7 Motion Planning for Radiation Sources during High-Dose-Rate Brachytherapy

High-dose-rate (HDR) brachytherapy is a type of radiation treatment for cancer. In this procedure, a physician guides radioactive sources through catheters that have been inserted inside or near the cancerous tumors. The goal is to provide a high radioactive dose to treat the tumor while not significantly damaging surrounding healthy tissues. HDR brachytherapy has been successfully used for treating many types of cancer, including prostate cancer [176], cervical cancer [145], and breast cancer [107].

When treating cancer using radiation, physicians desire dose distributions that conform to patient anatomy and satisfy dose prescriptions for the tumor target and nearby critical organs [106]. Using medical images of patient anatomy and estimates of tumor location, physicians prescribe radiation dose requirements for cancerous tumors and surrounding tissues. A sample slice of a CT scan used for this purpose for a prostate cancer patient case is shown in figure 7.1. The goal is then to move the radioactive source inside the catheters to generate a dose distribution that satisfies the clinical criteria as best as possible. This goal can be formulated as an optimization-based motion planning problem: how should we move the radioactive seed through the catheters such that the dose delivered to the patient minimizes the deviation from the prescribed dose?

We draw on linear programming to develop a fast and exact method to optimize radioactive source locations and dwell times for HDR brachytherapy cancer treatment. The method uses the objective and clinical criteria framework of Inverse Planning by Simulated Annealing (IPSA), an approach developed by Lessard and Pouliot in 2000 that has been used in the treatment of over a thousand patients [144, 145, 146]. By formulating the HDR brachytherapy dose optimization problem as a linear program, we enable the fast computation of mathematically optimal solutions.

In this chapter, we present our linear programming formulation and apply the method to a sample of 20 prostate cancer patient cases [13, 14]. We then quantitatively compare the mathematically optimal dwell times solutions for HDR brachytherapy treatment obtained using the LP method to the solutions currently being obtained clinically using simulated annealing (SA), a probabilistic

R. Alterovitz and K. Goldberg: Motion Planning in Medicine, STAR 50, pp. 91–106, 2008.
springerlink.com

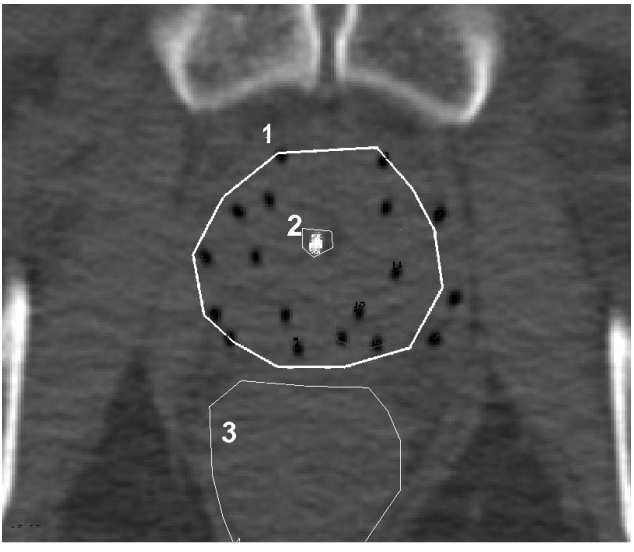

Fig. 7.1. Transverse slice of a CT scan with white contours of the prostate (1), urethra (2), and rectum (3). The catheters are marked with black dots.

method that is not guaranteed to return an optimal solution in finite computation time. We show that the LP method resulted in significantly improved objective function values compared to SA, but the dose distributions produced by the dwell times solutions were clinically equivalent as measured by standard dosimetric indices with a 2% threshold.

7.1 Introduction to HDR Brachytherapy and Dose Optimization

The specifics of the the HDR brachytherapy procedure depend on the cancer site. For the case of the HDR brachytherapy for prostate cancer, the physician commonly implants 14 to 18 catheters in the prostate through the perineum under ultrasound guidance. The physician obtains an image (usually CT scan or MRI) of the catheters and the surrounding tissue, which is used to specify dose prescriptions for the patient anatomy. The catheters are then attached to an HDR Remote Afterloader for treatment delivery. The afterloader, a type of robot, moves a single radioactive source, typically 4.5 mm long and 0.9 mm in diameter containing ^{192}Ir, inside each catheter.

In clinical practice, the seed is generally not moved at a continuous speed through the catheters but rather is temporarily stopped at predetermined dwell positions. Between stops, the seed moves at high speed. The use of predetermined dwell locations converts the problem from a motion planning problem in continuous space to a discrete optimization problem. By adjusting the length

of time (dwell time) that the source remains at any location within a catheter (dwell position), it is possible to generate a wide variety of dose distributions.

To address the dose optimization problem, Lessard and Pouliot developed Inverse Planning by Simulated Annealing (IPSA) [144, 145, 146]. IPSA has been used in the treatment planning of over a thousand patients at UCSF since 2000 and has been independently evaluated by several American and European institutions [54, 64, 132, 150, 154, 197].

A complete description of IPSA and its clinical applications was recently published [177]. Only the elements required for the present work are described here. Using hand-segmented boundaries of the dominant intraprostatic lesions and nearby organs [176], the software generates a discrete sample of dose calculation points inside and on the boundary of the tissue types. For dose calculation points of each tissue type, IPSA permits the physician to prescribe unique dose ranges as well as penalty costs that grow linearly when actual dose violates the prescribed dose ranges. Setting dwell times to minimize dose penalty costs rather than using rigid dose constraints guarantees that the method will find an achievable solution. IPSA defines an objective function equal to a weighted sum of penalty costs at dose calculation points given the dwell times. In the IPSA framework, the mathematically optimal solution is the solution of dwell times that globally minimizes the objective function. IPSA's single objective function assumes that the clinician has specified desirable dose penalty costs and generates a single dwell times solution, in contrast to multi-objective optimization formulations that consider the weights as variables and generate a Pareto front of solutions [134, 135].

The current version of IPSA software uses simulated annealing (SA) to compute dwell times to minimize the objective function. The computation time for a typical case is about 10 seconds on PC with a 3.6 GHz Intel Xeon processor (Nucletron's Masterplan Station). The computation time includes the automatic selection of the active dwell positions, the generation of the dose calculation points, the generation of a look-up dose-rate table, and 100,000 simulated annealing iterations. SA applies a random search with the ability to escape local minima and offers a statistical guarantee to converge asymptotically to the global minimum [1, 89, 205]. The longer the SA algorithm searches for a solution, the higher the probability that the optimal solution is found. Although this method has worked well in clinical practice using 100,000 iterations, there previously was no general quantitative information available regarding the closeness to mathematical optimality of the solutions obtained using simulated annealing, a probabilistic method that cannot guarantee the achievement of a global minimum within a finite computation time.

7.2 Linear Programming for HDR Brachytherapy

Our primary contribution is to take the well-established dose optimization problem defined by IPSA and show that it can be exactly formulated as a linear programming (LP) problem. Because the global minimum for an LP problem

can be computed exactly and deterministically using pre-existing algorithms, this formulation provides strong performance guarantees for cost minimization: one can rapidly find the minimum cost solution for any patient case and clinical criteria parameters. LP does not require setting parameters specific to the optimization method, such as stopping criteria or pseudo-temperatures for SA or mutation probabilities for GA [134, 135, 221]. This allows clinicians to customize dose prescriptions and penalty costs based on medical considerations without concern about their effect on the convergence of the optimization method. Unlike other deterministic algorithms such as local search [146], the LP method will never be trapped at sub-optimal solutions of IPSA's objective function. Since the LP solution is guaranteed to globally minimize the objective function, it provides a precise baseline for evaluating solutions currently being obtained clinically by probabilistic methods such as SA.

Our second contribution is to quantitatively compare the dwell times solutions for HDR treatment currently being obtained clinically using simulated annealing (SA) to the mathematically optimal solutions obtained using LP. With a sample of 20 prostate cancer patient cases, we show that the LP method resulted in significantly improved objective function values compared to SA, but the dose distributions produced by the dwell times solutions were clinically equivalent as measured by standard dosimetric indices.

A linear programming problem is defined by an objective function and constraints that are linear functions of the variables. An LP problem can be solved using the Simplex algorithm, a global deterministic optimization method that considers the geometric polyhedron defined by the linear constraints and systematically moves along edges of the polyhedron to new feasible solutions (represented as vertices of the polyhedron) with successively better values of the objective function until the optimum is reached [161]. In 1990, Renner et al. was the first group to propose a linear programming formulation for HDR brachytherapy dose optimization. Their method minimizes the time the source is irradiating tissue subject to a minimum dose constraint for a set of points in the target volume [179]. Kneschaurek et al. extended this method to permit the specification of dose ranges using rigid constraints for both minimum and maximum dose[123]. Jozsef et al. also used rigid constraints on dose range and minimized the maximum deviation from a prescribed dose constant at dose calculation points [114]. However, a solution of dwell times that results in a dose distribution that satisfies the rigid constraints may not be physically realizable. By defining the dwell times as variables and defining rigid linear constraints on dose, these previous approaches formulated the LP problem in a manner that does not guarantee the output of a solution since no feasible solution may exist. Finding a clinically realizable solution in such cases necessitates arbitrarily removing some rigid dose constraints, which requires substantial human intervention.

Our new linear programming (LP) formulation combines the advantages of IPSA's cost functions and extensive clinical validation with the benefits of deterministic global optimization for cost minimization. We show that the new LP

method computes in finite time the mathematically optimal solution for dwell times to generate the best achievable dose distribution given the clinical objectives and the pre-optimization data generated by IPSA (active dwell positions, dose calculation points, and dose rate look-up table). We applied both SA and the new LP method to 20 prostate cancer patient cases and evaluated improvement of results using objective function values and standard dosimetric indices.

7.2.1 Patient Data Input

The input to the method are 3-D images of the tissues surrounding the tumor. We assume anatomical structures corresponding to b tissue types are segmented, including the clinical target volume (CTV) and critical organs (CO). We also assume the catheters are segmented. From the segmented anatomical structures, we use IPSA to select the active dwell positions and generated a set of m dose calculation points for which the optimization methods will calculate dose. The dose calculation points are distributed based on the anatomy and the implant in order to represent an accurate measurement of the clinical objectives [144]. For each contoured volume, IPSA uses two categories of dose calculation points: "surface" and "volume." This results in $q = 2b$ dose calculation point types: "surface" and "volume" for the b segmented tissue types. For each tissue type, adjusting the dose to "surface" dose calculation points controls the dose coverage and conformality while adjusting the dose to "volume" dose calculation points controls the dose homogeneity [106].

7.2.2 Dose Calculation

Dwell positions are defined as points along catheters at which a source can be placed for a non-zero interval of time. The n active dwell positions were selected by IPSA. We define the dwell time of a source at dwell position j by t_j. A dwell time of 0 corresponds to skipping past a dwell position. The dwell times t_j are the variables that will be set to produce a dose distribution that satisfies the clinical criteria as best as possible.

We calculate the dose-rate contribution d_{ij} of a dwell position j to a dose calculation point i as specified in the AAPM TG-43 dosimetry protocol [162, 180]. The dose-rate contribution is the energy imparted by the radioactive source into an absorbing material (the tissue) per unit time and has units cGy/sec, where 1 gray (Gy) equals 1 joule per kilogram. The dose-rate contribution is a function of r_{ij}, the distance between the dwell position j and the dose calculation point i. It also depends on the radioactive material used in the source, which in this case was ^{192}Ir. Since small differences in the dose calculation may affect the outcome of the optimization, we use the look-up dose-rate table calculated by IPSA as an input for the LP method.

The dose contribution of a dwell position j to a dose calculation point i is computed by multiplying the dose-rate contribution d_{ij} by the dwell time t_j.

The dose D_i at a dose calculation point i, which has units of cGy, is calculated by summing the dose contribution from each dwell position.

$$D_i = \sum_{j=1}^{n} d_{ij} t_j.$$

The dose D_i has units of cGy, which describes the energy imparted by radiation into a unit mass of tissue.

7.2.3 Clinical Criteria

After contouring, the physician prescribes dose ranges for each anatomical structure. The dose ranges used in this study, listed in Table 7.1, are typical values clinically used at the UCSF Comprehensive Cancer Center for treating prostate cancer [106]. This includes the minimum dose D_s^{min} and maximum dose D_s^{max} for each dose calculation point type s. For a dose calculation point i of type s, the desired dose D_{si} should satisfy $D_s^{min} \leq D_{si} \leq D_s^{max}$.

Table 7.1. Clinical criteria parameters for dose penalty cost functions for a typical prostate cancer case

s	Dose calculation point type	D_s^{min} (cGy)	M_s^{min}	D_s^{max} (cGy)	M_s^{max}
1	Prostate (surface)	950	100	1425	100
2	Prostate (volume)	950	100	1425	30
3	Urethra (surface)	950	100	1140	30
4	Urethra (volume)	950	100	1140	30
5	Rectum (surface)	0	0	475	20
6	Rectum (volume)	0	0	475	20
7	Bladder (surface)	0	0	475	20
8	Bladder (volume)	0	0	475	20

In practice, it may not be physically possible to provide a radioactive dose in the physician specified range for every dose calculation point in the 3-D volume. Hence, the physician also specifies a "penalty" for any point for which the clinical criteria is not satisfied. If the actual dose is below or above the prescribed range, the penalty increases linearly at rates M_s^{min} and M_s^{max}, respectively. Adjustment of M_s^{min} and M_s^{max} sets the relative importance of dose range satisfaction between anatomical structures. The penalty weights M_s^{min} and M_s^{max} listed in Table 7.1 are typical values used at the UCSF Comprehensive Cancer Center for prostate cancer cases [106]. The penalty w_{si} at a dose calculation point i of type s can be described in mathematical form using a cost function.

$$w_{si} = \begin{cases} -M_s^{min}(D_{si} - D_s^{min}) & \text{if } D_{si} \leq D_s^{min} \\ M_s^{max}(D_{si} - D_s^{max}) & \text{if } D_{si} \geq D_s^{max} \\ 0 & \text{if } D_s^{min} < D_{si} < D_s^{max} \end{cases} \tag{7.1}$$

Figure 7.2 plots the cost functions (penalty as a function of dose) for the prostate cancer clinical criteria in Table 7.1.

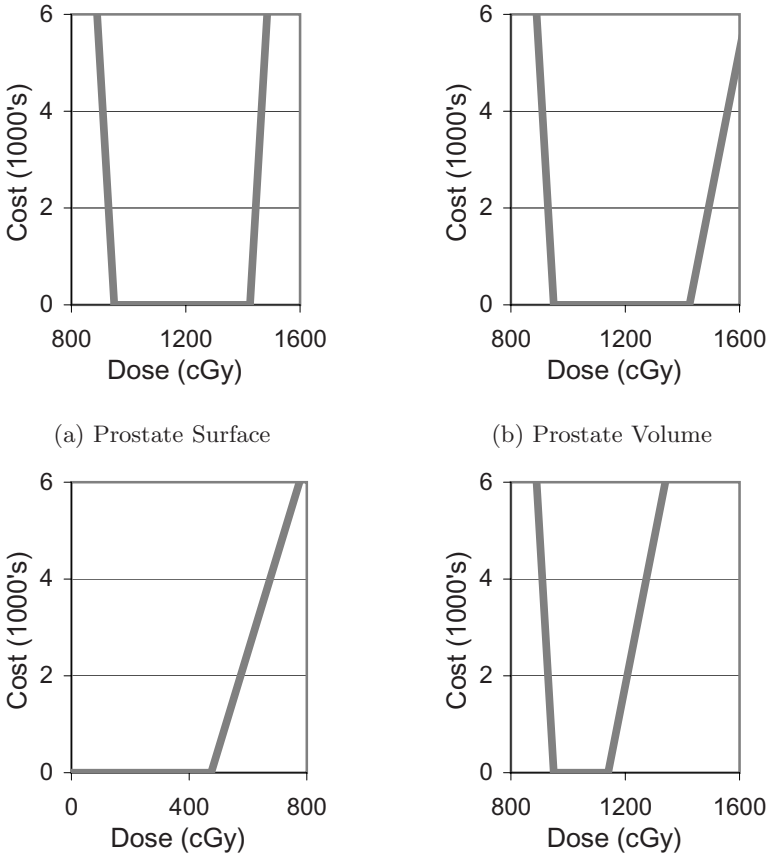

(a) Prostate Surface

(b) Prostate Volume

(c) Urethra

(d) Rectum and Bladder

Fig. 7.2. The clinical criteria, plotted here for a typical prostate cancer case, are specified using cost functions which define penalty as a function of dose for each dose calculation point type

7.2.4 Linear Programming Formulation

The objective is to satisfy the clinical criteria as best as possible by computing dwell times that minimize the net dose penalty costs. Equation (7.1) from section 7.2.3 defines the cost function for an individual dose calculation point i of type s based on the clinical criteria for that point. For each type s, we define the penalty cost E_s as the average penalty cost per point:

$$E_s = \sum_{i=1}^{m_s} \frac{w_{si}}{m_s} \tag{7.2}$$

where m_s is the number of dose calculation points of type s. The objective function E is effectively a weighted sum of the average cost for each tissue type s, where the relative weights are determined by the costs M_s^{min} and M_s^{max}. The global objective function is to minimize the sum of the penalty costs for the q dose calculation point types:

$$E = \sum_{s=1}^{q} E_s = \sum_{s=1}^{q} \sum_{i=1}^{m_s} \frac{w_{si}}{m_s}. \tag{7.3}$$

This objective function is identical to the objective function used by IPSA [146].

The objective function E is not linear because it is composed of nonlinear functions w_{si}. However, each function w_{si} is piece-wise linear. We can formulate this problem as a linear program by creating artificial variables c_{si} to represent cost and defining the following constraints:

$$\begin{aligned} c_{si} &\geq -M_s^{min}(D_{si} - D_s^{min}) \\ c_{si} &\geq M_s^{max}(D_{si} - D_s^{max}) \\ c_{si} &\geq 0. \end{aligned} \tag{7.4}$$

Because w_{si} is a piece-wise linear and convex function, the constraints above guarantee that $c_{si} \geq w_{si}$ for all i, s. Furthermore, we redefine the global objective function to

$$E = \sum_{s=1}^{q} \sum_{i=1}^{m_s} \frac{c_{si}}{m_s}. \tag{7.5}$$

For minimized E where the costs c_{si} satisfy the inequalities (7.4), we are guaranteed $c_{si} = w_{si}$ for all s, i. We show this by proving the contrapositive ($c_{si} \neq w_{si}$ implies E not minimized), which is logically equivalent [198]. If $c_{si} \neq w_{si}$, then $c_{si} > w_{si}$ for some s, i and there will exist a cost c'_{si} such that $c_{si} > c'_{si} \geq w_{si}$. Since c'_{si} will not violate any constraint in inequalities (7.4), it is feasible. We define E' exactly as E except using c'_{si} instead of c_{si}. Hence, $E' < E$ and no cost variables used to compute E' violate a constraint, which implies E is not minimized. Hence, for minimized E, we are guaranteed $c_{si} = w_{si}$.

Table 7.2. HDR dose optimization LP formulation constants, variables, and functions

Constants:	
m_s	Number of dose calculation points of type s.
N	Number of dwell positions.
d_{sij}	Dose-rate contribution from dwell position j to dose calculation point i of type s.
Variables:	
t_j	Source dwell time for dwell position j.
c_{si}	Penalty cost at dose calculation point i of type s.
Objective:	
E	Global cost function.

We summarize constants, variables, and the objective function for the LP formulation in Table 7.2. In equation (7.6), we explicitly define the linear program in canonical form [161] by plugging into the constraints the dose distribution D_{si} at point i of type s due to dwell times t_j.

$$\begin{aligned}
&\text{Minimize} \quad E = \sum_{s=1}^{q} \sum_{i=1}^{m_s} \frac{c_i}{m_s} \\
&\text{Subject to:} \\
&c_{si} + \sum_{j=1}^{n} M_s^{min} d_{sij} t_j \geq M_s^{min} D_s^{min} \qquad s = 1, \ldots, q; i = 1, \ldots, m_s \\
&c_{si} - \sum_{j=1}^{n} M_s^{max} d_{sij} t_j \geq -M_s^{max} D_s^{max} \quad s = 1, \ldots, q; i = 1, \ldots, m_s \\
&c_{si} \geq 0 \qquad\qquad\qquad\qquad\qquad\qquad\quad s = 1, \ldots, q; i = 1, \ldots, m_s \\
&t_j \geq 0 \qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad j = 1, \ldots, n
\end{aligned} \tag{7.6}$$

A (non-optimal) feasible solution for the LP formulation can be trivially found by setting $t_j = 0$ for all j and setting

$$\begin{aligned}
c_{si} = \max\{ \; & -M_s^{min} \left(\left(\sum_{j=1}^{n} d_{sij} t_j \right) - D_s^{min} \right), \\
& 0, \\
& M_s^{max} \left(\left(\sum_{j=1}^{n} d_{sij} t_j \right) - D_s^{max} \right) \}
\end{aligned}$$

for all i and s.

Because of the properties of the artificial variables c_{si} shown above for minimized E, the optimal solution obtained for the linear program in equation (7.6) will be the same as the optimal solution to the nonlinear formulation based on the objective function in equation (7.3) with the cost functions in equation (7.1). We effectively transformed the nonlinear IPSA optimization problem in equation (7.3) (for which deterministic optimization algorithms such as local search could be trapped in sub-optimal solutions [146]) to a higher dimensional space with artificial variables in which an equivalent linear formulation (7.6) can be minimized deterministically to find the global optimal solution using the Simplex algorithm.

7.3 Application to Prostate Cancer Treatment

We implemented software using C++ to read patient specific parameters from IPSA and output the linear program (7.6) in the file format of AMPL (A Mathematical Programming Language) [85]. We solved the linear program specified

in each AMPL file using ILOG CPLEX 9.0, an advanced implementation of the Simplex algorithm [161] designed for large industrial optimization problems [109]. Computation was performed on a 3.0 GHz Pentium IV computer running the Linux operating system.

7.3.1 Patient Data Sets

We applied the LP method retrospectively to 20 prostate cancer patient cases. The prostate volumes ranged from 23 cc to 103 cc. For these patients, the physician implanted 14 to 18 catheters in the prostate with transrectal ultrasound (TRUS) guidance while the patient was under epidural anesthesia. Then Flexiguide catheters (Best Industries, Inc., Flexi-needles, 283-25 (FL153-15NG)), which are 1.98 mm diameter hollow plastic needles through which the radioactive source moves, were inserted transperineally by following the tip of the catheter from the apex of the prostate to the base of the prostate using ultrasound and a stepper. A Foley catheter was inserted to help visualize the urethra.

After catheter implantation, a treatment planning pelvic CT scan was obtained for each patient. Three-millimeter-thick CT slices were collected using a spiral CT. The clinical target volume (CTV) and critical organs (CO) including bladder, rectum, and urethra were contoured using the Nucletron Plato Version 14.2.6 (Nucletron B.V., Veenendaal, The Netherlands). The CTV included only the prostate and no margin was added. When segmenting the bladder and rectum, the outermost mucosa surface was contoured. The urethra was defined by the outer surface of the Foley catheter. Only the urethral volume within the CTV was contoured. The CO's were contoured on all CT slices containing the CTV and at least two additional slices above and below. Implanted catheters were also segmented. A slice of a 3-D CT scan with contoured anatomical structures (prostate, urethra, rectum, and bladder) and catheters is shown in figure 7.1.

For the 20 cases, the number of dose calculation points m ranged from 1781 to 3510. Since the selection of the active dwell positions and dose calculation points affects the outcome of optimization [133], we use those generated by IPSA as input for the LP method. For the prostate cases, the images contained $q = 8$ dose calculation point types: "surface" and "volume" for the four contoured tissue types (prostate, urethra, bladder, and rectum). The clinical criteria used in this study are shown in 7.1.

All patients were treated at UCSF Comprehensive Cancer Center using dosimetric plans generated by the current version of IPSA. We used imaging and dosimetry records from those treatments to compare SA with LP.

7.3.2 Evaluation Metrics

We recorded the dwell times and the objective function value E for the solutions obtained using SA and LP. We evaluated the resulting dose distributions using standard dosimetric indices, including prostate V100 and V150 (the percentage

of the prostate receiving over 100% and over 150% of the prescribed dose, respectively). As dose inside the prostate should fall between 100% (D^{min}) and 150% (D^{max}) of prescribed dose, ideally V100 should be 100% and V150 should be 0%. Similarly, we also evaluated V100 and V150 for the urethra. Dosimetric indices for normal structures (non-cancerous tissues) include the rectum (V50 and V100) and the bladder (V50 and V100). Because normal structures should be spared radioactive dose, these indices ideally should be close to 0%. We also computed dosimetric indices in absolute dose, including the prostate D90 (the maximal dose that covers 90% of prostate volume), urethra D10 (the maximal dose that covers 10% of urethra volume), and rectum and bladder D2cc (the maximal dose that covers 2 cc of the organ volume).

7.3.3 Results

ILOG CPLEX solved for the optimal solution to the linear programming formulation in an average time of 9.00 seconds per case with a standard deviation of 3.77 seconds for the 20 prostate cancer patient cases. The times ranged from 3.68 seconds to 14.63 seconds. The Simplex algorithm in ILOG CPLEX required an average of 1653 iterations with a standard deviation of 341 iterations.

The average objective function value for the 20 prostate cancer patient cases was 3.27 for the LP method compared to 3.33 for SA. The percent difference in objective function value between the solution found using SA and the optimal solution found using LP for each individual patient case is shown in figure 7.3.

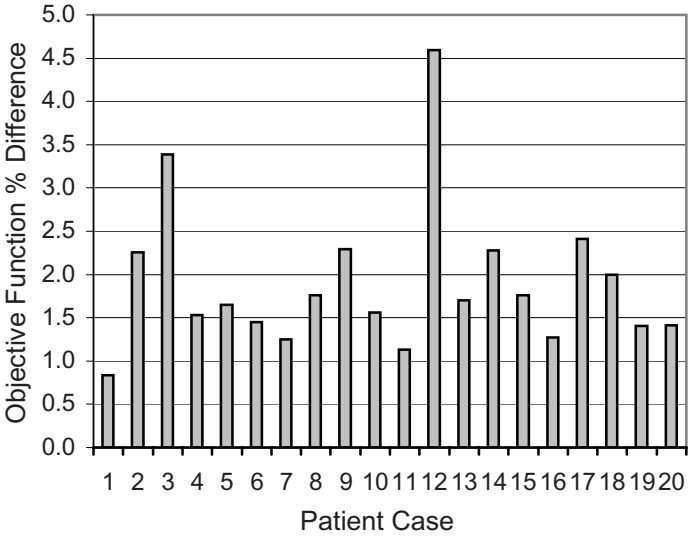

Fig. 7.3. The percent difference in objective function value between the optimal solution (found using the LP method) and the solution found by SA for 20 prostate cancer patient cases. The difference is statistically significant ($P = 1.54 \times 10^{-7}$).

Table 7.3. Improvement of LP solutions over SA solutions for 20 prostate cancer patient cases calculated as the absolute difference in dosimetric index percent values. Negative values indicate deterioration in the dosimetric index. The significance P of the differences was computed using paired t-tests.

Dosimetric Index	Maximum Improvement	Minimum Improvement	Mean Improvement	99% CI	Significance P
Prostate V100	0.95	-0.49	0.13	(-0.10, 0.37)	0.1644
Prostate V150	1.65	-1.63	0.51	(-0.02, 1.04)	0.0217
Urethra V100	1.52	-1.50	0.12	(-0.33, 0.57)	0.4858
Urethra V150	0.11	-0.05	0.00	(-0.01, 0.02)	0.7621
Rectum V50	0.50	-0.81	-0.17	(-0.36, 0.02)	0.0344
Rectum V100	0.03	0.00	0.01	(-0.00, 0.01)	0.0289
Bladder V50	0.75	-0.48	0.03	(-0.17, 0.23)	0.7042
Bladder V100	0.13	-0.02	0.02	(-0.00, 0.04)	0.0225

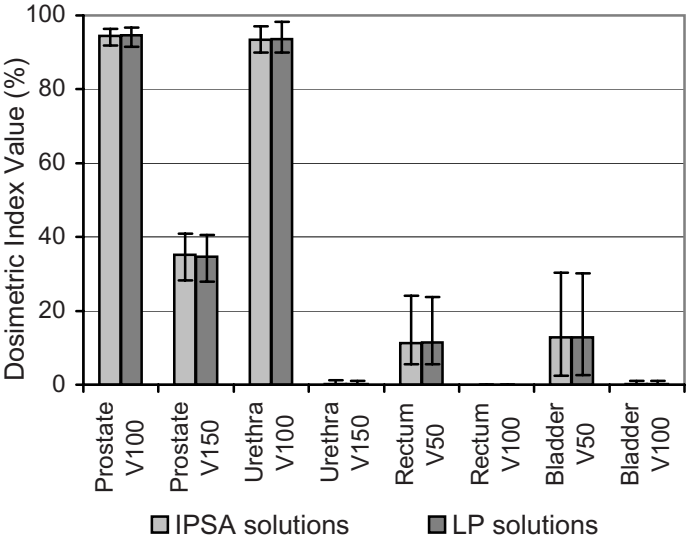

Fig. 7.4. Mean dosimetric index results for the SA and LP methods for 20 prostate cancer patient cases. Error bars indicate maximum and minimum values for the 20 patient cases.

Improvement varies from a minimum of 0.84% to a maximum of 4.59%. We performed paired t-tests to determine the statistical significance ($P < 0.01$) of the results and found that the improvement in objective function value using the LP method compared to SA was statistically significant ($P = 1.54 \times 10^{-7}$).

Figure 7.4 displays the standard dosimetric indices for both the SA and LP solutions. The bars indicate the mean indices as percents and the error bars indicate the maximum and minimum indices obtained for the 20 prostate cancer

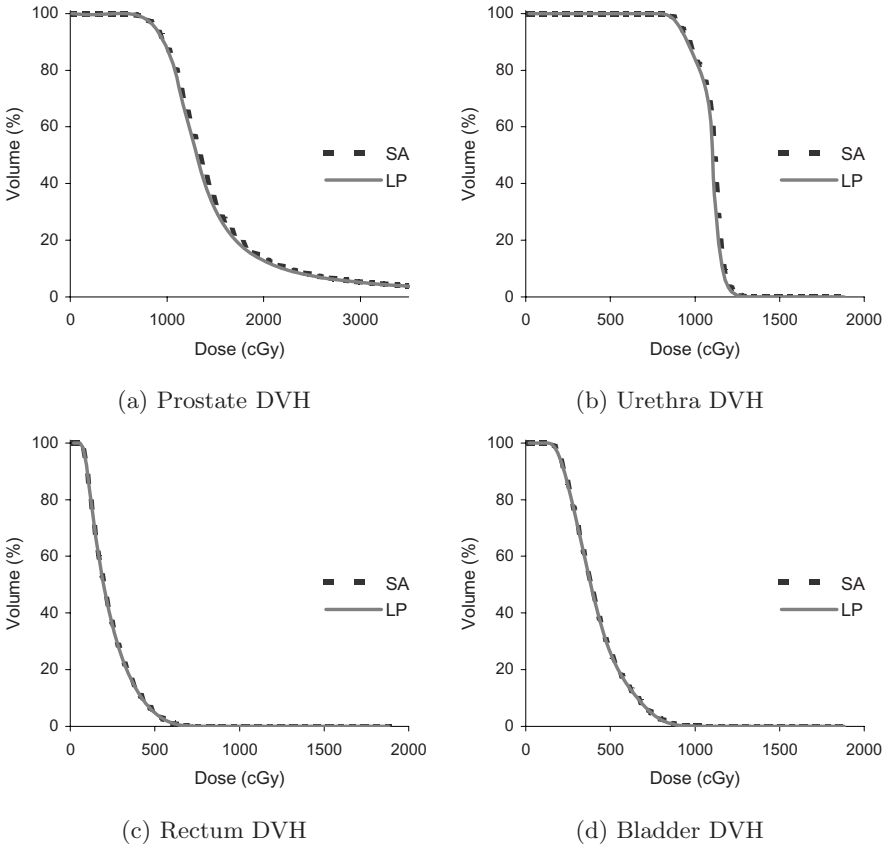

Fig. 7.5. Dose-volume-histogram (DVH) plots for the prostate (a), urethra (b), rectum (c), and bladder (d) for the patient case with greatest difference in dosimetric indices between the LP and SA solutions. For dose less than D^{min} for each tissue type, the desired volume is 100%. For dose greater than D^{max}, the desired volume 0%.

patient cases. Based on these dosimetric indices, the difference between the dose distributions generated by SA and LP was small. None of the dosimetric indices indicated a statistically significant ($P < 0.01$) difference between the dose distributions generated by SA and LP. The largest improvement for the prostate D90, the rectum D2cc, and the bladder D2cc were lower than 1%. The largest improvement for the urethra D10 was 2%. The urethra V150 was zero for both LP and SA method for this case. Additional dosimetric indices are shown in Table 7.3 where positive values indicate improvement and negative values indicate deterioration. The deterioration of one dosimetric index is sometimes traded for the improvement of other dosimetric indices and the improvement of the global solution. The maximum improvement of LP over SA was a reduction of 1.65% for the prostate V150 index. However, for the same patient, LP resulted in a reduction of 0.38% of the prostate V100. Similarly, the maximum deterioration

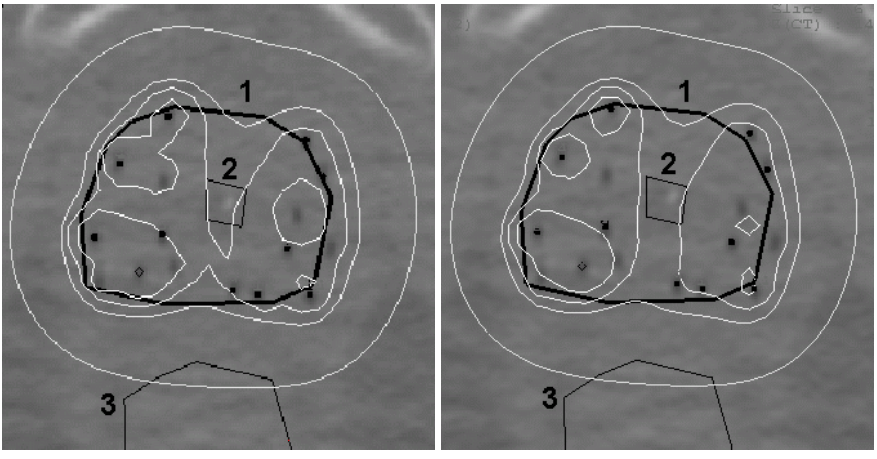

(a) SA Solution (b) LP Solution

Fig. 7.6. Isodose curves for the SA (a) and LP (b) solutions for the patient case with greatest difference in dosimetric indices. The prostate (1), urethra (2), and rectum (3) are contoured in black. Catheters are shown as black dots. Isodose curves for 50%, 100% (D^{min}), 120%, and 150% (D^{max}) of prostate minimum prescribed dose are plotted in white.

of LP over SA was an increase of 1.63% for the prostate V150 index inducing an improvement of 0.84% of the prostate V100. Even with these two extreme cases, the LP and SA methods provide two different solutions that are difficult to distinguish clinically. Figure 7.5 plots the dose-volume-histogram (DVH) for each tissue type for the patient case with the greatest magnitude improvement in a dosimetric index between the SA and LP solutions. Figure 7.6 displays a CT scan of the same patient with overlaid isodose contours for both solutions.

7.4 Discussion

The dosimetric index results are not significantly different from those of the current version of IPSA, which was previously shown to be superior to the commonly used method of geometric optimization followed by manual adjustment [106, 132]. The small variances observed for the prostate and urethra in figure 7.4 show the consistency of the treatment plan quality for both the SA and LP methods. The larger variances for the prostate V150, the rectum, and the bladder are due to differences between patients in anatomy, prostate volume, and distances between the prostate and organs at risk.

The LP and SA methods are both based on IPSA's objective function for the HDR brachytherapy dwell time optimization problem. The only difference is the optimization algorithm used, simulated annealing versus an equivalent linear programming formulation that can be solved using the Simplex algorithm. As

simulated annealing is a probabilistic method, it is only guaranteed to converge to an optimal solution after an infinite amount of computation time. Standard termination criteria, such as stopping the algorithm after a fixed number of iterations, can result in sub-optimal solutions. During the development phase of the current version of IPSA, a large number of cases were run using a very large number of iterations (>1 million) and no significant improvements in the dosimetric indices were found compared to the values found after 100,000 iterations. However, the closeness to mathematical optimality of the solutions of the current version of IPSA could not be guaranteed for every new clinical case.

Because the LP formulation of IPSA's objective function can be solved deterministically to find the solution that globally minimizes costs, the LP method solution provides a precise baseline for evaluating solutions obtained by probabilistic methods such as SA. The LP method computed a solution with a better objective function value compared to SA for every patient case. The improvement in objective function values of LP compared to SA was statistically significant. However, the effect size of the objective function improvement was not sufficient to result in statistically significant differences in standard dosimetric indices for our sample of 20 prostates with volume ranging from 23 cc to 103 cc. We observe that the DVH plots for the patient case with the largest difference in dosimetric indices are similar for both methods (figure 7.5) while differences are observable on the isodose curves (figure 7.6). The hot spots (prostate V150) have different shapes and the prostate V120 curve is at a different location. This indicates that the local dose distribution (isodose) is different while global dose delivered to the organs (DVH) and critical dose delivered to the organs (dosimetric indices) are equivalent. This quantitatively indicates that the dose distributions generated by SA are clinically equivalent to the best achievable dose distributions based on the current IPSA objective function with dose constraints and penalty weights selected for prostate cancer cases.

7.5 Conclusion and Open Problems

HDR brachytherapy requires that clinicians solve a motion planning problem: how should a radioactive source move through pre-implanted catheters to deliver the best achievable dose to the patient? This problem can be formulated as an optimization-based motion planning problem: set dwell times for the radioactive source at dwell positions along the catheters such that the resulting dose distribution minimizes the deviation from physician-specified dose prescriptions. The primary contribution of this chapter is to take the well-established dwell times optimization problem defined by Inverse Planning by Simulated Annealing (IPSA) developed at UCSF and exactly formulate it as a linear programming (LP) problem. Because LP problems can be solved exactly and deterministically, this formulation provides strong performance guarantees: one can rapidly find the dwell times solution that globally minimizes IPSA's objective function for any patient case and clinical criteria parameters. For a sample of 20 prostate

cancer patient cases, the new LP method optimized dwell times in less than 15 seconds per case on a standard PC.

We quantitatively compared the dwell times solutions currently being obtained clinically using simulated annealing (SA), a probabilistic method, to the mathematically optimal solutions obtained using the LP method. The LP method resulted in significantly improved objective function values compared to SA, but none of the dosimetric indices indicated a statistically significant difference. The results indicate that solutions generated by the current version of IPSA are clinically equivalent to the mathematically optimal solutions.

IPSA's objective function with dose constraints and penalty weights covers all organs and all clinical objectives so they can be optimized simultaneously. The physician can adjust the objectives for each optimization. However, if a particular set of objectives generates the desired results, then the same set of objectives can be used for optimization of clinically similar cases (i.e. prostate) without further adjustments. This set of objectives, commonly called a class solution, can be used as a starting point for every patient, significantly reducing the time needed to plan individual patient treatments.

Although we focused on the application of our linear programming formulation to prostate cancer patient cases, the mathematical formulation can also be applied to other cancer types for which HDR brachytherapy is used by incorporating different clinical parameters. The method can also be extended to support any piece-wise linear convex cost functions, not solely the 3-piece cost functions presented above. Recent developments in magnetic resonance spectroscopy imaging and image registration introduce a new clinical criterion, a dose boost to the tumor volume within the prostate [8, 11, 122]. Although we do not explicitly consider that dose calculation point type, the mathematical formulation we defined can be extended to incorporate it by adding a tumor volume tissue type. A potential advantage of the LP method for each of these extensions is that it will use the well-established framework of IPSA and deterministically compute mathematically optimal dwell time solutions for all patient cases.

Although we used linear programming purely as an optimization method in this chapter, linear programming brings with it a vast literature of tools and extensions. This includes well-established tools like sensitivity analysis [161], as well as newer extensions like robust optimization that considers uncertainty in input parameters [33]. In future work, we plan to explore these tools and extensions to provide clinicians with patient-specific information about the trade-offs between feasible treatment plans.