

MIP Models and BB Strategies in Brachytherapy Treatment Optimization

Robert R. Meyer (rrm@cs.wisc.edu)*
Computer Sciences Department, University of Wisconsin-Madison, USA

Warren D. D'Souza (wdsouza@mail.mdanderson.org)
Department of Radiation Physics, University of Texas M. D. Anderson Cancer Center, Houston, TX 77030

Michael C. Ferris (ferris@cs.wisc.edu)†
Computer Sciences Department, University of Wisconsin-Madison, USA

Bruce R. Thomadsen (brthomad@facstaff.wisc.edu)
Department of Medical Physics, University of Wisconsin-Madison

Abstract. Brachytherapy (brachy being derived from a Greek word meaning short) is the treatment of cancer by means of radioactive sources that are placed at short distances from the target cells. This form of therapy is becoming common in the treatment of early stage prostate cancer, the most common cancer and the second leading cause of cancer deaths among American males. We consider the use of mixed-integer programming (MIP) models and branch-and-bound (BB) methods to optimize the placement within the prostate of the radioactive “seeds” used in this procedure. Several different optimization models are considered along with a number of branch-and-bound strategies. With appropriate combinations of modelling and solution strategies, near-optimal seed placements can be generated for each two-dimensional ultrasound section of the prostate in less than five minutes on a 333 MHz workstation. The original three-dimensional problem can then be solved by considering an appropriately interrelated sequence of these two-dimensional problems.

Keywords:

brachytherapy, prostate cancer, branch-and-bound, optimization, integer programming, treatment planning.

* Research supported by NSF grant DMI-0100220

† Research supported by NSF grant CCR-9972372 and AFOSR grant F49620-98-1-0417 and Microsoft Corporation



1. Introduction

Prostate cancer is the most common cancer and the second leading cause of cancer death in American males. It is increasingly being treated with radiation therapy. Radiation therapy is the treatment of malignant tumors with radiation (photons, electrons, and heavy charged particles). Radiation can kill both normal and cancerous cells or prevent them from growing and dividing. As these particles travel through cells, they deposit energy through various particle interactions. DNA molecules in human cells are double stranded and helical in nature. The energy deposited by radiation can cause DNA breaks, and this is what is believed to be primarily responsible for killing the cell. The repair mechanism of cancer cells is less efficient than that of normal cells making them more susceptible to radiation.

Broadly, radiation therapy is sub-divided into teletherapy or external beam therapy and brachytherapy. Brachytherapy is the clinical use of small encapsulated radioactive sources at a short distance from or directly in the target volume for irradiation of tumors. The goal of radiation therapy is to deliver adequate dose to the tumor region while simultaneously sparing sensitive structures and normal tissue. Compared to conventional external beam therapy, the physical advantages of brachytherapy result from a superior localization of dose to the tumor volume. Interstitial brachytherapy is the implantation (temporarily or permanently) of radioactive sources directly in the tumor volume, and is used to treat prostate, breast, tongue, and gynecological cancers.

In this paper we focus on the application of interstitial brachytherapy to prostate cancer using permanently implanted Iodine-125 radioactive sources (commonly referred to as “seeds”). This is one of the most common applications of brachytherapy; in fact, NIH (the National Institutes of Health) estimates that the number of prostate permanent radioactive implants performed each year will increase from around 6,000 in 1995 to more than 110,000 in the year 2005 given the current rate of increase of detected cases and the aging population (BBIN, 1996).

Nearly all early stage prostate cancer patients can be appropriately treated with either brachytherapy alone or brachytherapy combined with conventional external beam radiation therapy. A detailed description of the seed implantation procedure can be found elsewhere (Blasko et al, 1987; Grimm et al., 1994). Patients in whom the cancer has spread to the capsule surrounding the prostate fare better with a combination of seed implants and conventional radiation therapy. Other treatment options include prostatectomy (removal of prostate by surgery). The advantages of seed implantation over prostatectomy include the con-

venience of an outpatient procedure, less normal tissue morbidity and the preservation of sexual function.

Effective implants require careful treatment planning based on ultrasound or computed tomography images. Traditionally, interstitial brachytherapy treatment planning has been a trial and error process in which a dosimetry expert tries to improve an initial treatment plan (derived from geometrical properties of the prostate) by iteratively changing (on the basis of expert judgement) the configuration of radioactive sources within the target volume.

From a mathematical viewpoint, one can seek “optimal” seed configurations that take into account a number of criteria, including radiation levels in the target volume and in the nearby normal tissues. Optimization in seed implant brachytherapy thus involves spatially distributing the seeds in order to try to achieve at least a prescribed dose level over the target region while ensuring that the organs at risk and normal tissues receive doses well below levels that may cause radiation injury.

The initial ultrasound “volume study” of the prostate generates 10-15 2-D transverse images (“slices”) and is followed by development of the corresponding brachytherapy treatment plan several days before the procedure. However, in a clinical environment in which multiple patients’ treatment are planned for different disease sites with limited computing power, only an hour or so can be devoted to the determination of a single treatment plan. A second set of ultrasound images is obtained on the day of the procedure. These images may not correlate perfectly with the images acquired previously. Hence, the treatment plan may need to be reoptimized at the time of the implantation procedure. In this context the speed of the optimization algorithm is again a key and necessary feature. The focus of this work was to construct MIP models and to determine node selection, variable selection, branching direction, and scaling strategies for branch-and-bound (BB) search that would result in near seed placement for typical 2-dimensional planar slices in less than 5 minutes of CPU time. The successful extension of this methodology to the 3-D case via optimization in successive 2-D slices is described briefly below and in (D’Souza et al., 2001).

2. Geometry

The prostate gland is a chestnut-shaped structure about 40 mm across and 30 mm thick that surrounds the beginning of the urethra, and is located below the urinary bladder (Hole, 1993). Information on prostate volume and shape is obtained during the “volume study” which is carried out with transrectal ultrasonography. Transverse images (corre-

sponding to planar slices) of the prostate are obtained at 5 mm intervals from the base to the apex of the prostate. The number of transverse images acquired is typically 10-15. Next, the target volume (prostate) is outlined on each ultrasound image along with organs at risk (urethra and rectum) by the radiation oncologist.

In this paper we demonstrate the robustness of the BB approach to optimum seed placement by considering both an idealized phantom geometry and actual patient data (additional patient data considered in (D'Souza et al., 2001) yields similar behavior in all cases). Because of the mechanics of the seed implantation procedure, the possible locations of the seeds correspond to the centers of a 5 mm^2 grid that is imposed on each transverse image. The dose calculation grid is finer than the seed placement grid and for the phantom is an 80×80 region of 1 mm squares. Over this fine grid, an 80×80 dose distribution matrix is pre-computed for each possible source location using the American Association of Physicists in Medicine (AAPM) Task Group 43 report (Nath et al., 1995) formalism for a point source (further details are given below). The grid structure of the problem coupled with radiation pre-computation allows replacement of the continuous nonlinear radiation model by the linear-integer model that we employ in the BB approach (further details are given below). In most cases, the total number of seeds implanted over all of the 2-D sections can vary between 90 and 110 depending on the size of the prostate.

In the phantom shown in Figure 1, the target is simulated by a circular structure with a radius of 20 mm. Located posterior to the target is a structure 13 mm in radius and represents the rectum. Within the circular target area is a circular structure, 3 mm in radius simulating the urethra. In addition, we introduced a subset of the target, referred to as the transition region, 8 mm in radius which is concentrically placed around the urethral region in the phantom. (Sometimes, it is desired that the target dose be raised above the normally prescribed level. While most of the target (namely, the partial target, which is the target less the transition region) receives this higher dose, it may be required that the dose to the portion of the target near the urethra remains unchanged. The transition region is thus added so that it receives less dose than the partial target region but enough to meet the minimum dose criteria for the entire target. Thus, the purpose for the introduction of the transition region is two-fold: (i) to ensure a more gradual fall-off in dose from regions in the target further away from the urethra to regions in the target closer to the urethra, and (ii) to ensure that seeds are not placed too close to the region representing the urethra. The "full" target refers to the partial target plus the transition region.)




Figure 1. Simulated and clinical geometries

Figure 1 shows the actual anatomical configuration from a patient who was subsequently treated with iodine-125 seed implants. This data was acquired from a transverse ultrasound image towards the central part of the prostate. The prostate, along with the organs at risk (urethra and rectum) were contoured. We also contoured a transition region around the urethra. (The rectum appears as a semi-circle in Figure 1 instead of a completed circle, because the ultrasound transducer, while placed in the rectum, images only the upper half of the anatomy.) In this clinical data, the dose calculation grid is $116 \times 87 \text{ mm}^2$. The target in this case is somewhat larger than for the phantom data, so the solutions require more seeds. However, as discussed below, our computational experience in terms of evaluating solution strategy alternatives was similar in the cases of phantom and clinical data (and several other geometries that we considered), so the strategy choices we determined are robust with respect to alternative geometries.

3. Radiation Dose Calculation

The Interstitial Collaborative Working Group (ICWG 1990) and AAPM Task Group 43 have devised a dose calculation formalism for the dosimetry of interstitial brachytherapy (Nath et al., 1995). In reality, iodine sources are cylindrical in shape with a physical length of 4.5 mm and a diameter of 0.8 mm. The radioactive iodine is adsorbed on a silver rod 3 mm in length and encapsulated in a titanium capsule. In the point

source approximation, the dose depends only on the radial distance from the center of the source. The unit of dose is Gray (Gy), where 1 Gy is defined as the deposition of 1 Joule of energy in a 1 kg mass.




Figure 2. Total radiation dose to each point (i,j) in the grid is obtained by summing contributions from the seeds

The dose rate at a distance r from a single source using the point source approximation is given by:

$$\dot{D}(r) = \frac{S_k \Lambda}{r^2} g(r) \Phi_{an} \quad (1)$$

where $\dot{D}(r)$ is the initial dose rate (cGy hr^{-1}) of the source, S_k is the (air kerma) strength (U) of the source measured in units of U, where $1 \text{ U} = 1 \mu\text{Gy m}^2 \text{ hr}^{-1}$ as recommended by the AAPM Report No. 21 (AAPM, 1987), Λ is the dose rate constant of the source ($\text{cGy hr}^{-1} \text{ U}^{-1}$), r is the radial distance (cm) of a point of interest from the source, $g(r)$ is the radial dose function that accounts for the radial dependence of dose on the transverse axis due to photon absorption and scatter in the medium, Φ_{an} is the anisotropy factor and is approximated by a constant in the point source dose formulation (Nath, 1995), (Nath et al., 1995). For Iodine-125 the values of these parameters are: $\Lambda = 0.88 \text{ cGy hr}^{-1} \text{ U}^{-1}$, $\Phi_{an} = 0.993$. For purposes of simplicity we used a source strength S_k of 1 U. The radial dose function used for I-125 is a fitted fifth-order polynomial of the form:

$$g(r) = a_0 + a_1 r + a_2 r^2 + a_3 r^3 + a_4 r^4 + a_5 r^5 \quad (2)$$

and the coefficients of the fit are $a_0 = 1.01376$, $a_1 = 0.122747$, $a_2 = -0.173025$, $a_3 = 0.0402378$, $a_4 = -0.00385227$, $a_5 = 0.000134283$. The dose rate calculated using the point source approximation on the transverse axis (perpendicular to the axis of the cylindrical source) in

the medium is lower than the actual dose rate by 3% to 9% for most interstitial brachytherapy sources (Nath et al., 1995).

The total dose over time to tissue from a permanent radioactive implant is given by integrating the dose rate, taking into account the half-life $T_{1/2}$ of the radioactive source:

$$D(r) = \int_0^{\infty} \dot{D}(r) \exp\left(\frac{-0.693t}{T_{1/2}}\right) dt \quad (3)$$

where $T_{1/2}$ is the the time it takes for the source to exponentially decay to one-half of its original strength (59.6 days for Iodine-125). Even though sources are permanently implanted within the tumor volume most of the radiation dose is delivered during the first five half-lives (~ 300 days for Iodine-125). Carrying out the integration in Equation 3, the total dose from a permanently implanted source is:

$$D(r) = 1.44T_{1/2}\dot{D}(r) \quad (4)$$

4. MIP Models for Prostate Seed Implants

Methods for optimization of prostate implants include fast simulated annealing (SA) (Pouliot et al., 1996), genetic algorithms (GA) (Yu et al.(s), 1996; Ezzell(s), 1996; Yang, 1998) and branch and bound (Gallagher et al., 1997; Lee et al., 1999). Fast simulated annealing has been used by Pouliot *et al.* to optimize the dose distribution by finding the best seed distribution through the minimization of a cost function which includes constraints on the dose at the periphery of the planned target volume and on the dose uniformity within this volume. Yu *et al.* have developed a genetic algorithm for prostate implants in which the objective function is comprised of separable cardinal utility terms. SA and GA techniques cannot provide the solution guarantees that are possible with BB. Gallagher and Lee's model is three-dimensional but imposes constraints only at sampled points in the region. Here we consider quasi-independent two-dimensional slices and impose constraints at all relevant points.

The focus of this work is mixed integer models for prostate seed implant brachytherapy using the branch and bound technique and strategies influencing the outcome of the solution process. The Mixed Integer Program (MIP) model that we used can be represented as follows:

$$\begin{aligned} & \text{minimize} && cx + dy \\ & \text{subject to} && Ax + By \geq b \\ & && y \text{ binary,} \end{aligned}$$

where each binary variable y_P corresponds to the placement or non-placement of a seed in a particular location. The possible seed locations are designated by the use of the ultrasound template. Seeds can only be placed within the target. The placement of a seed in a particular location contributes radiation dose to each point in the dose computation grid.

If y_P is the binary (0/1) variable indicating the presence or absence of a seed in a particular grid location P within the full target then the total dose D_{ij} delivered by all seeds to a point (i,j) in the grid is given by:

$$D_{ij} = \sum_{P=1}^N y_P D_{ij}^P \quad (5)$$

where D_{ij}^P is the dose contribution from seed at grid location P, and N is the number of possible seed positions (in our examples N varies from 50 to 90). Consider Figure 2 for example. If seeds are present in the pixel locations as shown in Figure 2, then the dose to a pixel (i,j) in the 2-dimensional matrix is the sum of the doses from individual seeds 1, 2, and 3 to pixel (i,j).

The primary goal in prostate seed implant brachytherapy is to deliver a certain minimum dose to the prostate gland while minimizing the dose to organs at risk such as the urethra and rectum. Hence a lower bound, T_l is placed on the dose to the (partial) target while upper bounds U_u and R_u are placed on the dose to the urethra and the rectum. In brachytherapy, there is a strong dose gradient, over the tumor volume i.e. the dose is very high at points close to the source and decreases rapidly as the distance from the source increases. This is because the dose contribution to a point is inversely proportional to the square of the distance from the source. Hence, the r^2 term dominates in (1) when r is less than 1 mm. Thus, there is no upper bound on the dose to the target. The above mentioned constraints can be stated as

$$\begin{aligned} D_{ij} &\geq T_l & \forall (i, j) \in T \\ D_{ij} &\geq S_l & \forall (i, j) \in S \\ D_{ij} &\leq U_u & \forall (i, j) \in U \\ D_{ij} &\leq R_u & \forall (i, j) \in R \end{aligned}$$

where T, S, U, and R represent the (partial) target, transition region (remainder of the target), urethra and rectum respectively. The values used for the upper and lower bounds (in units of Gy) are as follows: $T_l=145$, $S_l=140$, $U_u=240$, and $R_u=112$.

We considered several objective functions in our research. In the first objective function (see model 1 below), the mean dose to the urethra was minimized subject to the above constraints. The urethra is

considered the primary organ at risk and therefore minimizing the dose to the urethra is of great importance. The mean dose to the urethra can be stated as:

$$meandose = \frac{\sum_{i,j \in U} D_{ij}}{|U|}, \quad (6)$$

where $|U|$ denotes the number of pixels in the urethra. (While this is mathematically equivalent to minimizing total dose to the urethra, mean dose is of greater clinical interest.)

The second objective function (see model 2 below) involved minimizing the total underdose (relative to the prescription target dose). This can be represented as

$$\min \sum_i \sum_j z_{ij} \quad (7)$$

with the additional constraint

$$D_{ij} + z_{ij} \geq T_l \quad \forall (i,j) \in T \quad (8)$$

where T_l is the dose prescribed to the target and z_{ij} is a continuous non-negative variable representing underdose (relative to the prescription dose). The optimization process ensures that a penalty is introduced only if the target dose at a point is less than the prescription dose. This is thus essentially a feasibility problem, since (in most cases) we are able to find solutions with all $z_{ij}=0$. (In principle, a penalty can be assigned if the target dose exceeds a certain threshold level (e.g., $1.5 \times T_l$) in order to achieve a certain degree of dose homogeneity over the target region. However, it is not clear if dose homogeneity is preferred in prostate seed implants or if dose heterogeneity provides better tumor control, so here we penalize only underdose.)

An MIP model was also constructed seeking to minimize the total number of seeds utilized to achieve the desired dose distribution. Minimizing the number of seeds is desirable in order to reduced the level of trauma induced in the prostate gland. However, this resulted in solution times being too large in most cases. Therefore, instead of minimizing the total number of seeds in the objective function, an upper bound was placed on the number of seeds:

$$\sum_P y_P \leq \maxseed, \quad (9)$$

where \maxseed is the maximum number of seeds allowed. If this upper bound is not set too low, reasonable solution times are obtained as will be discussed below.

The constraints and the objective functions for the two principal models used are formulated below. Model 1 (urethra objective, hard

bound on target):

$$\begin{aligned}
& \min && \frac{\sum_{(i,j) \in U} D_{ij}}{|U|} \\
\text{subject to} &&& D_{ij} = \sum_{P=1}^N y_P D_{ij}^P \quad \forall (i,j) \\
&&& D_{ij} \geq T_l \quad \forall (i,j) \in T \\
&&& D_{ij} \geq S_l \quad \forall (i,j) \in S \\
&&& D_{ij} \leq U_u \quad \forall (i,j) \in U \\
&&& D_{ij} \leq R_u \quad \forall (i,j) \in R \\
&&& \sum_P y_P \leq \text{maxseed} \\
&&& y_P \in 0,1 \quad \forall P
\end{aligned}$$

Model 2 (feasibility objective, soft bound on target):

$$\begin{aligned}
& \min && \sum_i \sum_j z_{ij} \quad \forall (i,j) \in T, S \\
\text{subject to} &&& D_{ij} = \sum_{P=1}^N y_P D_{ij}^P \quad \forall (i,j) \\
&&& D_{ij} + z_{ij} \geq T_l \quad \forall (i,j) \in T \\
&&& D_{ij} + z_{ij} \geq S_l \quad \forall (i,j) \in S \\
&&& D_{ij} \leq U_u \quad \forall (i,j) \in U \\
&&& D_{ij} \leq R_u \quad \forall (i,j) \in R \\
&&& \sum_P y_P \leq \text{maxseed} \\
&&& y_P \in 0,1 \quad \forall P \\
&&& z_{ij} \geq 0 \quad \forall (i,j) \in T
\end{aligned}$$

Finally, we also experimented with a composite objective function constructed by weighting objective 1 by a factor $\alpha < 1$ and objective 2 by $1 - \alpha$ and adding the two weighted objective functions, using the “soft” underdose constraint defining the variables z_{ij} as opposed to a strict lower bound on the target dose. Thus the composite objective function is α x objective 1 + $(1 - \alpha)$ x objective 2. As discussed below, an appropriate choice of α produced good results (from the standpoint of both objectives) in a reasonable time.

5. Branch-and-Bound Strategies

Branch-and-bound is a global optimization technique that recursively partitions relaxations of the feasible set. For the problem of prostate seed implant optimization, a number of different strategies were explored in conjunction with the General Algebraic Modeling (GAMS) language (Brooke et al., 1997) and the CPLEX mixed integer programming solver. We provide a brief description of the various GAMS/CPLEX strategies explored. Table I summarizes these strategy options and the

combinations of options that we found effective. Strategy 1 provided good results in all cases and is the focus of most of the discussion below.

Branching Direction: By default, CPLEX will use the magnitude of the branching variable's integer infeasibility to decide whether to process next the up or down branch. Alternatively, this option forces the consistent selection of either the up or down branch.

Node Selection: We investigated three node selection strategies: 1) depth-first search, 2) best-bound search, and 3) best-estimate search. The depth-first search strategy tends to limit the selection of the next node to one of the two descendants of the current node in the branch and bound tree. The successor node is more efficiently defined since only a single additional lower or upper bound need be specified. While depth-first search results in faster per-node processing times, each branch may be exhaustively searched to the deepest level before fathoming it in favor of better branches, thus consuming considerable time.

The best-bound or best-node search selects the unfathomed node with the best objective function value (or "best-bound"). The best-estimate search uses estimates of the objective function value that would be obtained if all integer variables were forced to integer values. Both best-bound and best-estimate strategies result in more breadth-first tree development.

Variable Selection: Three variable selection strategies were explored: 1) branching on the variable with minimum infeasibility, 2) branching based on pseudo-reduced costs, and 3) strong branching. Branching on a variable with minimum infeasibility may lead more quickly to a first integer feasible solution, but will usually be slower overall to reach the optimal integer solution. The pseudo-reduced costs represent an estimate of the penalties associated with integerizing variables, and thereby tend to identify "important" variables for branching. Strong branching attempts several branches at a node and proceeds with the "best" such branching.

Scaling: The influence of various types of scaling on the problem matrix was noted. The model was run with no scaling, standard scaling and modified, more aggressive scaling. Standard scaling uses an equilibrium scaling method that is generally very effective. Aggressive scaling can produce improvements on some problems and can be used if a particular problem has difficulty staying feasible during the solution process.

6. Software Interfaces

Using MATLAB, the set of dose distribution matrices corresponding to possible seed locations was pre-computed in two dimensions over an 80 x 80 mm² grid (for the phantom data case, with a slightly different grid being used for clinical data) using a pixel resolution of 1 mm in each direction. We selected a pixel resolution of 1 mm because the dose drops off very rapidly in brachytherapy. Using a coarser resolution would result in the loss of dose gradient information which will result in a dramatic difference in the optimization.

GAMS is a high-level modeling system for mathematical programming problems. It consists of a language compiler and a suite of integrated high-performance solvers. GAMS is tailored for complex, large scale modeling applications, and allows the construction of large maintainable models that can be adapted quickly to new situations. The data generated within MATLAB was transferred to GAMS using MATLAB-GAMS interfacing software (Ferris, 1998). The solver used in conjunction with GAMS was CPLEX. The CPLEX MIP algorithm is an implementation of a branch-and-bound search with many algorithmic options to be considered below. The solution obtained was returned into the MATLAB workspace where it was visualized.

Branching direction, scaling, node selection and variable selection strategies were varied and the problem was run with different combinations of branch-and-bound solving options. The two objective functions and their corresponding constraints as described above were considered. The model was run on a SunSPARC UltraTM 10 workstation which has a clock speed of 333 MHz and 256 MB RAM. In each case the CPU time, the branch-and-bound tree node count, the best integer solution, and the lower bound of the relaxation were recorded. In addition, the maximum, minimum and mean doses to the target region and organs at risk and the total number of seeds were also noted. The relative optimality gap for model 1 (which minimizes urethral dose) was pre-set to 3%. The relative gap (rg) is defined as the ratio of the absolute value of the difference between OBJ, the value of the best integer solution found and LB, the relaxation (lower bound) over the lower bound,

$$rg = \frac{|OBJ - LB|}{LB} \quad (10)$$

We investigated the effect of the relative gap on the solution time only for model 1. Model 2 (minimizing target underdose) was run to optimality in all cases.

Based on the final dose distribution, isodose lines were constructed. Isodose line displays are analogous to topographical maps. They are

formed by connecting points on a 2-dimensional grid that receive equal dose. With the help of an isodose line display, the radiation oncologist is able to confirm how well the prescription dose conforms to the target boundary while also examining the doses received by the urethra and the rectum. It is also possible to spatially identify regions within the target or organs at risk that receive excessively high or unacceptably low doses.

7. Results and Discussion

Both optimization models led to clinically interesting and somewhat different results, and additionally displayed rather different computational behavior. This section summarizes these differences.

The MIP relative gap tolerance has a significant impact on the solution time of model 1 (minimizing urethra dose) as can be seen from Figure 3. Although this was noted for all strategies, only three are illustrated in Figure 3. The strategies shown utilize best bound node selection, pseudo-cost variable selection, aggressive scaling and only vary in the choice of the branching direction. As the relative gap is reduced, the solution time increases in a non-linear manner but rate of increase varies depending on the solving strategy used. Dose computation algorithms in brachytherapy, such as the one used in the work described here are accurate to within 5%. Termination with a non-zero relative gap of a few per cent is appropriate. For most of the cases, we set the relative gap within the GAMS modeling language to 3%. A further decrease in this parameter was considered unnecessary due to the accuracy limits of the dose calculation method and the tremendous increase in time for a non-significant increase in the quality of the solution (see Figure 3). For testing purposes, model 1 was actually run to a proven optimal value of 136.9032 (phantom data), but this required 4864 cpu seconds.

In general, we found that the best-bound node selection strategy and pseudo-cost variable selection strategy provided the best results. In the tables listed, we have only included solution strategies that generally performed well with the models previously described and with different stopping criteria (relative gap size) for model 1. Overall, the depth-first search and best-estimate search node selection strategies did not provide solutions in reasonable amounts of time although the best-estimate search did perform well in certain individual instances. Variable selection strategies such as branching on a variable with minimum infeasibility and strong branching generally resulted in an enormous cpu time.




Figure 3. Effect of the relative gap between the current objective function value and the lower bound of the LP relaxation of the MIP problem. The time as a function of the relative gap is shown for three different strategies all of which use best bound node selection strategy, pseudo-cost variable selection strategy, aggressive scaling but which vary in the branching direction strategy at each node. Strategy 1 - algorithm decides branching direction, Strategy 2 - down branch selected first, Strategy 3 - up branch selected first

Table II shows the results obtained with with clinical data. Included in the table is a symbolic description of the strategy employed, the best integer solution (mean urethra dose), lower bound of the LP relaxation, cpu time, number of nodes generated and the relative gap at the time of termination of the solution process. Table III lists the results obtained by minimizing the total underdose to the target (model 2) with clinical data (similar results were obtained with phantom data).

Model 2 was run to optimality while model 1 was terminated when the relative gap was less than or equal to 3%. For this reason there are no entries for the best integer solution, lower bound and relative gap for model 2 in III. The solutions for model 2 are obtained in far less computing time than the near-optimal solutions for model 1. Placing an upper bound on the number of seeds in model 1 increased the cpu time relative to no upper bound. Due to the amount of information present in the tables we selected a strategy (which uses best-bound node selection, pseudo-cost variable selection, aggressive scaling and branching direction decided by the algorithm) referred to as strategy 1 below to illustrate the salient points.

Strategy 1 uses 272 cpu seconds and 4288 nodes to arrive at a feasible solution using model 1 (phantom data) with an upper bound of 18 placed on the total number of seeds. The relative gap between the final

integer solution and the corresponding LP solution is 1.8%. Recall that model 1 was run to optimality in this case and the optimal solution obtained was 136.9032. In the clinical case, (see Table II) where an upper bound of 25 is placed on the number of seeds, a good solution is achieved in 288 cpu seconds with a node count of 1865 and a relative gap of 2.6%. The target in the clinical data is larger than in the phantom and a greater number of seeds are needed in order to meet the minimum dose criterion within the target region.

It was found that scaling affects the speed with which a feasible solution is achieved. Aggressive scaling led to a quicker solution with the best-bound node selection search, pseudo-cost variable selection and branching direction at each node decided by the CPLEX MIP algorithm. Similarly, an equilibrium scaling method provides a solution quickly with the best-estimate node selection search, pseudo-cost variable selection and downward branching direction selected first at each node while other scaling methods fail with this strategy.

The branching direction can significantly impact the solution times when using model 1. The down branch resulted in undesirably long solution times for clinical data and therefore is not listed in Table II. From Table III, it can be seen that the aggressive scaling method fairs extremely well when the down branch is selected first while the equilibrium scaling method does better when the up branch is first picked. Figure 4 summarizes graphically the computing times for the two models and the phantom and clinical anatomies.

It was found that the maximum dose to the structure simulating the urethra in the phantom (OR1) and the urethra itself in the clinical data was considerably lower (20-40%) for model 1 as compared with model 2. Model 1 does in fact seek to drive the dose to the urethra (OR1) as far down as possible subject to the minimum dose constraint on the target. The emphasis in model 2 is on pushing the dose to the target above the minimum criterion while maintaining an upper limit on the dose to the critical organs. Hence, in model 2 there is no objective term forcing down the dose to the critical organs.

Again, we will select strategy 1 to elaborate on the above mentioned points. Since the minimum dose criteria is satisfied by both models and the minimum dose to the full target is almost the same, we will not focus on the target dose in this discussion. It is important to mention that the minimum full target dose is 0-7 Gy higher when using model 2 as compared with model 1. We will place the emphasis on the difference in the maximum doses received by the organs at risk between the two models. Table IV lists the doses to the regions in the phantom and clinical data. Figures 5 and 6 contrast isodose plots (dose distribution maps) for a solution based on urethra dose minimization with a solution




Figure 4. Graphical representation of computing times achieved through viable branch-and-bound search strategies in brachytherapy treatment optimization

based on target underdose minimization (using in both cases strategy 1 with an upper bound of 25 placed on the number of seeds). The prescription dose for treatment of prostate cancer using permanent seed implant brachytherapy is 140 Gy; shown on the isodose plots are the 280, 168, 140, 98, 70 and 14 Gy isodose lines corresponding to 200%, 120%, 100%, 70%, 50%, and 10% of the prescription dose. Also indicated on these isodose plots are the positions of the seeds. With isodose plots it is possible to identify “cold spots” as well as “hot spots” within the anatomical region. “Hot spots” and “cold spots” are not absolute but relative terms used in radiation therapy. A “hot spot” refers to an area receiving much higher dose when compared to other regions while a “cold spot” refers to an area receiving a much lower dose relative to other regions or relative to other areas within the same region. In general “hot spots” near critical organs are to be avoided and “cold spots” interior to the target are undesirable. Figures 5 and 6 show “cold spots” in and surrounding the urethra and extending into the transition region but the shape of these “cold spots” is different in the two treatment plans. The “cold spot” in the urethra dose minimization case (Figure 5) extends beyond the transition region and into the target region, but it is important to reiterate that the full target region (partial target and transition region) does receive at least prescription dose.

For clinical data with urethra dose minimization (model 1), the maximum and mean doses to the urethra are 144.4 Gy and 139.8 Gy respectively, whereas for target underdose minimization (model 2), the maximum and mean doses to the urethra are 208.1 Gy and 167.6 Gy respectively. With the composite objective function described above, which is a weighted combination of the objectives of model 1 and model 2, we found that a weight of 0.8 yielded a solution that came acceptably close to the target dose bound (see Figure 7), yet came close to the low mean urethra dose in model 1. The solution time for this objective was 70 seconds, which was in between the times for the two individual objectives. Note that care is needed in the choice of α in order to obtain high quality solutions within a clinically acceptable time frame. Composite objectives that include both of the above terms plus additional terms such as the number of seeds are considered in (D'Souza et al., 2001).

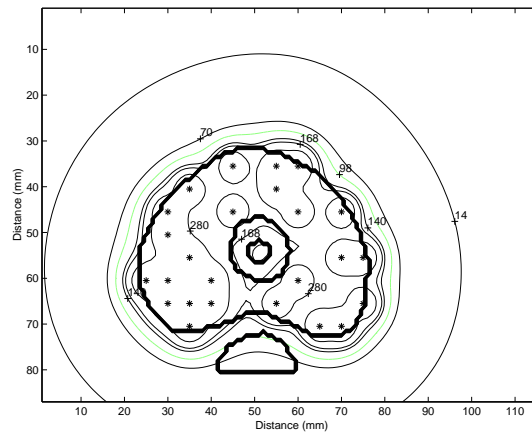


Figure 5. Isodose line display for clinical data with an upper bound of 25 on the number of seeds (mean urethra dose minimization)

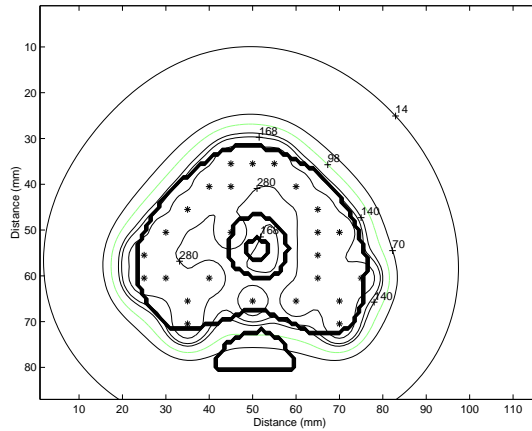




Figure 6. Isodose line display for clinical data with an upper bound of 25 on the number of seeds (target underdose minimization)



a. Solution times (composite objective)



b. Target underdose (composite objective)

Figure 7. Effect of composite objective function of the form $\alpha \times \text{obj1} + (1 - \alpha) \times \text{obj2}$ where obj1 is the objective function that seeks to minimize the mean dose to the urethra and obj2 seeks to minimize the underdose to the target. Strategy 1 and Strategy 2 both employ best bound node selection, pseudo-cost variable selection, aggressive scaling. Strategy 1 allows the algorithm to choose the branching direction while Strategy 2 selects the down branch at each node. (Note that the times are plotted on a log scale)

8. Summary and Conclusions

Mixed integer programming models for ultrasound guided prostate seed implant brachytherapy were presented and successfully solved via branch-and-bound methods. Model 1 focuses on minimizing the dose to organs at risk and is computationally more expensive. Model 2 focuses on target dose, and is faster to solve, but the dose received by the critical structure is higher than with model 1. A composite objective with carefully chosen relative weights for the two objective terms lead to low doses to the critical structure in a computing time that was in between the times required for models 1 and 2. Since the time taken for the optimization process is an important parameter in a clinical setting, one has to weigh the results presented in this paper and decide if a faster solution time is worth the extra dosage to the organs at risk.

From the standpoint of solution quality, introducing a “transition region” between the urethra (OR1) and the rest of the target region did help in reducing the dose gradually across from the target to the urethra. More complex models for constraining the dose gradient between these two regions are now under consideration. Even though the prescribed target dose was achieved using both models, the location of the seeds used to deliver the desired dose distribution is different and depends on the model used and the number of seeds allowed in the target region. Lowering the number of seeds below a certain level may result in undesirable solutions (“hot spots” in the critical structures) and may be computationally expensive.

We have presented a summary of the solution strategies that performed well with all models. In general, the best-bound node selection search and pseudo-cost variable selection search provided the best results. Other node selection strategies such such best-estimate search and depth-first search did not lead to good solution in a desirable amount of time. Variable selection strategies such as branching on a variable with minimum infeasibility and strong branching fared poorly. In all cases, the following strategy was successful: best-bound node selection, pseudo-cost variable selection, aggressive scaling and branching direction decided by the algorithm. However, as seen from the results, several other viable strategies also emerged.

While the results here focus on the application of BB to single 2-dimensional slices, we have demonstrated (D’Souza et al., 2001) that this approach may be successfully extended to three dimensions by using the techniques of this paper on an appropriate sequence of 2-D problems. That is, for each 2-D slice, we first generate a radiation contribution to that plane from the remaining planes, and then optimize the seed placements in the current plane taking into account the

radiation already provided by interplane contributions. Seed positions are thus computed for successive slices until the variation in seed positions becomes sufficiently small. As described in (D'Souza et al., 2001) this sequential approach, coupled with appropriate initializations and controls over the 2-D optimizations, yields full 3-D solutions in about 30 minutes on a 333 MHz processor.

References

- AAPM Report No. 21, Recommendations of AAPM Task Group 32: Specification of Brachytherapy Source Strength, (American Institute of Physics, New York, 1987).
- Biomedical Business International Newsletter (1996) 72-75.
- J.C. Blasko, H. Ragde, D. Schumacher, Transperineal Percutaneous Iodine-125 Implantation for Prostatic Carcinoma Using Transrectal Ultrasound and Template Guidance, *Endocurie./Hypertherim. Oncol.* (1987) 131-139.
- J.C. Blasko, K. Wallner, Brachytherapy Alone for Early Stage Prostate Cancer - Part I: Patient Selection and Technique, (39th Annual Meeting, American Society for Therapeutic Radiology and Oncology, Orlando, FL, 1997).
- A. Brahme, Treatment Optimization Using Physical and Radiobiological Objective Functions, in *Radiation Therapy Physics*, ed. A.R. Smith, pp. 209-246 (Springer-Verlag, New York, NY, 1995).
- A. Brooke, D. Kendrick, A. Meeraus, R. Raman, GAMS: A User's Guide, (The Scientific Press, San Francisco, CA, 1997).
- Warren D. D'Souza, Robert Meyer, Bruce R. Thomadsen, and M.C Ferris, "An iterative sequential mixed-integer approach to automated prostate brachytherapy treatment optimization" *Physics in Medicine and Biology*, 46, 297-322, 2001.
- G.A. Ezzell, Genetic and geometric optimization of three-dimensional radiation planning, *Medical Physics* (1996) 293-305.
- M.C. Ferris, MATLAB and GAMS: Interfacing Optimization and Visualization Software, Mathematical Programming Technical Report 98-19, Computer Sciences Department, University of Wisconsin-Madison, WI, 1998.
- R.J. Gallagher, E.K. Lee, Mixed Integer Programming Optimization Models for Brachytherapy Treatment Planning, *Proceedings/AMIA Annual Fall Symposium* (1997) 278-282.
- P.D. Grimm, J.C. Blasko, H. Ragde, Ultrasound-Guided Transperineal Implantation of Iodine-125 and Palladium-103 for the Treatment of Early Stage Prostate Cancer, *Atlas Urol. Clin. North Am.* (1994) 113-125.
- J.H. Hole Jr., *Human Anatomy and Physiology*, Chap. 22 (Wm. C. Brown Publishers, Dubuque, IA, 1993).
- E.K. Lee, R.J. Gallagher, D. Silvern, C.S. Wu, M. Zaider, Treatment planning for brachytherapy: an integer programming model, two computational approaches and experiments with permanent prostate implant planning, *Phys. Med. Biol.* (1999) 145-165.
- R. Nath, Physical Properties and Clinical Uses of Brachytherapy Radionuclides, in *Brachytherapy Physics*, eds. J.F. Williamson, B.R. Thomadsen, R. Nath, pp. 7-37 (Medical Physics Publishing, Madison, WI, 1995).
- R. Nath, L.L. Anderson, K.A. Weaver, J.F. Williamson, A.S. Meigooni, Dosimetry of interstitial brachytherapy sources: Recommendations of the AAPM Radiation Therapy Committee Task Group No. 43, *Medical Physics* (1995) 209-233.
- J. Pouliot, D. Tremblay, J. Roy, S. Filice, S., Optimization of permanent ^{125}I prostate implants using fast simulated annealing, *Int. J. Rad. Oncol. Biol. Phys.* (1996) 711-720.
- G. Yang, L.E. Reinstein, S. Pai, Z. Xu, A new genetic algorithm technique in optimization of permanent ^{125}I prostate implants, *Medical Physics* (1998) 2308-2315.
- Y. Yu, M.C. Schell, A genetic algorithm for the optimization of prostate implants, *Medical Physics* (1996) 2085-2091.

Table I. Summary of BB options and effective strategies. All strategies listed make use of the pseudo-cost variable selection search.

Node selection (ns)	Variable selection (vs)	Branching direction (br)	Scaling (s)
best-bound (bb)	minimum infeasibility (mi)	up branch (up)	equilibrium scaling (0)
best-estimate (be)	pseudo-cost (pc)	down branch (dn)	aggressive scaling (1)
depth-first (df)	strong branching (sb)	algorithm (alg)	-
Strategy 1	ns=bb, br=alg, s=1		
Strategy 2	ns=bb, br=dn, s=0		
Strategy 3	ns=bb, br=dn, s=1		
Strategy 4	ns=bb, br=up, s=0		
Strategy 5	ns=bb, br=up, s=1		
Strategy 6	ns=be, br=dn, s=0		

Table II. Branch-and-bound solution strategies (relative gap set at 3%) yielding near optimal solutions with **clinical** data in less than 15 minutes of CPU time. The objective is to minimize the mean urethra dose (model 1). Listed in the table is the integer solution at the time of termination.

<i>B-and-B Strategy</i>	<i>Integer So- lution</i>	<i>Lower Bound</i>	<i>CPU time (sec)</i>	<i>Node count</i>	<i>Relative gap</i>
Strategy 1	139.7917	136.3	288	1865	2.6%
Strategy 4	139.1145	136.3	849	8014	2.0%
Strategy 5	139.1145	136.3	871	8014	2.0%

Table III. Branch-and-bound solution strategies yielding optimal solutions with **clinical** data. The objective is to minimize the total target underdose (model 2).

<i>B-and-B Strategy</i>	<i>CPU time (sec)</i>	<i>Node count</i>
Strategy 1	122	1240
Strategy 2	637	6393
Strategy 3	47	447
Strategy 4	25	289
Strategy 5	652	10402
Strategy 6	51	647

Table IV. Dose to regions within phantom and clinical anatomy; results of Strategy 1 applied to model 1 (urethral dose minimization) and model 2 (target underdose minimization).

<i>Trial</i>	<i>Minimum Full Target Dose (Gy)</i>	<i>Maximum OR1/Urethra Dose (Gy)</i>	<i>Mean OR1/Urethra Dose (Gy)</i>	<i>Maximum OR2/Rectum Dose (Gy)</i>
model 1 (phantom)	141.4	144.2	138.3	109.0
model 2 (phantom)	146.2	213.4	162.9	104.2
model 1 (clinical data)	141.9	144.4	139.8	109.5
model 2 (clinical data)	145.6	208.1	167.6	111.9

