact of the presence of the bound ΔS on the rate of change in gantry speed on treatment time.

Treatment time (s)	Case				
	1	2	3	4	5
Without bound	86.0	88.0	84.3	85.4	86.2
With bound	128.1	128.4	121.7	121.7	122.5

4. Conclusions

We have developed an efficient algorithm for generating VMAT treatment plans. Our approach incorporates various physical constraints, including bounds on dose rate, gantry speed and rate of change in the gantry speed during treatment. Moreover, the approach is flexible enough to allow for additional constraints, including MLC delivery constraints. We have shown that our algorithm is capable of generating treatment plans of high quality on five prostate cancer cases. The VMAT plans produced by our algorithm can be delivered in around 2 min. Moreover, a GPU implementation allows us to complete the optimization within 1 min. Along with additional 15 s for computing the dose deposition coefficients before the optimization and 15 s for computing the final dose distribution after the optimization using a GPU-based dose calculation engine (Gu *et al* 2009a, 2009b), this work provides a major step toward a clinical application of adaptive VMAT therapy.

References

- Bedford J 2009 Treatment planning for volumetric modulated arc therapy *Med. Phys.* **36** 5128–38
- Bortfeld T and Webb S 2009 Single-arc IMRT? *Phys. Med. Biol.* **54** N9–20
- Bzdusek K, Friberger H, Eriksson K, Hårdemark B, Robinson D and Kaus M 2009 Development and evaluation of an efficient approach to volumetric arc therapy planning *Med. Phys.* **36** 2328–39
- Cameron C 2005 Sweeping-window arc therapy: an implementation of rotational IMRT with automatic beam-weight calculation *Phys. Med. Biol.* **50** 4317–36
- Cao D, Afghan M, Ye J, Chen F and Shepard D 2009 A generalized inverse planning tool for volumetric-modulated arc therapy *Phys. Med. Biol.* **54** 6725–38
- Craft D, McQuaid D, Wala J, Chen W, Salari E and Bortfeld T 2011 Multicriteria VMAT optimization *Med. Phys.* **39** 686–96
- Crooks S, Wu X, Takita C, Watzich M and Xing L 2003 Aperture modulated arc therapy *Phys. Med. Biol.* **48** 1333–44
- Earl M, Shepard D, Naqvi S, Li X and Yu C 2003 Inverse planning for intensity-modulated arc therapy using direct aperture optimization *Phys. Med. Biol.* **48** 1075–89
- Gözbasi H 2010 Optimization approaches for planning external beam radiotherapy *PhD Thesis* Georgia Institute of Technology
- Gu X, Choi D, Men C, Pan H, Majumdar A and Jiang S 2009a GPU-based ultra-fast dose calculation using a finite size pencil beam model *Phys. Med. Biol.* **54** 6287–97
- Gu X, Jelen U, Li J, Jia X and Jiang S 2009b A GPU-based finite-size pencil beam algorithm with 3D-density correction for radiotherapy dose calculation *Phys. Med. Biol.* **56** 3337–50
- Luan S, Wang C, Cao D, Chen D, Shepard D and Yu C 2008 Leaf-sequencing for intensity-modulated arc therapy using graph algorithms *Med. Phys.* **35** 61–9
- Men C, Gu X, Choi D, Majumdar A, Zheng Z, Mueller K and Jiang S 2009 GPU-based ultra-fast IMRT plan optimization *Phys. Med. Biol.* **54** 6565–73
- Men C, Jia X and Jiang S 2010a GPU-based ultra-fast direct aperture optimization for online adaptive radiation therapy *Phys. Med. Biol.* **55** 4309–19
- Men C, Romeijn H, Jia X and Jiang S 2010b Ultra-fast treatment plan optimization for volumetric modulated arc therapy (VMAT) *Med. Phys.* **37** 5787–91
- Otto K 2008 Volumetric modulated arc therapy: IMRT in a single gantry arc *Med. Phys.* **35** 310–7

- Otto K 2009 Letter to the editor on 'single-arc IMRT?' *Phys. Med. Biol.* **54** L37–41
- Palma D, Vollans E, James K, Nakano S, Moiseenko V, Shaffer R, McKenzie M, Morris J and Otto K 2008 Volumetric modulated arc therapy for delivery of prostate radiotherapy: comparison with intensity-modulated radiotherapy and three dimensional conformal radiotherapy *Int. J. Radiat. Oncol. Biol. Phys.* **72** 996–1001
- Preciado-Walters F, Rardin R, Langer M and Thai V 2004 A coupled column generation, mixed integer approach to optimal planning of intensity modulated radiation therapy for cancer *Math. Program.* **101** 319–38
- Radiation Therapy Oncology Group 2004 H-0126: a phase iii randomized study of high dose 3D-CRT/IMRT versus standard dose 3D-CRT/IMRT in patients treated for localized prostate cancer <http://www.rtog.org/members/protocols/0126/p0126.pdf>
- Romeijn H, Ahuja R, Dempsey J and Kumar A 2005 A column generation approach to radiation therapy treatment planning using aperture modulation *SIAM J. Optim.* **15** 838–62
- Shepard D, Earl M, Li X, Naqvi S and Yu C 2002 Direct aperture optimization: a turnkey solution for step-and-shoot IMRT *Med. Phys.* **29** 1007–18
- Shepard D M, Cao D, Afghani M and Earl M 2007 An arc-sequencing algorithm for intensity modulated arc therapy *Med. Phys.* **34** 464–70
- Verbakel W, Cuijpers J, Hoffmans D, Bieker M, Slotman B and Senan S 2009 Volumetric intensity-modulated arc therapy versus conventional IMRT in head-and-neck cancer: a comparative planning and dosimetric study *Int. J. Radiat. Oncol. Biol. Phys.* **74** 252–9
- Wang C, Luan S, Tang G, Chen D, Earl M and Yu C 2008 Arc-modulated radiation therapy (AMRT): a single-arc form of intensity-modulated arc therapy *Phys. Med. Biol.* **53** 6291–303
- Yu C 1995 Intensity-modulated arc therapy with dynamic multileaf collimation: an alternative to tomotherapy *Phys. Med. Biol.* **40** 1435–49