



CIRRELT

Centre interuniversitaire de recherche
sur les réseaux d'entreprise, la logistique et le transport

Interuniversity Research Centre
on Enterprise Networks, Logistics and Transportation

Simultaneous Delivery Time and Aperture Shape Optimization for the Volumetric-Modulated Arc Therapy (VMAT) Treatment Planning Problem

Mehdi Mahnam
Michel Gendreau
Nadia Lahrichi
Louis-Martin Rousseau

February 2017

CIRRELT-2017-09

Bureaux de Montréal :
Université de Montréal
Pavillon André-Aisenstadt
C.P. 6128, succursale Centre-ville
Montréal (Québec)
Canada H3C 3J7
Téléphone : 514 343-7575
Télécopie : 514 343-7121

Bureaux de Québec :
Université Laval
Pavillon Palasis-Prince
2325, de la Terrasse, bureau 2642
Québec (Québec)
Canada G1V 0A6
Téléphone : 418 656-2073
Télécopie : 418 656-2624

www.cirrelt.ca

Simultaneous Delivery Time and Aperture Shape Optimization for the Volumetric-Modulated Arc Therapy (VMAT) Treatment Planning Problem

Mehdi Mahnam*, Michel Gendreau, Nadia Lahrichi, Louis-Martin Rousseau

Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT) and
Department of Mathematics and Industrial Engineering, Polytechnique Montréal, P.O. Box 6079, Station
Centre-Ville, Montréal, Canada H3C 3A7

Abstract. In this paper, we propose a novel heuristic algorithm for the volumetric-modulated arc therapy (VMAT) treatment planning problem, optimizing the trade-off between delivery time and treatment quality. We present a new mixed integer programming model in which the multileaf collimator leaf positions, gantry speed, and dose rate are determined simultaneously. Our heuristic is based on column generation; the aperture configuration is modeled in the columns and the dose distribution and time restriction in the rows. To reduce the number of voxels and increase the efficiency of the master model, we aggregate similar voxels using a clustering technique. The efficiency of the algorithm and the treatment quality are evaluated on a benchmark clinical prostate cancer case. The computational results show that a high-quality treatment is achievable in about 2.64 minutes using a four-thread CPU. Finally, we analyze the effects of the various parameters and two leaf-motion strategies.

Keywords: Radiation therapy treatment planning, VMAT, column generation, heuristic.

Acknowledgements. This project was supported by the Vanier Canada Graduate Scholarship program and the Natural Sciences and Engineering Research Council of Canada (NSERC), No. DVC134707. The authors would like to thank Prof. David Craft, Assistant Professor at Harvard Medical School, and Martin Hinse, Medical Physicist at Laval Integrated Cancer Centre (CICL), who helped us to provide and evaluate the case data.

Results and views expressed in this publication are the sole responsibility of the authors and do not necessarily reflect those of CIRRELT.

Les résultats et opinions contenus dans cette publication ne reflètent pas nécessairement la position du CIRRELT et n'engagent pas sa responsabilité.

* Corresponding author: mehdi.mahnam@cirrelt.ca

1. INTRODUCTION

Volumetric-modulated arc therapy (VMAT) is a novel form of external radiotherapy which incorporates rotation of the beam around the patient's body, while the beam is on. In this technology, the beam shape, rotation speed, and dose rate change continuously (Otto 2008). This minimizes the adverse effects of radiation to normal tissues and matches the radiation to the profile of the tumor. In general, the main advantages are a better dose distribution around the patient, lower radiation to normal tissues, faster treatment, and decreased patient discomfort (Yu & Tang 2011).

Similar to intensity modulated radiation therapy (IMRT), the beams are formed by a multi-leaf collimator (MLC). This equipment has a finite number of leaf pairs, each consisting of a leading and a trailing leaf. The treatment planning problem in VMAT involves selecting sequences of MLC shapes and determining the optimal dose rate and gantry rotation speed around the patient (Romeijn & Dempsey 2008).

An initial version of VMAT, intensity modulated arc therapy (IMAT), was proposed by Yu (1995) as a rotational alternative to the IMRT delivery technique in tomotherapy. Unfortunately, IMAT was not widely adopted because of its time inefficiency. Otto (2008) proposed a simple heuristic as the first study for VMAT treatment planning problem. Several optimization approaches have subsequently been developed for this large-scale problem; a comprehensive review is provided by Unkelbach et al. (2015).

The methods proposed to model MLC restrictions in VMAT treatment planning can be classified into two groups: (1) arc-based and (2) control-point-based. In the arc-based approach, leaf motions during an arc are determined using a traditional two-stage method. In the first stage, fluence map optimization (FMO), the profile of beams for all the sectors are determined. A deliverable arc-sequencing, typically based on unidirectional leaf motion, is then designed such that the dose distribution is similar to the ideal fluence map (Cao et al. 2009, Craft et al. 2012). Wala et al. (2012) improved this method by generating efficient partial-arc plans via an iterative heuristic called PMERGE. Salari et al. (2012) replaced the merging heuristic with an exact discrete bicriteria optimization. In the same line of research, Papp & Unkelbach (2014) proposed an arc-based model to deliver the ideal fluence map and directly optimize the unidirectional leaf trajectories. Although the treatment plan resulting from the arc-based approach is reliable, the required computation is very time-consuming.

The control-point-based approach considers a finite number of points around the patient, the so-called control points or sectors, and each is associated with an aperture. The apertures are designed taking into account the MLC constraints, and the model optimizes the corresponding intensity. For this approach, heuristics based on column generation (CG) have been developed. These algorithms sequentially generate new aperture shapes in the pricing subproblem (PSP) and add them to the master model, thus ensuring that each new aperture is compatible with the previous ones (Men et al. 2010, Peng et al. 2012). Moreover, to take advantages of GPU computing, effective parallel processing algorithm have been implemented. The goal of this approach is to develop online adaptive radiotherapy techniques that can handle the inter-fraction variation of the patient's geometry.

Recently, Peng et al. (2015) extended the control-point-based approach to the case with constant gantry speed and dose rate (VMATc) using a heuristic framework. VMATc is beneficial in developing countries without access to specialized treatment equipment which allows the parameters to vary continuously. Akartunali et al. (2015) proposed a unified mixed-integer programming model for VMAT, tomotherapy, and CyberKnife. Although the formulation was improved by considering polyhedral analysis and valid inequalities, this model is large, complex, and difficult to solve. The authors developed two Lagrangian relaxation (LR) heuristics, a centring-

based heuristic, and a guided variable neighborhood search; the last approach was the most efficient. Nguyen et al. (2016) developed a nongreedy approach using an L2-norm fidelity term in the objective and a level set function to optimize the fluence intensity and aperture shapes simultaneously. Recently, Balvert & Craft (2016) investigated the trade-off between delivery time and fluence map matching. They optimized the leaf trajectories and dose rates for a given delivery time and obtained the complete trade-off curve for delivery time and plan quality by solving the model heuristically with a sequence of delivery times.

Most models in the literature fix the key parameters—delivery time and dose rate—at different levels of the treatment planning. In this paper, we propose an efficient algorithm that simultaneously optimizes the aperture shapes, delivery time, and dose rate. The method incorporates the advantages of both arc-based and control-point-based approaches. We formulate a new mixed integer programming model. We form successive generations of apertures using graph algorithms for the sequencing and a gradient-based model with column-and-row generation to optimize the intensities. To increase the computational efficiency and also the treatment quality, we avoid solving the large 180-FMO optimization model, and evaluate the treatment plans directly using voxels rather than fluence maps. Our model could easily be used to restrict the number of arcs or to take into account multiple arcs. However, in this paper we focus on single-arc therapy to determine high-quality plans with a minimum delivery time. In summary, this paper makes the following contributions to VMAT treatment planning problem:

1. The new formulation simultaneously optimizes the gantry speed, dose rate, and leaf trajectories.
2. We optimize the total delivery time and the treatment quality. The gantry speed and consequently the delivery time at each sector around the patient are considered dynamic.
3. The algorithm is based on simultaneous column-and-row generation to relate the MLC apertures and the delivery time. Moreover, the feasibility of the movement constraints is taken into account in the PSP rather than the master problem, which increases the efficiency of the algorithm.
4. To decrease the computational time, we propose a new model of down-sampling based on the geometry of the voxels in the patient’s body and clustering algorithms from data mining. This method aggregates similar voxels while considering the effect of all the voxels in the objective function of the model.

The remainder of the paper is organized as follows. Section 2 introduces the basic notation and the VMAT formulation in our framework. In Section 3, we propose our CG-based heuristic including the master model, PSP, greedy heuristic, and post-optimization. Section 4 presents computational experiments and results for a prostate case. Finally, Section 5 provides concluding remarks and discusses future research.

2. VMAT TREATMENT PLANNING FORMULATION

To model the problem, we discretize each structure s of the patient into small cubic volume elements called *voxels*, \mathcal{V}_s , and denoted \mathcal{V}_T and \mathcal{V}_N for tumorous and normal tissues. Each beam is decomposed into a rectangular grid of *beamlets*, I . A beamlet is on if it is not blocked by either the leading or trailing leaf, and the relative motion of the leaves controls the intensity. The estimated dose received by voxel j from beamlet i at unit intensity is denoted D_{ij} , in Gy/MU. $D = [D_{ij}]$ is called the dose-influence matrix. We assume that the dose absorbed by a given voxel can be determined by adding the dose from each of the individual beamlets comprising the aperture. The dose received at voxel j from all the beamlets of aperture A is denoted $D_j(A)$.

In VMAT treatment planning, the continuous dose delivery is discretized over a finite number of sectors, H , by changing the angle of the beam, with typically 180 2° -spaced beam angles. This discretization is necessary because estimating the dose-influence matrix in each sector, D_{ij}^h , by dosimetry techniques is time-consuming. Moreover, each sector is associated with an aperture A regarding the position of the MLC trailing (left) and leading (right) leaves, the dose rate ρ (in MU/s), and the gantry speed v_g (in deg/s). We assume that the aperture, dose rate, gantry speed, and dose-influence matrices are constant for adjacent sectors. This approximation is accurate when the angular distance between the two sectors is small (Otto 2008, Peng et al. 2012). In addition, the gantry speed is the inverse of the sector time, and we use these terms interchangeably.

In this paper, the 360 degrees around the patient are covered by the set of arcs K , and each arc k of length δ^k contains a finite number of equispaced sectors, H^k , from h_s^k to h_f^k . Moreover, A_h^k denotes the aperture corresponding to sector h in arc k . To model the feasible treatment plans, we restrict the machine characteristics as follows:

- MLC constraints: The leading and trailing leaves are taken into account in the MLC ranges in the leaf motion constraints, and no overlap is allowed. Moreover, compatibility between adjacent MLC apertures, Δ_A , is required based on the maximum MLC leaf speed, v_l^U .
- Gantry speed constraints: We impose lower and upper bounds v_g^L and v_g^U on the gantry speed to guarantee that the machine is in motion and within the feasible range.
- Dose rate constraints: We assume that the dose rate is bounded above by R , and a zero dose rate is allowed to avoid useless apertures.
- Adjacency constraints: We require compatibility of adjacent sectors in terms of the dose rate and the sector time in addition to the aperture shape, Δ_A . Thus, we restrict the rate of change of the dose rate, Δ_ρ , and the sector time, Δ_t .

The model for the VMAT problem is as follows, where $y^k = 1$ if and only if arc k is selected in the solution:

$$\text{GP: } \min_{y, \rho, t} F(\mathbf{z}) + w T \quad (1)$$

$$\text{s.t. } z_j = \sum_{k \in K} \sum_{h \in H_k} D_{jh}(A_h^k) y^k \rho_h t_h \quad \forall j \in \mathcal{V} \quad (2)$$

$$\sum_{k \in K} a_h^k y^k \leq 1 \quad \forall h \in H \quad (3)$$

$$|\rho_{h+1} - \rho_h| \leq \Delta_\rho \quad \forall h \in H \setminus \{|H|\} \quad (4)$$

$$0 \leq \rho_h \leq R \quad \forall h \in H \quad (5)$$

$$\sum_{k \in K} \tau_{h,h+1}^k y^k \leq t_h \quad \forall h \in H \quad (6)$$

$$|t_{h+1} - t_h| \leq \Delta_t \quad \forall h \in H \setminus \{|H|\} \quad (7)$$

$$\underline{T} \leq t_h \leq \bar{T} \quad \forall h \in H \quad (8)$$

$$\sum_{h \in H} t_h \leq T \quad (9)$$

$$y^k \in \{0, 1\} \quad \forall k \in K \quad (10)$$

Here \mathbf{z} is the dose absorbed by the voxels, and $F(\mathbf{z})$ is a voxel-based convex function. T is the total delivery time, which is penalized by a small value of w . The objective is to maximize the treatment quality and minimize the delivery time. The dose absorbed by voxel j , z_j , is computed in Equation (2) based on the sectors covered by arc k , H_k ; the dose-influence matrix D ; the aperture shape A_h^k ; the arc selection variable y^k ; and the sector time t_h . Also, a_h^k indicates the coverage of sector h by arc k , and thus Constraint (3) ensures that each sector h is covered by at most one arc. Constraints (4, 5) enforce lower and upper bounds on the dose rate and limit the change in the

rate between adjacent sectors to at most Δ_ρ . Constraints (6) ensure that the time spent at each sector t_h is sufficient to allow the leaf movements to reach the aperture shape positions in the next sector, A_{h+1}^k , given the maximum leaf speed v_l^U . Constraint (7) limits the rate of change in the sector time between adjacent sectors to at most Δ_t . Finally, Constraint (8) enforces the lower and upper bounds on the delivery time, and Constraint (9) determines the total delivery time T .

Because of the nonlinearity in Equation (2) and the binary variable y , this is a large-scale mixed integer nonlinear programming (MINLP) model. It is difficult to solve the proposed MINLP model, especially in real-world cases even by commercial softwares. In the next section we propose an efficient heuristic to produce high-quality treatment plans.

3. PROPOSED METHOD

In the column generation (CG) technique, instead of generating all the feasible columns in advance, we decompose the model into two subproblems: (1) the restricted master problem (RMP), which contains a subset of the columns, and (2) the PSP, in which promising new columns are generated, expecting to improve the RMP objective function. In our model, each column k is an arc from sector h_s^k to h_f^k , made up of the apertures of the leading and trailing leaves in a sequence of sectors. Without loss of generality and to simplify the method, we assume that the number of arcs and their lengths are fixed. The procedure starts with a subset of initial arcs completely surrounding the patient. We then repeatedly solve the following two subproblems until no new improving column is available:

- RMP: Given a large set of arcs ($K' \subseteq K$), we select the best subset that covers all the sectors, optimizing the intensities, gantry speed, and dose distribution z^* . We transfer the dual values of the voxels to the PSP.
- PSP: Given the dual values from RMP, we identify promising arcs using graph theory and add them to the pool of generated arcs (K').

To see the benefits of CG in VMAT, consider an (n_r, n_c) beam and h sectors. The total number of aperture shapes is about $((n_r/2)(n_c + 1)(n_c + 2))^h$ while the optimal solution has only h apertures. In an instance with a $(5, 10)$ beam and 100 sectors, there are 7.1×10^{251} possible apertures. In real cases, the number of beamlets and sectors is usually much larger. Therefore, CG generates good apertures while avoiding those that are unhelpful, similar, or infeasible. Furthermore, CG is a linearization method that converts the nonlinear GP model into two linear models.

3.1. Master Model

The master model optimizes the dose rate and gantry speed during the LINAC rotation around the patient given the generated arcs. The total number of potential arcs K is large, so a subset $K' \subset K$ is considered in the RMP. This model is solved several times during the CG procedure, so the algorithm must be efficient. To decrease the computational time, we have developed a convex model with a convex cost function, linearized the solution space, and reduced the number of constraints by aggregating similar voxels.

3.1.1 Cost function

Several objective functions have been proposed in radiotherapy treatment planning, including quadratic objectives, nonlinear radiobiological objectives, e.g., TCP and NTCP, and the equivalent uniform dose (EUD) (Ehrgott et al. 2010). We use the following *quadratic voxel-based least square*

penalty function to measure the treatment quality:

$$F(\mathbf{z}) = \sum_{s \in S} \sum_{j \in V_s} \underline{w}_s [d_j - z_j]_+^2 + \overline{w}_s [z_j - \bar{d}_j]_+^2, \quad (11)$$

where $[\bullet]_+$ denotes $\max\{0, \bullet\}$ and \underline{d}_j and \bar{d}_j are prespecified lower and upper dose thresholds for voxel j . This objective is well known because of its convexity and the quality of the resulting treatment plans (Romeijn et al. 2005, Peng et al. 2012, Papp & Unkelbach 2014). Its disadvantage is the large number of constraints required to linearize the $[\bullet]_+$ operator for the relevant voxels.

3.1.2 Linearization

In model (GP), the dose received by each voxel j depends on the dose intensity and the time in the sectors, which is the so-called sector-dependent approach. This approach is nonlinear, and we make two simplifying assumptions to linearize the model. First, we replace the product of the decision variables ρ and t in Constraint (2) by a new variable called the *fluence rate*, defined to be $\gamma = \rho t$. We can determine the optimal dose rate for each sector by simultaneously optimizing the fluence rate and the sector time. Second, the nonlinear term $\gamma_h y^k$ gives the fluence rate of sector h in arc k if selected. To avoid this product, we assume a constant fluence rate through the arc, γ^k , as in the arc-dependent approach. Constraint (2) is thus replaced by Equation (12), where the dose absorbed by voxel j depends only on the fluence rate of arc k , γ^k , and aperture k , A^k :

$$z_j = \sum_{k \in K} D_{jh}(A_h^k) \gamma^k \quad \forall j \in V_s \quad (12)$$

The linear feasible space of the arc fluence rate, γ^k , with regard to the arc variable, y^k , the sector time t_h , the maximum dose rate R , and the linearizing relaxations (Adams & Sherali 1990) is as follows:

$$\gamma^k \leq R t_h \quad \forall k \in K, \forall h \in H^k \quad (13)$$

$$0 \leq \gamma^k \leq y^k R \bar{T} \quad \forall k \in K \quad (14)$$

$$\gamma^k \geq 0 \quad \forall k \in K \quad (15)$$

To make the model more efficient, we replace Constraints (13) by

$$\sum_{k \in K} a_h^k \gamma^k \leq R t_h \quad \forall h \in H \quad (16)$$

Although the fluence rate is fixed for any given arc, it can change between arcs. Moreover, the gantry speed, v_g^h , can change from one sector to another. Finally, when we have found the final arcs, we use a post-optimization model to determine the sector-based fluence rates (see Section 3.4). In addition, restricting the leap of fluence rate between adjacent sectors is taken into account in the post-optimization model.

The proposed master model is a mixed integer programming model with binary arc variables, y^k . At the beginning of the algorithm, we relax the binary variables in the relaxed restricted master problem (RRMP):

$$\text{RRMP : } \min_{y, \gamma, t} \quad \mathbf{F}(z) + w T \quad (17)$$

$$\text{s.t. } z_j = \sum_{k \in K'} D_{jh}(A_h^k) \gamma^k \quad \forall j \in V_s \quad (18)$$

$$\sum_{k \in K'} a_h^k y^k \leq 1 \quad \forall h \in H \quad (19)$$

$$\sum_{k \in K'} \tau_{h,h+1}^k y^k \leq t_h \quad \forall h \in H \quad (20)$$

$$\sum_{h \in H} t_h \leq T \quad (21)$$

$$\sum_{k \in K'} a_h^k \gamma^k \leq R t_h \quad \forall h \in H \quad (22)$$

$$\gamma^k \leq y^k R \bar{T} \quad \forall k \in K' \quad (23)$$

$$\underline{T} \leq t_h \leq \bar{T} \quad \forall h \in H \quad (24)$$

$$y^k, \gamma^k \geq 0, \quad \forall k \in K' \quad (25)$$

In this relaxation, the upper bound $y^k \leq 1$ is dropped because it is induced by (19). In addition to the quadratic objective function, the *column-dependent rows* in Constraints (23) and the large number of Constraints (18) are the two main challenges of this model. We discuss the first challenge in Section (3.1.3) and develop an objective function for the PSP based on an equivalent relaxation. To overcome the second challenge, we develop a down-sampling algorithm inspired by clustering methods to aggregate similar voxels (Section 3.1.4).

3.1.3 Column-dependent rows

In standard CG, the effect of a new column is computed by minimizing its *reduced cost* in the objective function of the PSP (Desaulniers et al. 2005):

$$\min_{k \in K \setminus K'} z^* = c_k - \pi^T a_k, \quad (26)$$

where c and a are the objective and constraint coefficients, and π is the vector of the dual values of the corresponding constraints in the master model. If $z^* < 0$, introducing the new column into the solution of the master problem would decrease the objective. However, in our model, there are two columns y^k and γ^k for each arc k , determining the selection value and fluence rate, respectively. These columns are linked by Constraint (23), and the linking constraints are not in the RRMP a priori; thus, the reduced costs of the columns may be computed incorrectly in the PSP because no dual information for the missing constraints is available.

General approaches for *simultaneous column-and-row generation* have been developed (Muter et al. 2013). However, our model has the following simplifying properties:

Property 1. For each arc, columns y and γ are linked through a single linking constraint.

Property 2. γ cannot take a positive value unless y is positive.

Property 3. Adding new linking constraints does not violate primal feasibility.

In this case, analyzing the reduced costs of columns y and γ separately would lead to incorrect results. In particular, column y , considered independently of γ , cannot have a negative reduced cost because it does not affect the dose distribution $F(z)$. Moreover, suppose that for a given arc, the reduced cost of γ is negative while the reduced cost of y is positive. Although column γ is eligible to enter the basis, it is forced to be zero and the resulting iteration will be degenerate because of Property (2). To avoid this, we must generate the two columns simultaneously together with the corresponding linking constraint (23).

In Appendix A we present a new model, RRMP', based on only column γ . It is shown that any optimal solution of RRMP can be transformed into an equivalent optimal solution in RRMP'.

Then, models RRMP and RRMP' are equivalent, and we can derive the objective function of the PSP from RRMP' in Section (3.2.2). To determine the best combination of arcs in the integer model, we generate column y and the linking constraint (23) based on column γ and add them to the model.

3.1.4 Voxel aggregation

As mentioned earlier, the master model is time-consuming because of the large number of voxels. In the literature, random down-sampling is often implemented to simplify the problem. In a new technique, Küfer et al. (2003) and Scherrer et al. (2005) clustered the voxels via heuristics based on hierarchy construction in IMRT optimization; they developed an acceptable approximation that could be solved ten times faster.

In this paper, we propose an efficient method inspired by the K-means algorithm, a well-known technique to classify data by fixing a certain number of clusters a priori (MacQueen 1967). We treat the voxels as points that must be clustered into $|C|$ groups. To define the set of observations, we consider full open radiation (open beamlets in all sectors, i.e., $\sum_{h \in H} D_h$). Each voxel is then a vector in R^n with n beamlets. We first randomly assign the voxels to the clusters and define the centroid of each cluster to be the mean of the observations. The original K-means method computes the distance of each voxel from the centroid of all the clusters and assigns it to the nearest one. We instead calculate the distance of the voxel from its *neighboring* clusters; this is because we expect a high degree of similarity among neighbors. We also take into account the average Euclidean distance of the voxels from the corresponding centroid as a measure of the quality of the clustering. This objective is nonincreasing, and the process continues until there is no significant improvement.

This method uses the geometric information about the voxels and the fact that neighboring voxels usually have similar dosimetry profiles. The weight of each cluster in the objective is proportional to the number of embedded voxels, so all the voxels are taken into account despite the down-sampling. Algorithm 1 gives the pseudocode for the aggregation method. We perform the aggregation at the beginning of the CG, and the sets of aggregated voxels are then fixed until the post-optimization.

3.2. Pricing subproblem

In the PSP, we use a graph approach to determine the arrangement of the leading and trailing leaves in the δ sectors from h_s to h_f . This approach was inspired by Boland et al. (2004) and Romeijn et al. (2005), who proposed a network model to produce the aperture shapes in IMRT, and also by Luan et al. (2008), who extended it to IMAT. In our model, each node represents a pair of MLC leaves in a sector, and the arcs represent feasible movements in the graph.

3.2.1 Graph algorithm

Let graph $G = (V, E)$ be a layered acyclic digraph where V is the set of nodes and E the set of arcs. Each node in the h^{th} layer (denoted by L_h) defines a feasible arrangement of MLC leaves in each row of sector h . The arrangement of trailing and leading leaves (t, l) in row m and sector h is denoted by $(h, t, l)_m$. Figure (1) shows example positions for a pair of leaves. Moreover, E contains the multi-layer arcs (u, v) where if $u \in L_h$ then $v \in L_{h+1}$. Our goal is to cover all the sectors in an arc, even possibly with closed apertures if required. We define two dummy nodes D and D' : the start node D is linked to the first layer, and the nodes of the last layer are linked to D' . Smaller partial arcs could easily be generated by linking D and D' to all the layers.

Algorithm 1 Voxel Aggregation

```

1: Input: Dose-influence matrices in all sectors  $D^h$ , number of clusters  $|C|$ 
2: Output: List of voxels in each cluster  $c^*$ 
3: Calculate  $D^F$  with full radiation:  $D^F = \sum_{h \in H} D^h$ 
4: Randomly assign voxels to clusters
5: while stop criteria not satisfied do
6:   for voxel  $j = 1$  to  $|V|$  do
7:      $\bar{d}_{j,\bar{c}_j}$  = distance from voxel  $j$  to its current cluster  $\bar{c}_j$ 
8:     for neighboring voxel  $j' = 1$  to  $|N_j|$  do
9:        $d_{j,\bar{c}_{j'}} = \|D_j^F - D_{\bar{c}_{j'}}^F\|_2^2$ 
10:      if  $d_{j,\bar{c}_{j'}} < d_{j,\bar{c}_j}$  then
11:         $\bar{c}_j^* = \bar{c}_{j'}$ 
12:      end if
13:    end for
14:    if  $\bar{c}_j^* \neq \bar{c}_j$  then
15:      Move voxel  $j$  to cluster  $c_j^*$ 
16:    end if
17:  end for
18: end while
19: stop

```




Figure 1: Leaf position with $(t, l) = (1, 4)$. The trailing (left) leaf is located in beamlet 1, and the leading (right) leaf is located in beamlet 4. The beamlets between these two are open and the others are blocked. Columns 0 and 5 are dummy blocked beamlets demonstrating the range of the leaf movement.

Our network takes into account the MLC constraints by eliminating infeasible nodes and arcs. As the first constraint, the range of leaf movements must be within certain limits. In another common constraint, collisions of trailing and leading leaves must be avoided, i.e., $t + 1 \leq r$. Moreover, leaf movements should be compatible between adjacent sectors, i.e., leaf movements between adjacent sectors are restricted based on the maximum leaf speed, v_l^U . We remove the infeasible arcs from the graph to guarantee the compatibility of adjacent sectors in the PSP. To consider the effect of the adjacency constraints, we transfer the minimum required time for the leaf movement at each sector h of arc k to the master model with the parameter τ_h^k . Figure (2) shows an instance of the pricing graph with three layers and maximum leaf movement $\Delta_l = 1$. An advantage of the graph approach is that increasing the number of technical constraints decreases the size of the graph.

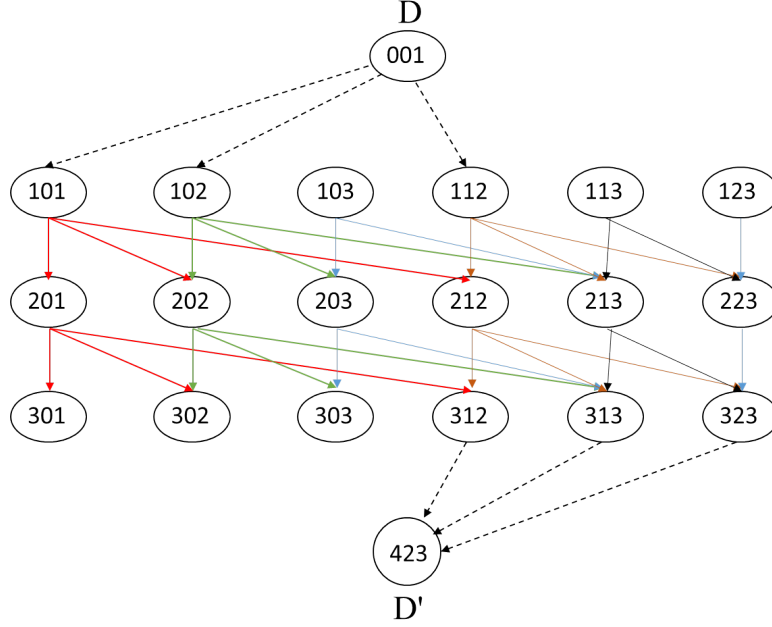


Figure 2: A pricing graph with three layers. We assume that the leaf movement is unidirectional from left to right and the maximum leaf movement is $\Delta_l = 1$ beamlet. Δ_l depends on the leaf and gantry speed parameters. For instance, if $v_l^U = 3$ cm/s, $v_g = 6^\circ/\text{s}$, $\delta = 2^\circ$, and the beamlet width $w_b = 1$ cm, Δ_l is 1. This value affects the arc feasibility during the building of the graph. For example, movement from node (101) to (213) is infeasible, and the corresponding arc is removed.

3.2.2 Pricing objective function

As explained in Section (3.1.3) and (A), the PSP is based on the reduced cost of RRMP':

$$\min_{k \in K \setminus K'} \left\{ \sum_{h \in H} \sum_{j \in \mathcal{V}} D_{jh}^k(A^k) \hat{\pi}_j - \sum_{h \in H} (\tau_{h,h+1}^k \hat{\phi}^h + a_h^k \hat{\sigma}_h) \right\}, \quad (27)$$

where $\hat{\pi}_j, \hat{\phi}_h \leq 0$, and $\hat{\sigma}_h \leq 0$ are the dual values of Equations (42), (43), and (44) in RRMP', respectively. Therefore, it suffices to solve the PSP with the objective function (27) and to add columns to the master model if the objective value is negative. To transfer this objective into the graph, we define the weight of each node $(h, t, l)_m$ as follows:

$$\sum_{j \in \mathcal{V}} D_{jh}^{(t,l)_m} (A^{(t,l)_m}) \hat{\pi}_j - \hat{\sigma}_h. \quad (28)$$

The weight of each arc connecting node $(h, t, l)_m$ in sector h to $(h+1, t', l')_m$ in sector $h+1$ is computed as follows:

$$-\tau_{(h,t,l)_m, (h+1,t',l')_m} \hat{\phi}^h, \quad (29)$$

where $\tau_{(h,t,l)_m, (h+1,t',l')_m}$ is the minimum time for vertex $(h, t, l)_m$ in sector h to reach $(h+1, t', l')_m$. Obviously, if a node is on the shortest path, $a_h^k = 1$ and the corresponding cost according to its sector and leaf position is taken into account; otherwise, the cost is not considered and $a_h^k = 0$.

3.2.3 Implementation details

The pricing subproblem, determining the best leaf pattern for a given range of sectors is solved by a polynomial-time shortest-path algorithm. According to the directed acyclic graph G in this problem, we use the topological sorting technique, which runs in $O(|V| + |E|)$ time. Although the run time is negligible, the time to refeed the graph weights regarding the large number of voxels is considerable. We solve the PSPs in parallel for new arcs to reduce the computational time. To this end, a large graph with 180 layers (sectors) is defined and a slave PSP for each new arc k is implemented. For instance, if the length of the arcs is four, the PSPs are PSP_1 (1–4), PSP_2 (5–8), \dots , PSP_{45} (177–180).

The independent arcs around the patient and the shared memory of the problem allow us to implement pricing in parallel using open multi-processing (openMP). This approach is easy to implement and flexible (Kiessling 2009). It starts with a single thread (the master thread) that creates sets of parallel worker threads (forks) in the parallel region. The threads are executed in parallel, and they synchronize and join the master thread at the end of the parallelization. In our work, the graph is shared among all the threads, and the threads simultaneously generate disjoint arcs in different parts of the graph.

3.3. Greedy Heuristic

The optimal solution for the arc selection variable y^k in the RRMP is usually fractional. The next step is 1) a branch-and-price search on the set of columns generated or 2) the use of a simple heuristic to round and fix the best columns. The first approach is too time-consuming, so we use a greedy heuristic to provide a feasible solution. At each CG iteration, we fix the lower bound of the y^k variable with the largest fractional part to 1. This method is fast but greedy. To accelerate the process, if there are other arcs with high fractional values, e.g., $y^k > 0.85$, we fix them also to 1 and skip the pricing problem for arcs that contain some or all of the corresponding sectors. Algorithm (2) gives the pseudocode for the greedy column-and-row generation heuristic (GCRGH).

Algorithm 2 GCRGH procedure

```

1: Initial random heuristic columns (arcs)
2: while stopping criteria not satisfied do
3:   Master Model (Arc-based)
4:   Fix the best arc to 1
5:   Pricing subproblem
6: end while
7: Post-Optimization (Sector-based)
8: Re-optimize gantry speed & intensities
9: stop

```

3.4. Post-Optimization

The proposed column generation-based heuristic returns the final aperture shapes and the arc-dependent fluence rates, γ^k . To improve the treatment plan, we propose a post-optimization model with three features. First, the arc-dependent approach is replaced by sector-dependent fluence rates. Second, we consider all the voxels and restrict the rate of change in the fluence rate, Δ_γ ,

and the sector time, Δ_t , between adjacent sectors. Third, the MLC leaf positions at the intersection points of adjacent arcs must satisfy the maximum leaf-movement constraint. If necessary, a simple heuristic removes any inconsistency with as little change as possible. Let the set of arcs found by the CG be K^* . The post-optimization model is as follows:

$$\text{Post-Opt : } \min_{\gamma, t} \quad \mathbf{F}(z) + w T \quad (30)$$

$$\text{s.t. } z_j = \sum_{h \in H} D_j^h(A^{K^*}) \gamma_h \quad \forall j \in V \quad (31)$$

$$\gamma_h \leq R t_h \quad \forall h \in H \quad (32)$$

$$\tau_{h,h+1} \leq t_h \leq \bar{T} \quad \forall h \in H \quad (33)$$

$$t_{h+1} - t_h \leq \Delta_t \quad \forall h \in H \setminus \{|H|\} \quad (34)$$

$$\gamma_{h+1} - \gamma_h \leq \Delta_\gamma \quad \forall h \in H \setminus \{|H|\} \quad (35)$$

$$\sum_{h \in H} t_h \leq T \quad (36)$$

$$t_h, \gamma_h \geq 0 \quad \forall h \in H \quad (37)$$

4. EXPERIMENTS AND RESULTS

We implemented the GCRGH algorithm in C++ and used IBM ILOG CPLEX 12.6.2 for the mathematical models. If not stated otherwise, the computational experiments were run on four threads of a 3.07 GHz Intel(R) Xeon(R) X5675 Linux workstation. The quadratic RRMP is solved at each iteration using the CPLEX Barrier Optimizer, which is based on an interior-point method. To evaluate the quality of the treatment plan, we use the dose volume histogram (DVH), the most common method in practice. These diagrams indicate the percentage of a volume that receives at least a certain dose. We export our plans to Computational Environment for Radiotherapy Research (CERR) to visualize the results and obtain DVHs (Deasy et al. 2003).

4.1. Data and benchmark

To evaluate the efficiency of the proposed heuristic, we considered a challenging prostate case in the CORT dataset, with 180 sectors, 25,404 beamlets, and 699,864 voxels (Craft et al. 2014). There are two tumors, PTV-56 and PTV-68, and seven healthy tissues including the bladder, rectum, left and right femoral heads, prostate bed, and penile bulb. The parameters are listed in Table (5). The fraction treatment is the standard 2 Gy per fraction, the prescribed radiation dose is 68 Gy for PTV-68, and the treatment plans were optimized for 34 fractions. Furthermore, the machine parameters of the VMAT problem, typically considered in the literature, are listed in Table (2) (Craft et al. 2014). In proportion to the range of the gantry speed, v_g , the delivery time is between 1 and 6 minutes.

Table 1: Prostate case parameters.

Total # beamlets	25,404
Beamlet size (cm)	1×1
Voxel resolution (cm)	3, 3, 3
# Target voxels	9491
# Body voxels	690,373

Table 2: VMAT machine parameters

Maximum leaf speed, v_l^U	3 cm/s
Gantry speed, v_g	1–6°/s
Maximum dose rate, ρ^U	600 MU/min
Maximum fluence change, Δ_γ	2 MU/s
Maximum change of sector time, Δ_t	2 s

4.2. Delivery time


In this section, we measure the performance of our GCRGH heuristic on the prostate case. The best delivery time found is 2.79 minutes (Plan-Opt). In Figure (3) we investigate the effect of delivery time on solution quality by comparing the DVH diagram for Plan-Opt to that for Plan-2 (maximum delivery time 2 minutes) and Plan-6 (maximum delivery time 6 minutes). Plan-2 has a common delivery time, and Plan-6 has the delivery time for a minimum 1°/s gantry speed. Table (3) compares the critical measures of these three plans. In Plan-Opt all the healthy tissues are far from the thresholds, and the tumor tissues are sufficiently irradiated. Similarly, the healthy tissues are protected in Plan-2, but the dose in the tumor tissues are substantially lower, especially for PTV-68. Plan-6 provides enough time for the treatment delivery, but its treatment plan is close to that for Plan-Opt: the mean absolute difference in the critical measures in Table (3) is about 1.44. Therefore, our algorithm finds a trade-off between the delivery time and the dose distribution.

A challenge in this case is *multi-organ voxels*: a voxel could be in more than one organ, especially at boundaries. This could introduce bias in the objective function and DVH diagrams. In these cases, if a voxel belongs to both tumor and healthy tissues, we penalize it twice in the objective function; if it is in two similar structures, we penalize the organ with higher priority. For example, if a voxel belongs to both PTV-68 and PTV-56, we choose PTV-68. Anyway the calculations for the treatment measures and DVH diagrams take all the tissues into account. As shown in Figure (3), the tail of the rectum and the bladder have some overlapping parts and tumors. The more accuracy in the boundaries and the fewer the multi-organ voxels, the greater the accuracy in the DVHs and plans.


The computational time is about 2.64 minutes, which compares well with the literature, especially since we use only four threads. Note that this does not include the time required to read the data and plot the diagrams.

Table 3: Treatment plan dosimetry measures

Organ	Measure	Plan-Opt	Plan-2	Plan-6
PTV-56	$V56 \geq 95\%$	95.83	97.75	97.59
PTV-68	$V68 \geq 95\%$	96.08	89.46	94.51
Rectum	$V30 \leq 80\%$	68.57	75.53	73.41
	$V50 \leq 50\%$	22.39	22.90	22.51
	$V65 \leq 25\%$	9.13	6.97	8.42
Bladder	$V40 \leq 70\%$	53.74	63.33	53.63
	$V65 \leq 50\%$	19.96	22.54	19.00



(a)



(b)

Figure 3: DVH plots for Plan-Opt (2.61 minutes), Plan-2 (2 minutes), and Plan-6 (6 minutes) obtained for the prostate case using GCRGH. (a) Comparison of Plan-Opt and Plan-2, and (b) comparison of Plan-Opt and Plan-6. The solid lines indicate Plan-Opt; the dashed lines indicate Plan-2 and Plan-6.

4.3. Arc length

In the proposed approach, we assume that the arcs are of equal length. There is a trade-off between long arcs with less computational time and short arcs with better treatment quality. To evaluate the impact of this parameter, we tested four plans with arcs of lengths 8, 12, 20, and 30 degrees; Figure (4) shows the results, and Table (4) gives the computational and delivery times. As expected, longer arcs reduce the computational time but also the treatment quality. Plans 20 and 30 are not acceptable because of the overdose in the normal tissue, especially in the rectum. Despite the higher computational time of Plan 8, there was no significant improvement in the treatment quality. Plan 12 offers the best trade-off between computational time and solution quality.

Table 4: Analysis with different arc lengths

Arc length ($^{\circ}$)	Delivery time (min)	Comp. time (min)
8	3.49	3.59
12	2.61	2.79
20	2.26	2.31
30	2.03	1.81

4.4. Leaf movement strategy

In this study, we assume that the MLC leaves movement are bidirectional, i.e., the leaves are free to move in any direction. However, some MLCs have only unidirectional leaf sequencing. In this technology, the leaves are aligned at one edge of the field at the beginning of the arc and at the opposite edge of the field at the end of the arc. In this section we compare these two approaches. Table (5) presents the results for arc lengths of 12° and 20° to evaluate these two strategies for small and large arc lengths.


The unidirectional approach has less leaf motion time; however, since the aperture opens and closes several times during the rotation, the treatment time is higher. With fewer arcs, the delivery time decreases. On the other hand, since there is only one direction and a smaller search space, the unidirectional approach is computationally faster. Figure (5) shows the DVH diagrams for the 12° arcs. The diagrams are similar, showing that both strategies are acceptable.

Table 5: Comparison of leaf-movement strategies for two arc lengths


Arc length ($^{\circ}$)	Leaf strategy	Delivery time (min)	Comp. time (min)
12	Bidirectional	2.79	2.64
12	Unidirectional	5.23	2.21
20	Bidirectional	2.31	2.19
20	Unidirectional	3.79	1.83

5. DISCUSSION AND CONCLUSION

In this paper, we have developed a simultaneous column-and-row generation algorithm embedded in a heuristic for VMAT treatment planning. The algorithm simultaneously optimizes the aperture shapes, dose rates, and gantry speeds in an arc-based approach. The problem is decomposed into



(a)



(b)

Figure 4: DVH curves for treatment plans with different arc lengths. (a) Comparison of lengths 8° (solid line) and 12° (dashed line), and (b) comparison of 20° (solid line) and 30° (dashed line).




Figure 5: Comparison of unidirectional and free leaf-movement strategies for 12° arc lengths.

the PSP, which generates possible aperture shapes, and the RMP, which optimizes the dose rate and delivery time. The proposed RMP is a large-scale integer programming model that minimizes the quadratic dose deviation from the prescribed lower and upper bounds. In addition to voxel-based dose distribution constraints, the columns are also linked to delivery time constraints. The heuristic repeatedly solves the relaxed restricted master problem and fixes the best arcs. In the PSP, graph algorithms form the aperture shapes around the patient in polynomial time. One of the advantages of the graph approach is that we can consider most of the MLC leaf-movement constraints in the PSP rather than the RMP, which makes the algorithm more efficient. This method can handle both bidirectional and unidirectional leaf motion. In addition, we can better evaluate the impact of new arcs on the voxel-based dose distribution calculation by using the pricing rather than an FM approach.

To reduce the number of voxel-based constraints in the master model, we propose an aggregation algorithm inspired by the K-means technique in clustering. This aggregation takes about 10 s; it greatly reduces the number of constraints with a negligible effect on the treatment quality.

Finally, we have evaluated our algorithm on a prostate cancer benchmark to demonstrate the efficiency of our model. We solved this benchmark in about 2.64 minutes using only four threads, whereas some approaches in the literature require GPU computing. Moreover, we analyzed the effect of unidirectional and bidirectional leaf movement.

The restricted master problem is the main bottleneck of the algorithm, requiring 60% of the computational time. To decrease this time and improve the voxel sampling, a dynamic voxel aggregation could be developed. The sampling would be more precise, fewer voxels would be required, and the computational time would decrease. Moreover, the columns generated at each iteration from different sectors are not disjoint and have similar effects; this could be investigated in future research. Another possibility would be to use GPU computing to accelerate the method.

ACKNOWLEDGEMENTS

This project was supported by the Vanier Canada Graduate Scholarship program and the Natural Sciences and Engineering Research Council of Canada (NSERC), No. DVC134707. The authors would like to thank Prof. David Craft, Assistant Professor at Harvard Medical School, and Martin Hinse, Medical Physicist at Laval Integrated Cancer Centre (CICL), who helped us to provide and evaluate the case data.

Appendix A. Equivalent RRMP' model

In this appendix, we show the RRMP' is an equivalent model for the RRMP with a single column γ . We use this model to provide a simultaneous reduced cost as the objective function of the PSP.

As mentioned in Section (3.1.3), there are two columns y and γ in the RRMP. Any optimal solution $(\tilde{y}, \tilde{\gamma}, \tilde{t})$ of the RRMP in which $\tilde{\gamma}^k < \tilde{y}^k R \bar{T}$ for at least one index k can be transformed into an equivalent optimal solution via

$$y^k = (R \bar{T})^{-1} \gamma^k \quad \forall k \in K'. \quad (38)$$

Therefore, we can consider the equivalent problem RRMP' in which Constraint (23) is removed and y^k is replaced by $\gamma^k / (R \bar{T})$ in (19) and (20):

$$\sum_{k \in K'} a_h^k \gamma^k \leq R \bar{T} \quad \forall h \in H \quad (39)$$

$$\sum_{k \in K'} \tau_{h,h+1}^k \gamma^k \leq R \bar{T} t_h \quad \forall h \in H \quad (40)$$

Constraint (39) is dominated by Equations (19) and (24). This leads to a model based only on column γ :

$$\text{RRMP': } \min_{\gamma, t} \quad \mathbf{F}(z) + w T \quad (41)$$

$$\text{s.t. } z_j = \sum_{k \in K'} D_{jh}(A_h^k) \gamma^k \quad \forall j \in V_s \quad (42)$$

$$\sum_{k \in K'} \tau_{h,h+1}^k \gamma^k \leq R \bar{T} t_h \quad \forall h \in H \quad (43)$$

$$\sum_{k \in K'} a_h^k \gamma^k \leq R t_h \quad \forall h \in H \quad (44)$$

$$\sum_{h \in H} t_h \leq T \quad (45)$$

$$\underline{T} \leq t_h \leq \bar{T} \quad \forall h \in H \quad (46)$$

$$\gamma^k \geq 0 \quad \forall k \in K' \quad (47)$$

References

- Adams, W. P. & Sherali, H. D. (1990), 'Linearization strategies for a class of zero-one mixed integer programming problems', *Operations Research* **38**(2), 217–226.
- Akartunali, K., Mak-Hau, V. & Tran, T. (2015), 'A unified mixed-integer programming model for simultaneous fluence weight and aperture optimization in VMAT, tomotherapy, and cyberknife', *Computers & Operations Research* **56**, 134–150.
- Balvert, M. & Craft, D. (2016), 'Fast approximate delivery of fluence maps: The VMAT case', *arXiv preprint arXiv:1608.03153*.

- Boland, N., Hamacher, H. W. & Lenzen, F. (2004), 'Minimizing beam-on time in cancer radiation treatment using multileaf collimators', *Networks* **43**(4), 226–240.
- Cao, D., Afghan, M. K., Ye, J., Chen, F. & Shepard, D. M. (2009), 'A generalized inverse planning tool for volumetric-modulated arc therapy', *Physics in Medicine and Biology* **54**(21), 6725.
- Craft, D., Bangert, M., Long, T., Papp, D. & Unkelbach, J. (2014), 'Shared data for intensity modulated radiation therapy (IMRT) optimization research: The CORT dataset', *GigaScience* **3**, 37–48.
- Craft, D., McQuaid, D., Wala, J., Chen, W., Salari, E. & Bortfeld, T. (2012), 'Multicriteria VMAT optimization', *Medical Physics* **39**(2), 686–696.
- Deasy, J. O., Blanco, A. I. & Clark, V. H. (2003), 'CERR: A computational environment for radiotherapy research', *Medical Physics* **30**(5), 979–985.
- Desaulniers, G., Desrosiers, J. & Solomon, M. M. (2005), *Column Generation*, Springer.
- Ehrgott, M., Guler, C., Hamacher, H. W. & Shao, L. (2010), 'Mathematical optimization in intensity modulated radiation therapy', *Annals of Operations Research* **175**(1), 309–365.
- Kiessling, A. (2009), An introduction to parallel programming with OpenMP, Technical report, The University of Edinburgh.
- Küfer, K.-H., Scherrer, A., Monz, M., Alonso, F., Trinka, H., Bortfeld, T. & Thieke, C. (2003), 'Intensity-modulated radiotherapy—a large scale multi-criteria programming problem', *OR Spectrum* **25**(2), 223–249.
- Luan, S., Wang, C., Cao, D., Chen, D. Z., Shepard, D. M. & Yu, C. X. (2008), 'Leaf-sequencing for intensity-modulated arc therapy using graph algorithms', *Medical Physics* **35**(1), 61–69.
- MacQueen, J. (1967), Some methods for classification and analysis of multivariate observations, in 'Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics', University of California Press, Berkeley, Calif., pp. 281–297.
- Men, C., Romeijn, H. E., Jia, X. & Jiang, S. B. (2010), 'Ultrafast treatment plan optimization for volumetric modulated arc therapy (VMAT)', *Medical Physics* **37**(11), 5787–5791.
- Muter, I., Birbil, Ş. İ. & Bülbül, K. (2013), 'Simultaneous column-and-row generation for large-scale linear programs with column-dependent-rows', *Mathematical Programming* **142**(1–2), 47–82.
- Nguyen, D., Lyu, Q., Ruan, D., O'Connor, D., Low, D. A. & Sheng, K. (2016), 'A comprehensive formulation for volumetric modulated arc therapy planning', *Medical Physics* **43**(7), 4263–4272.
- Otto, K. (2008), 'Volumetric modulated arc therapy: IMRT in a single gantry arc', *Medical Physics* **35**(1), 310–317.
- Papp, D. & Unkelbach, J. (2014), 'Direct leaf trajectory optimization for volumetric modulated arc therapy planning with sliding window delivery', *Medical Physics* **41**(1), 011701–01–011701–10.
- Peng, F., Jia, X., Gu, X., Epelman, M. A., Romeijn, H. E. & Jiang, S. B. (2012), 'A new column-generation-based algorithm for VMAT treatment plan optimization', *Physics in Medicine and Biology* **57**(14), 4569–88.

- Peng, F., Jiang, S. B., Romeijn, H. E. & Epelman, M. A. (2015), 'VMATc: VMAT with constant gantry speed and dose rate', *Physics in Medicine and Biology* **60**(7), 2955.
- Romeijn, H. & Dempsey, J. (2008), 'Intensity modulated radiation therapy treatment plan optimization', *TOP* **16**(2), 215–243.
- Romeijn, H. E., Ahuja, R. K., Dempsey, J. F. & Kumar, A. (2005), 'A column generation approach to radiation therapy treatment planning using aperture modulation', *SIAM Journal on Optimization* **15**(3), 838–862.
- Salari, E., Wala, J. & Craft, D. (2012), 'Exploring trade-offs between VMAT dose quality and delivery efficiency using a network optimization approach', *Physics in Medicine and Biology* **57**(17), 5587–5600.
- Scherrer, A., Küfer, K.-H., Bortfeld, T., Monz, M. & Alonso, F. (2005), 'IMRT planning on adaptive volume structures—a decisive reduction in computational complexity', *Physics in Medicine and Biology* **50**(9), 20–33.
- Unkelbach, J., Bortfeld, T., Craft, D., Alber, M., Bangert, M., Bokrantz, R., Chen, D., Li, R., Xing, L., Men, C. et al. (2015), 'Optimization approaches to volumetric modulated arc therapy planning', *Medical Physics* **42**(3), 1367–1377.
- Wala, J., Salari, E., Chen, W. & Craft, D. (2012), 'Optimal partial-arcs in VMAT treatment planning', *Physics in Medicine and Biology* **57**(18), 5861–5874.
- Yu, C. X. (1995), 'Intensity-modulated arc therapy with dynamic multileaf collimation: An alternative to tomotherapy', *Physics in Medicine and Biology* **40**, 1435–1449.
- Yu, C. X. & Tang, G. (2011), 'Intensity-modulated arc therapy: Principles, technologies and clinical implementation', *Physics in Medicine and Biology* **56**(5), R31–54.