Published in final edited form as:

Phys Med Biol. 2012 September 21; 57(18): 5861–5874. doi:10.1088/0031-9155/57/18/5861.

Optimal partial-arcs in VMAT treatment planning

Jeremiah Wala, Ehsan Salari, Wei Chen, and David Craft

Department of Radiation Oncology, Massachusetts General Hospital, 30 Fruit Street, Boston, MA 02114, USA

Abstract

Purpose—To improve the delivery efficiency of VMAT by extending the recently published VMAT treatment planning algorithm *VMERGE* to automatically generate optimal partial-arc plans.

Methods and materials—A high-quality initial plan is created by solving a *CONVEX* multicriteria optimization problem using 180 equi-spaced beams. This initial plan is used to form a set of dose constraints, and a set of partial-arc plans is created by searching the space of all possible partial-arc plans that satisfy these constraints. For each partial-arc, an iterative fluence map merging and sequencing algorithm (*VMERGE*) is used to improve the delivery efficiency. Merging continues as long as the dose quality is maintained above a user-defined threshold. The final plan is selected as the partial-arc with the lowest treatment time. The complete algorithm is called *PMERGE*.

Results—Partial-arc plans are created using *PMERGE* for a lung, liver and prostate case, with final treatment times of 127, 245 and 147 seconds. Treatment times using full arcs with *VMERGE* are 211, 357 and 178 seconds. The mean doses to the critical structures for the *VMERGE* and *PMERGE* plans are kept within 5% of those in the initial plan, and the target volume covered by the prescription isodose is maintained above 98% for the *PMERGE* and *VMERGE* plans. Additionally, we find that the angular distribution of fluence in the initial plans is predictive of the start and end angles of the optimal partial-arc.

Conclusions—The *PMERGE* algorithm is an extension to *VMERGE* that automatically finds the partial-arc plan that minimizes the treatment time. VMAT delivery efficiency can be improved by employing partial-arcs without compromising dose quality. Partial-arcs are most applicable to cases with non-centralized targets, where the time savings are greatest.

Keywords

VMAT; treatment planning; partial-arcs; delivery efficiency

1. Introduction

Volumetric modulated arc therapy (VMAT) is an extension to intensity modulated radiation therapy (IMRT) that offers the ability to deliver highly conformal dose distributions in a fraction of the time taken by IMRT (Yu & Tang 2011). In both techniques, the delivered radiation is modulated by a multi-leaf collimator (MLC) that conforms the dose to the targets. However, unlike IMRT where radiation is delivered at a fixed set of beam angles, in arc therapy the gantry delivers radiation continuously as it moves in an arc around the patient. In VMAT, one of the more widely used forms of arc therapy, the gantry speed, dose rate, and MLC leaf positions are allowed to vary as the gantry moves around the couch (Otto

jeremiah_wala@hms.harvard.edu.

‡D matrix dimensions – lung: 21979×14400; liver 41508×28800; prostate: 38760×21780. Critical structures (except the PTV) were sampled at a 1:4 ratio, and the unclassified tissue was sampled at 1:16.

2008). The additional degrees of freedom afforded by VMAT have been shown to maintain or improve the dose conformity when compared with IMRT (Teoh et al. 2011, Palma et al. 2008, Cozzi et al. 2008, Johnston et al. 2011). However, in VMAT the radiation is delivered through moving apertures that are constrained to lie within a limited distance of each other, which makes planning much more complex.

One of the most compelling reasons for using VMAT over IMRT is the ability to reduce the treatment time, often below 5 minutes per session (Bedford 2009, Clivio et al. 2009, Rao et al. 2010, Zhang et al. 2010). Reducing treatment time and the number of monitor units delivered reduces errors from patient motion (Hoogeman et al. 2008), decreases the scattered dose responsible for secondary malignancies (Hall & Wu 2003, Ruben et al. 2008, Kry et al. 2005), and may have direct radiobiological benefits (Wang et al. 2003, Bese et al. 2007). In addition, it reduces the amount of time patients need to spend receiving their therapy, and may allow for the treatment of more patients.

Approaches in reducing treatment time have focused on, for example, optimizing the delivery parameters (e.g. dose rate) given an initial plan (Rangaraj et al. 2010, Boylan et al. 2011), improving the delivery control system (Bertelsen et al. 2011) and developing leaf sequencing algorithms that improve delivery efficiency (Luan et al. 2008). In our own work on VMAT, we introduced *VMERGE*, a method to facilitate the selection of the best balance between planning objectives and delivery efficiency, one of the main tradeoffs in VMAT treatment planning (Craft et al. 2012). Inherent to most approaches is the idea that in the tradeoff between dose quality and delivery efficiency, dose quality should be highly favored.

Arc therapy can be delivered as a partial-arc, a full single arc, or as multiple arcs. For complex cases where a large amount of beam modulation is required, delivery in a single arc requires significant slowing down of the gantry rotation and reduction of the dose rate to achieve dose distributions similar to IMRT (Bortfeld & Webb 2009). Instead, radiation may be delivered over multiple arcs, which provides the necessary degrees of freedom but will tend to increase the total treatment time and required number of monitor units (Guckenberger et al. 2009).

Alternatively, for cases with peripherally located targets, one might expect that appreciable beam modulation is only required over a critical angle range. In this case, the gantry might move at maximal speed and with limited fluence over the angles where the radiation produces little or no beneficial effect. Additionally, because of the limit on the minimum dose rate, the beam cannot be entirely turned off over non-beneficial angles. Rotating the gantry through these angles may increase the dose to critical structures, as well as increase the total number of monitor units. Eliminating this portion of the arc from the delivery would be expected to reduce the treatment time and may decrease the dose to healthy tissue. However, to the best of our knowledge, no VMAT algorithm explicitly addresses this situation. In this work, we present a method called *PMERGE* for automatically finding the optimal partial-arcs that eliminate the unneeded arc segments without significantly compromising dose quality.

Like the beam angle optimization problem for fixed beam IMRT, the partial-arc optimization problem is nonconvex. Nonconvexity means that a local minimum is not guaranteed to be a global minimum. In order to guarantee a global minimum for nonconvex problems, a global search needs to be performed. Given that the nonconvex space of the partial-arc optimization problem consists of only two variables, the start and end gantry angles, this global search can be done by discretizing the search space and enumerating all possible start/end angle pairs, which is the strategy we adopt. This strategy is inherently parallelizable, allowing for shorter computation times.

To solve the VMAT problem for each partial-arc we use our previously reported merging and sequencing algorithm called VMERGE (Craft et al. 2012). We call the complete process PMERGE , which employs both merging and sequencing as in VMERGE and an enumerable set of partial-arcs to reduce treatment time. Given an initial high quality plan as a dose “gold standard”, PMERGE will automatically determine the fluence maps, dose rate, gantry speed, and starting and ending gantry angles that minimize the treatment time while maintaining the dose distribution to within a user-defined quality threshold. In this paradigm, all dose optimization goals are included in the first step with the creation of the initial plan. The single goal of the partial-arc routine is then to match this dose distribution with a partial-arc plan that provides the lowest possible treatment time.

2. Materials and methods

We define the following terms: *initial plan*, *partial-arc plan*, *planning goals*, VMERGE *plan* and PMERGE *plan*. The *initial plan* refers to a high-quality plan that serves as the dose standard used to constrain all partial-arcs. The initial plan is created without regards to delivery efficiency, providing a standard for understanding the tradeoff between dose quality and delivery efficiency. Our strategy for creating the initial plan is described in Section 2.2, where we define a 180-beam IMRT problem with beams placed every 2° , and solve it using multi-criteria optimization (MCO).

Next, a set of *partial-arc plans* is created using a two-step process. First, we solve a set of restricted angle IMRT problems with various start and end gantry angles by finding feasible solutions that minimize the sum-of-positive gradients (SPG), subject to matching the initial dose quality (Section 2.3). We then use VMERGE to iteratively merge neighboring fluence maps to create deliverable VMAT plans with improved delivery efficiency. Merging is automatically halted when the merged plan violates a set of user-defined *planning goals*, which articulates how much the planner is willing to deviate from the initial plan in order to improve delivery efficiency. We refer to the VMERGE plan as the result of the merging algorithm on the 180-beam plan, and the PMERGE plan as the merged partial-arc with the lowest treatment time. The user-defined planning goals, which serve as the halting criteria for merging, are used in generating both the VMERGE and PMERGE plans. The full method is outlined in Figure 1.

2.1. Delivery parameters and leaf sequencing

There are two classes of methods for solving VMAT problems: one-stage and two-stage (Yu & Tang 2011). In the one stage method, known as *direct aperture optimization*, the paths of the MLC leaves are directly optimized, taking into account the limitations on the maximum leaf speed. Instead, we use a two-stage planning approach to solve the VMAT problem for any size arc. In the first step, an N -beam IMRT problem with beams spaced every 2° is solved to generate a set of N fluence maps ($f_{l,x}^1, f_{l,x}^2, \dots, f_{l,x}^N$), where each fluence map is an $L \times X$ matrix of L rows of MLC leaf-pairs and X columns. In the second step, a leaf sequencing algorithm determines the MLC leaf trajectories, gantry speed (ω), and dose rate (r) for each 2° arc portion, which we call a *gantry sector*. In leaf sequencing, the trajectories of the leaves are set so that the total accumulated fluence over a given gantry sector matches the input fluence from the static beam fluence map. To accomplish this, we employ a widely used dynamic MLC (dMLC) sequencer which creates a “sliding window” that moves leaves unidirectionally across the field for each gantry sector (Spirou & Chui 1994, Svensson et al. 1994). Thus, for each gantry sector, the leaves must move from left-to-right (or right-to-left), traversing the entire width of the field (W). In our implementation of the dMLC sequencer, the field width for each gantry sector is constant and is defined as the distance the leaves

must travel across the union of all of the fluence maps. The dose is then calculated from the apertures defined by the leaf positions, as discussed in Section 2.2.

In practice, the gantry speed, dose rate and maximum leaf speed are constrained by the limitations of the linac hardware, and will have a significant effect on the total treatment time (Boylan et al. 2011). We select a set of hardware capabilities (Table 1) that are representative of the values found on many current treatment machines.

When applied to VMAT, the total treatment time using the dMLC algorithm with a fixed dose-rate is controlled by three terms: 1) the total distance the leaves must travel during delivery, 2) the amount of beam modulation required to deliver the desired fluence map and 3) the maximum gantry speed. These three considerations can be combined to provide an equation to calculate the total treatment time for a plan with N gantry sectors of arc lengths $\Delta\theta_i$:

$$T = \sum_{i=1}^N \max \left[\frac{W}{v_{\max}} + \frac{\max_{\text{rows } l} \left[\sum_{x=1}^X \left(\frac{d f_{l,x}^i(x)}{dx} \right)^+ \right]}{r}, \frac{\Delta\theta_i}{\omega_{\max}} \right] \quad (1)$$

where the $(.)^+$ operator is shorthand for $\max(0, .)$. The first term in (1) is the field width divided by the maximum leaf speed, which is the minimum time required for the leaves to move from the left side of the field to the right (or vice versa) if no fluence is delivered. As fluence is added, the leaf speed will be modulated to allow enough time for the dose to accumulate. This time is described in the second term as the sum-of-positive gradients (SPG), a measure of the amount of variation in a field (Craft et al. 2007), divided by the dose rate. Fields that are highly varied (have a high SPG term) will take longer to deliver with dMLC. Finally, because of limitations on the maximum gantry speed, the time required to deliver a gantry sector cannot be less than the time required to move the gantry at maximum speed over the gantry sector. The total treatment time is the sum of the time required to deliver each of the N gantry sectors.

Given a fixed set of hardware capabilities, the total treatment time can be improved by either decreasing the total number of gantry sectors N or by decreasing the SPG. Decreasing the number of gantry sectors motivates the partial-arc strategy and the merging strategy $_{\text{VMERGE}}$ - by using partial-arc solutions, we begin with a lower number of fluence maps than in a full arc, and this is decreased further by combining gantry sectors in the merging routine. Thus, for each gantry sector that is merged or eliminated with a partial-arc, the treatment time will necessarily decrease because the leaves have to make one fewer trip across the field. The rare exception to this is when the gantry is operating at the maximum speed across neighboring gantry sectors, in which case there is no benefit to merging. To further reduce treatment time, we incorporate SPG minimization into our optimization step as described in Section 2.3.

2.2. Initial plan optimization

We use a multicriteria optimization (MCO) approach to solve the initial 180-beam IMRT problem, as described in Craft et al. (2012). Briefly, in MCO each structure is assigned an objective function, and a tradeoff surface of Pareto optimal plans is generated by the optimizer (Monz et al. 2008, Thieke et al. 2007). A Pareto optimal plan is one where no improvements could be made to any objective without worsening another objective (Cotrutz et al. 2001). The user navigates this tradeoff surface and selects the plan that best meets the treatment goals.

The general multicriteria IMRT formulation is the following:

$$\begin{aligned} & \text{optimize} && \{g_1(d), g_2(d), \dots, g_M(d)\} \\ & \text{subject to} && d = D_f \\ & && d \in C \\ & && f \geq 0 \end{aligned} \quad (2)$$

where g_i are the M objective functions describing the dose to a particular structure (e.g., mean dose to stomach), d is the vector of voxel doses, D is the dose-influence matrix, and f is a concatenation of all the fluence maps into a single beamlet fluence vector. The set C is a convex set of dose constraints (e.g. maximum dose to the target) that is designed to be met by all plans. We use the CERR 3.0 (Deasy et al. 2003) environment and its beamlet based quadratic infinite beam (QIB) dose computation to calculate the dose influence matrix, with a computation time of around 10 minutes per case[‡].

Dose calculations are performed as previously described for the v_{MERGE} algorithm (Craft et al. 2012). Briefly, following each merge the dMLC sequencer is applied and the fluence is calculated from the apertures created by the MLC leaves. The accumulated fluence is binned every 2° to create the f vector. Given the D matrix and the f vector, the dose distribution can be calculated in seconds.

The constraints on the planning target volumes (PTV) and the OARs for each site are listed in Table 2. In our MCO formulation, we specify that the optimizer should minimize the mean doses to the OARs. Additionally, we include as an OAR unclassified tissue (u.t.), which encompasses all voxels that do not belong to another structure. Optimizations are performed using the solver described by Chen et al. (2010). We then choose a plan for each treatment site by exploring the Pareto surface and selecting a plan that balances the dose to the OARs.

2.3. Generating the partial-arc plans

The set of partial-arc solutions is created by iteratively solving the 2° grid IMRT problem for a range of partial-arcs. In creating the plans, we respect the constraint that the gantry cannot rotate through 0° , a line from the couch to the floor. Plans are created with arc lengths every 40° ($360^\circ, 320^\circ, \dots$) until there are no feasible solutions. For each arc length, plans are created with arcs centered at intervals of 40° . Thus, there is one 360° plan, three 320° degree plans (beginning at $-40^\circ, 0^\circ, 40^\circ$), five 280° plans, etc.

Each partial-arc solution is created to maintain the same dose quality as the initial 180-beam solution. This is done with the same solver used to solve the original 180-beam MCO formulation, except run in feasibility mode. Here, the solver returns the first solution it finds which satisfies all of the hard constraints. Because treatment time scales with the complexity of the solution, we seek to minimize the SPG while maintaining dosimetric feasibility. Our formulation of the SPG minimization problem with dose constraints is the following:

$$\begin{aligned} & \text{minimize} && \sum_{i=1}^N \max_{\text{rows } l} \left[\sum_{x=1}^K \left(\frac{df_{lx}^l(x)}{dx} \right)^+ \right] \\ & \text{subject to} && d = D_f \\ & && d_i \geq B_L, \forall i \in \text{target} \\ & && d_i \leq B_U, \forall i \\ & && \text{mean}_j^{(P)}(d) \leq \text{mean}_j^{(I)}(d), \forall \text{OAR}_j \\ & && f \geq 0 \end{aligned} \quad (3)$$

where B_L and B_U is the target lower bound, B_U is the upper bound on the dose, and $\text{mean}_j^{(P)}(d)$ and $\text{mean}_j^{(I)}(d)$ are the mean doses to the j th OAR for the partial-arc and initial plans, respectively. This formulation ensures that for each partial-arc solution the target coverage is no less than, and the maximum dose and the mean doses to each OAR are no greater than, those of the initial plan. The above optimization problem is solved by progressively tightening the constraint on the total SPG until a feasible solution can no longer be found.

2.4. Merging and final plan selection

After the SPG minimization step, we have a set of IMRT solutions for various partial-arcs. To improve the delivery efficiency, we use a previously described fluence map merging strategy called vMERGE to combine gantry sectors (Craft et al. 2012). Briefly, vMERGE employs a greedy merging strategy to iteratively add similar neighboring fluence maps - for each iteration, the two neighboring fluence maps that are most similar are added together into a single fluence map. This reduces the number of gantry sectors and improves the delivery efficiency, at the expense of reducing the amount of fluence modulation and the angular resolution.

pMERGE automates the merging routine by utilizing a quality cutoff, which limits how much the dose can change and automatically stops merging when the quality drops below the planning goals. Here, the planner is free to choose how much the dose quality is allowed to change in order to improve delivery efficiency. We use the initial high-quality plan as a guide for how to set these goals. For the OARs that receive the most radiation, we limit the partial-arc mean dose to less than 105% of the initial mean doses for those OARs. Additionally, we ensure that the target volume covered by the prescription isodose stays above 98%, and the maximum dose is less than 102% of the initial maximum dose. Finally, we illustrate the use of site-specific planning goals, including limiting the maximum dose to the femoral heads in the prostate case and adding additional constraints on the left lung dose in the lung case. The full sets of planning goals are provided in Tables 3–5.

3. Results

3.1. Lung

We navigated to an initial 180-beam plan with mean doses and target coverage as shown in Table 3. The treatment time of the initial plan before merging was 771 s, and was reduced to 211 s with vMERGE . The results of the partial-arc routine for the lung case are shown in Figure 2. The vMERGE plan is shown with the outer ring of Figure 2(c). The vMERGE plan is notable for a large gantry sector of 116° on the contralateral side of the PTV where the gantry speed is at its maximum. As shown in Figure 2(d), there is little fluence delivered over this section in the initial 180-beam plan.

Partial-arc solutions were attempted for arcs with lengths between 80° and 360° , at intervals of 40° . There were 50 feasible solutions that met the dose constraints from Equation 3, and 26 of these had lower treatment times than the vMERGE plan. The final pMERGE plan had an arc length of 200° and treatment time of 127 s, and is shown with the inner ring of Figure 2(c) after merging. The total difference in treatment times between the vMERGE (211 s) and pMERGE (127 s) plans was 84 s, a 40% reduction. Optimization, merging and sequencing took about 40 minutes to complete all 64 plans, running in parallel on 8 processors. Computation time was similar across the lung, liver and prostate cases.

3.2. Liver

We navigated to an original 180-beam solution with mean OAR doses, target coverage and maximum dose as shown in Table 4. The results of the partial-arc routine for the liver case are shown in Figure 3. The treatment time before merging for the full arc solution was 1214 s and was reduced to 357 s with v_{MERGE} (Figure 3(c), outer ring). Partial-arc solutions were attempted for arcs lengths between 80° and 360°, at intervals of 40°, and 33 had feasible solutions. There were no solutions with arc lengths less than 160°. Of the feasible partial-arcs, 19 had lower treatment times than the v_{MERGE} plan.

The selected partial-arc had an arc length of 240° (Figure 3(c), inner ring) and treatment time of 245 s. The total difference in treatment times between the v_{MERGE} plan (357 s) and p_{MERGE} plan (245 s) was 112 s, a 31% reduction. In this case, the initial 180-beam solution was significant for two peaks of fluence, one at approximately 150° and one at 330°, as shown in Figure 3(d). The optimal partial-arc includes both of these peaks, and was able to eliminate the first 100° where the amount of required fluence was low.

3.3. Prostate

We navigated to an original 180-beam solution with mean OAR doses, target coverage and maximum dose as shown in Table 5. The results of the partial-arc routine for the prostate case are shown in Figure 4. The treatment time before merging for the full arc solution was 860 s and was reduced to 178 s with v_{MERGE} (Figure 4(c), outer ring). Unlike the lung and liver cases, the v_{MERGE} plan had regularly spaced gantry sectors of similar size, and the fluence was delivered fairly uniformly around the target as shown in Figure 4(d). This was reflected in the optimal partial-arc solution, which was highly similar to the initial plan through nearly the full arc.

Partial-arc solutions were created for arcs lengths between 80° and 360°, at intervals of 40°, and 45 had feasible solutions. Of the 45, only 3 had treatment times lower than the v_{MERGE} plan. There were no feasible partial-arcs at 80°. For plans with arc lengths of less than 280°, almost no merging was possible before the plan quality degraded beyond the stopping criteria. The selected partial-arc had an arc length of 300° (Figure 4(c), inner ring) and treatment time of 147 s. The total difference in treatment times between the v_{MERGE} (178 s) plan and p_{MERGE} plan (147 s) was 31 s, a 17% reduction.

4. Discussion and conclusions

We have presented a method for automatically creating efficient VMAT partial-arc plans from a high quality reference plan. By sampling a set of partial-arcs, we reduced treatment times for a lung, liver and prostate case. The final p_{MERGE} plans reflected the underlying geometry of the problem. For the lung and liver cases with non-centralized PTVs, a significant portion of the full-arc was eliminated with little effect on the dose quality. This allowed for the treatment time to be reduced by 40% and 31% respectively. The reduction in treatment time for the prostate case was more modest at 17%. As might be expected for a centrally located target, the fluence was distributed more evenly around the arc in the prostate case, and most of the partial-arcs were not able to improve the delivery efficiency while maintaining the dose quality to within the specified boundaries.

In p_{MERGE} , the dose quality of the final partial-arc is designed to be highly similar to the initial plan. However, the p_{MERGE} plan is limited only by the planning goals specified after the initial plan creation. Thus, the partial-arc plan could deviate from the initial plan in ways that are not specified by this quality threshold. Any number or type of constraints could easily be imposed on the final solution to ensure that the initial plan and p_{MERGE} plan were sufficiently similar. In practice, we did not find that it was necessary to impose many constraints in

order to match the solutions. Although we were free to use any number of dose-volume points or EUD functions, our planning goals were restricted to only one or two constraints per OAR, usually a limit on the mean or maximum dose. For all three cases, these were sufficient to closely match the doses across the DVH, even for OARs which were not explicitly specified.

In VMAT, as in IMRT, there are many different ways to deliver the fluence that will give highly similar dose distributions. Because the treatment time is dependent on the geometry of these fluence maps, the method by which the initial fluence maps are created is an important determinant of treatment time. Although PMERGE is not confined to any single optimization strategy, we found that using SPG smoothing was critically important to achieving highly efficient plans. Indeed, fluence map smoothing is important in IMRT (e.g. Matuszak et al. (2008)), where it is known that the complexity of the fluence maps – and hence the number of monitor units – can be greatly reduced without significantly impacting the dose distribution (Giorgia et al. 2007).

Common commercial hardware implementations of VMAT are not implemented with sliding window delivery techniques as we propose herein. This is likely due to the larger number of control points needed to specify sliding window plans, but also because the commercial optimization algorithms have tended to direct aperture optimization rather than the two-step fluence map and sequencing approach. Although sliding window delivery typically has smaller fields and can also take longer, due to the finite leaf speed, the problems that arise from this – more challenging dosimetry (Crop et al. 2007, Calvo et al. 2012) and larger delivery times with more control points – are surmountable, and in our opinion worth tackling due to the promise of superior dose distributions achieved by exact matching of optimal fluence maps (Calvo et al. 2012)

Although the partial-arc global search can be parallelized, we recognize that the high computation time is a limit to our method. We found that an arc sampling angle of 40° was sufficient for these three cases, and finer sampling provided little benefit for the additional computation time. However, for more difficult cases with a complex arrangement of critical structures, finer sampling may be required to achieve the maximum possible time savings.

The observation that the optimizer distributes fluence very similarly between the full-arc and partial-arc plans suggests a way of approximating the appropriate partial-arc without having to enumerate all of the possibilities. One strategy would be to plot the arrangement of fluence around the arc and have the planner select the partial-arc that captures the critical angles. For the lung and liver cases, there were several different partial-arcs that greatly improved the treatment time, and multiple plans near the global minimum were almost equivalent to the optimal plan (see Figures 2(a) and 3(a)). More importantly, we conclude that partial-arcs that cover certain critical angles are useful for reducing treatment time in VMAT, and PMERGE is a useful first-pass method at finding these optimal plans.

Acknowledgments

The project described was supported by Award Number R01CA103904 from the National Cancer Institute. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Cancer Institute or the National Institutes of Health. The authors would like to thank Dualta McQuaid for his contributions to the computational infrastructure used in this research.

References

- Bedford J. Medical Physics. 2009; 36:5128–5138. [PubMed: 19994523]
- Bertelsen A, Lorenzen E, Brink C. Medical Physics. 2011; 38:4802–4810. [PubMed: 21928653]

- Bese N, Hendry J, Jeremic B. *International Journal of Radiation Oncology Biology Physics*. 2007; 68(3):654–661.
- Bortfeld T, Webb S. *Physics in Medicine and Biology*. 2009; 54:N9–N20. [PubMed: 19075362]
- Boylan C, Rowbottom C, Mackay R. *Physics in Medicine and Biology*. 2011; 56:4119–4133. [PubMed: 21693790]
- Calvo O, Gutiérrez A, Stathakis S, Esquivel C, Papanikolaou N. *Journal of Applied Clinical Medical Physics*. 2012; 13(3)
- Chen W, Craft D, Madden T, Zhang K, Kooy H, Herman G. *Medical Physics*. 2010; 37:4938–4945. [PubMed: 20964213]
- Clivio A, Fogliata A, Franzetti-Pellanda A, Nicolini G, Vanetti E, Wytenbach R, Cozzi L. *Radiotherapy and Oncology*. 2009; 92(1):118–124. [PubMed: 19181409]
- Cotrutz C, Lahanas M, Kappas C, Baltas D. *Physics in Medicine and Biology*. 2001; 46:2161. [PubMed: 11512617]
- Cozzi L, Dinshaw K, Shrivastava S, Mahantshetty U, Engineer R, Deshpande D, Jamema S, Vanetti E, Clivio A, Nicolini G, et al. *Radiotherapy and Oncology*. 2008; 89(2):180–191. [PubMed: 18692929]
- Craft D, McQuaid D, Wala J, Chen W, Salari E, Bortfeld T. *Medical Physics*. 2012; 39(2):686–696. [PubMed: 22320778]
- Craft D, Süß P, Bortfeld T. *International Journal of Radiation Oncology Biology Physics*. 2007; 67(5): 1596–1605.
- Crop F, Reynaert N, Pittomvils G, Paelinck L, De Gersem W, De Wagter C, Vakaet L, De Neve W, Thierens H. *Physics in medicine and biology*. 2007; 52:3275. [PubMed: 17505102]
- Deasy J, Blanco A, Clark V. *Medical Physics*. 2003; 30:979–985. [PubMed: 12773007]
- Giorgia N, Antonella F, Eugenio V, Alessandro C, Filippo A, Luca C. *Radiat Oncol*. 2007; 2:42. [PubMed: 18036217]
- Guckenberger, M.; Richter, A.; Krieger, T.; Wilbert, J.; Baier, K.; Flentje, M. *Radiotherapy and Oncology*. Vol. 93. 2: 2009. p. 259-265.
- Hall E, Wu C. *International Journal of Radiation Oncology Biology Physics*. 2003; 56(1):83–88.
- Hoogeman M, Nuytens J, Levendag P, Heijmen B. *International Journal of Radiation Oncology Biology Physics*. 2008; 70(2):609–618.
- Johnston M, Clifford S, Bromley R, Back M, Oliver L, Eade T. *Clinical Oncology*. 2011; 23:503–511. [PubMed: 21397477]
- Kry S, Salehpour M, Followill D, Stovall M, Kuban D, White R, Rosen I. *International Journal of Radiation Oncology Biology Physics*. 2005; 62(4):1195–1203.
- Luan S, Wang C, Cao D, Chen D, Shepard D, Yu C, et al. *Medical Physics*. 2008; 35(1):61–69. [PubMed: 18293562]
- Matuszak M, Larsen E, Jee K, McShan D, Fraass B. *Medical Physics*. 2008; 35:1532–1546. [PubMed: 18491548]
- Monz M, Küfer K, Bortfeld T, Thieke C. *Physics in Medicine and Biology*. 2008; 53:985–998. [PubMed: 18263953]
- Otto K. *Medical Physics*. 2008; 35:310–317. [PubMed: 18293586]
- Palma D, Vollans E, James K, Nakano S, Moiseenko V, Shaffer R, McKenzie M, Morris J, Otto K. *International Journal of Radiation Oncology Biology Physics*. 2008; 72(4):996–1001.
- Rangaraj D, Odiraju S, Sun B, Santanam L, Yang D, Goddu S, Papiez L. *Medical Physics*. 2010; 37:4056–4067. [PubMed: 20879567]
- Rao M, Yang W, Chen F, Sheng K, Ye J, Mehta V, Shepard D, Cao D. *Medical Physics*. 2010; 37:1350–1359. [PubMed: 20384272]
- Ruben J, Davis S, Evans C, Jones P, Gagliardi F, Haynes M, Hunter A. *International Journal of Radiation Oncology Biology Physics*. 2008; 70(5):1530–1536.
- Spirou S, Chui C. *Medical Physics*. 1994; 21:1031–1041. [PubMed: 7968833]
- Svensson R, Kallman P, Brahme A. *Physics in Medicine and Biology*. 1994; 39:37–61. [PubMed: 7651999]

- Teoh M, Clark C, Wood K, Whitaker S, Nisbet A. *British Journal of Radiology*. 2011; 84(1007):967–996. [PubMed: 22011829]
- Thieke C, Küfer K, Monz M, Scherrer A, Alonso F, Oelfke U, Huber P, Debus J, Bortfeld T. *Radiotherapy and Oncology*. 2007; 85(2):292–298. [PubMed: 17892901]
- Wang J, Li X, D'Souza W, Stewart R. *International Journal of Radiation Oncology Biology Physics*. 2003; 57(2):543–552.
- Yu C, Tang G. *Physics in Medicine and Biology*. 2011; 56:R31–R54. [PubMed: 21297245]
- Zhang P, Happersett L, Hunt M, Jackson A, Zelefsky M, Mageras G. *International Journal of Radiation Oncology Biology Physics*. 2010; 76(5):1456–1462.

Figure 1.

Diagram of our three step process. 1) An initial IMRT plan is created to generate a set of dose constraints for $PMERGE$ and $VMERGE$. 2a) $PMERGE$ uses these constraints to generate a set of partial-arc solutions that matches the initial dose distribution. 2b) Each N -beam partial-arc solution is merged and sequenced, subject to the user-defined planning goals. 3) The final plan is selected as the one with the lowest treatment time.

Figure 2.

Lung: a) Final treatment times for each of the merged partial-arc plans (blue squares) with treatment times near or below the v_{MERGE} plan (red circle). b) DVH for the initial 180-beam plan (solid) and the p_{MERGE} plan (dashed). c) Initial and merged gantry sectors for the p_{MERGE} plan (inner rings) and 180-beam plan (outer rings) over the p_{MERGE} dose wash. d) Distribution of fluence for the initial solution and the optimal partial-arc before merging (100-beam).

Figure 3.

Liver: a) Final treatment times for each of the merged partial-arc plans (blue squares) with treatment times near or below the v_{MERGE} plan (red circle). b) DVH for the initial 180-beam plan (solid) and the p_{MERGE} plan (dashed). c) Initial and merged gantry sectors for the p_{MERGE} plan (inner rings) and 180-beam plan (outer rings) over the p_{MERGE} dose wash. d) Distribution of fluence for the initial solution and the optimal partial-arc before merging (120-beam).

Figure 4.

Prostate: a) Final treatment times for each of the merged partial-arc plans (blue squares) with treatment times near or below the v_{MERGE} plan (red circle). b) DVH for the initial 180-beam plan (solid) and the p_{MERGE} plan (dashed). c) Initial and merged gantry sectors for the p_{MERGE} plan (inner rings) and 180-beam plan (outer rings) over the p_{MERGE} dose wash. d) Distribution of fluence for the initial solution and the optimal partial-arc before merging (140-beam).

Table 1

VMAT delivery parameters and constraints

Delivery parameter	Symbol	Constraint
Maximum leaf speed	$v_{\max}$	2.5 cm/s
Maximum gantry speed	$\omega_{\max}$	6 deg/s
Minimum gantry speed	$\omega_{\min}$	0 deg/s
Maximum dose rate	$r_{\max}$	600 MU/min
Minimum dose rate	$r_{\min}$	50 MU/min

Table 2

Target coverage constraints and OARs used in the MCO optimization

Treatment site	PTV Min; Max (Gy)	OARs
Lung	50.0; 55.0	left lung, right lung, esophagus, heart, spinal cord, u.t.
Liver	50.0; 55.0	left lung, right lung, liver, spinal cord, stomach, left kidney, right kidney, u.t.
Prostate	79; 85.3	left femoral head, right femoral head, anterior rectum, bladder, u.t.

Table 3**Lung:** Planning goals, dose quality and treatment times

	Goal	Initial (I)	PMERGE (P)	VMERGE
u.t. dose	$\text{mean}_{UT}^{(P)} < 1.10 \times \text{mean}_{UT}^{(I)}$	2.8 Gy	3.0 Gy	2.9 Gy
Right lung	$\text{mean}_{RL}^{(P)} < 1.05 \times \text{mean}_{RL}^{(I)}$	0.6 Gy	0.5 Gy	0.6 Gy
Left lung	$\nabla V40_{LL}^{(P)} < 1.05 \times V40_{LL}^{(I)}$	9.3%	9.7%	8.9%
	$V50_{LL}^{(P)} < 1.05 \times V50_{LL}^{(I)}$	3.5%	3.5%	3.4%
Target coverage	$V50_{PTV}^{(P)} > 0.98 \times V50_{PTV}^{(I)}$	100%	98.2%	98.2%
Maximum dose	$\max^{(P)} < 1.02 \times \max^{(I)}$	55.0 Gy	56.0 Gy	55.2 Gy
(Start, end) angles		(0°, 360°)	(60°, 260°)	(0°, 360°)
Treatment time [*]		771 s	127 s	211 s

[†] V_x denotes the volume receiving at least x Gy

^{*}Treatment times for PMERGE and VMERGE are after merging

Table 4**Liver:** Planning goals, dose quality and treatment times

	Goal	Initial (I)	PMERGE (P)	VMERGE
u.t.	$\text{mean}_{UT}^{(P)} < 1.10 \times \text{mean}_{UT}^{(I)}$	5.7 Gy	5.6 Gy	5.7 Gy
Right lung	$\text{mean}_{RL}^{(P)} < 1.05 \times \text{mean}_{RL}^{(I)}$	9.1 Gy	9.2 Gy	9.1 Gy
Left lung	no constraint ($\text{mean}_{LK}^{(P)}$)	1.5	1.5 Gy	1.5 Gy
Liver	$\text{mean}_L^{(P)} < 1.05 \times \text{mean}_L^{(I)}$	23.3 Gy	23.3 Gy	23.4 Gy
Right kidney	$\text{mean}_{RK}^{(P)} < 1.05 \times \text{mean}_{RK}^{(I)}$	5.8 Gy	4.7 Gy	5.9 Gy
Left kidney	no constraint ($\text{mean}_{LK}^{(P)}$)	0.7 Gy	0.4 Gy	0.6 Gy
Target coverage	$V50_{PTV}^{(P)} > 0.98 \times V50_{PTV}^{(I)}$	100%	98.0%	98.4%
Maximum dose	$\max^{(P)} < 1.02 \times \max^{(I)}$	55.0 Gy	55.6 Gy	55.3 Gy
(Start, end) angles		(0°, 360°)	(100°, 340°)	(0°, 360°)
Treatment time [†]		1214 s	245 s	357 s

[†]Treatment times for PMERGE and VMERGE are after merging

Table 5**Prostate: Planning goals, dose quality and treatment times**

	Goal	Initial (I)	PMERGE (P)	VMERGE
u.t.	$\text{mean}_{UT}^{(P)} < 1.10 \times \text{mean}_{UT}^{(I)}$	15.5 Gy	16.4 Gy	15.5 Gy
Anterior rectum	$\text{mean}_{AR}^{(P)} < 1.10 \times \text{mean}_{AR}^{(I)}$	41.4 Gy	41.8 Gy	41.9 Gy
Bladder	$\text{mean}_B^{(P)} < 1.05 \times \text{mean}_B^{(I)}$	39.6 Gy	39.8 Gy	39.9 Gy
Left femoral head	$\text{max}_{LFH}^{(P)} < 1.10 \times \text{max}_{LFH}^{(I)}$	33.6 Gy	35.5 Gy	33.4 Gy
Right femoral head	$\text{max}_{RFH}^{(P)} < 1.10 \times \text{max}_{RFH}^{(I)}$	37.0 Gy	38.3 Gy	38.6 Gy
Target coverage	$V79_{PTV}^{(P)} > 0.98 \times V79_{PTV}^{(I)}$	100%	98.2%	98.2%
Maximum dose	$\text{max}^{(P)} < 1.02 \times \text{max}^{(I)}$	85.3 Gy	86.3 Gy	86.6 Gy
(Start, end) angles		(0°, 360°)	(40°, 320°)	(0°, 360°)
Treatment time [†]		859 s	147 s	178 s

[†]Treatment times for pmerge and vmerge are after merging