Coverage-based treatment planning: Optimizing the IMRT PTV to meet a CTV coverage criterion

J. J. Gordon^{a)} and J. V. Siebers

Department of Radiation Oncology, Virginia Commonwealth University, P.O. Box 980058, Richmond, Virginia 23298

(Received 30 May 2008; revised 5 December 2008; accepted for publication 28 December 2008; published 25 February 2009)

This work demonstrates an iterative approach—referred to as coverage-based treatment planning—designed to produce treatment plans that ensure target coverage for a specified percentage of setup errors. In this approach the clinical target volume to planning target volume (CTV-to-PTV) margin is iteratively adjusted until the specified CTV coverage is achieved. The advantage of this approach is that it automatically compensates for the dosimetric margin around the CTV, i.e., the extra margin that is created when the dose distribution extends beyond the PTV. When applied to 27 prostate plans, this approach reduced the average CTV-to-PTV margin from 5 to 2.8 mm. This reduction in PTV size produced a corresponding decrease in the volume of normal tissue receiving high dose. The total volume of tissue receiving ≥ 65 Gy was reduced on average by 19.3% or about 48 cc. Individual reductions varied from 8.7% to 28.6%. The volume of bladder receiving ≥ 60 Gy was reduced on average by 5.6% (reductions for individuals varied from 1.7% to 10.6%), and the volume of periprostatic rectum receiving ≥ 65 Gy was reduced on average by 4.9% (reductions for individuals varied from 0.9% to 12.3%). The iterative method proposed here represents a step toward a probabilistic treatment planning algorithm which can generate dose distributions (i.e., treated volumes) that closely approximate a specified level of coverage in the presence of geometric uncertainties. The general principles of coverage-based treatment planning are applicable to arbitrary treatment sites and delivery techniques. Importantly, observed deviations between coverage implied by specified CTV-to-PTV margins and coverage achieved by a given treatment plan imply a generic need to perform coverage probability analysis on a per-plan basis to ensure that the desired level of coverage is achieved. © 2009 American Association of Physicists in Medicine. [DOI: 10.1118/1.3075772]

Key words: radiotherapy, prostate, IMRT, setup errors, margins

I. INTRODUCTION

In margin-based treatment planning, the clinical target volume to planning target volume (CTV-to-PTV) margin is intended to ensure that the CTV is covered for a specified percentage of patient setup errors.^{1,2} Setup errors follow a probability distribution. Accommodating the most extreme errors—those falling in the tails of the distribution—would require an unacceptably large margin, leading to irradiation of a large volume of normal tissue. The accepted approach is therefore to select a CTV-to-PTV margin that will absorb a reasonable percentage of errors. For example, the van Herk margin formula³ estimates the margin required to absorb 90% of systematic setup errors. (The systematic error is the average error for the whole treatment course. Random errors are those occurring in each fraction.⁴)

The primary effect of random (per-fraction) setup errors is to blur the dose distribution, while the effect of systematic setup errors is to move the CTV with respect to the planned dose distribution. Consider an ideal dose distribution that conforms perfectly to the PTV (i.e., which has an isodose surface that hugs the PTV), and assume this dose distribution is blurred by random setup errors to account for per-fraction setup deviations. The CTV-to-PTV margin is typically selected so that the CTV remains inside this blurred dose dis-

tribution for the specified percentage of systematic setup errors. The plan is then created without further regard to CTV coverage. A key assumption in this approach is that the dose distribution conforms perfectly to the PTV, so that any movement of the CTV outside the (blurred) PTV results in loss of coverage.

However, the planned dose distribution rarely conforms perfectly to the PTV. Instead, it extends beyond the PTV in some directions, creating an additional dosimetric buffer around the PTV that can absorb further setup errors.⁵ This effect is illustrated in Fig. 1(a), which shows (in two dimensions) a CTV, PTV, and corresponding treated volume (TV). The ICRU defines the TV to be the “volume enclosed by an isodose surface, selected, and specified by the radiation oncologist as being appropriate to achieve the purpose of the treatment.”¹ In this work, the TV is defined to be the volume enclosed by the PTV minimum dose isodose surface. We further define the dosimetric margin to be the distance between the CTV and TV in direction (ϕ, θ) (see Fig. 1) and the dosimetric margin distribution (DMD) to be the distribution of dosimetric margins over all directions (ϕ, θ) .⁵

Figure 1(a) shows that CTV coverage is a function of the TV or DMD, not the PTV. As long as the CTV remains within the TV it will receive acceptable coverage. (In this

FIG. 1. Illustration of relationship between the clinical target volume (CTV), planning target volume (PTV), and treated volume (TV). The TV is the volume enclosed by the PTV minimum dose isodose surface. The dosimetric margin $D(\phi, \theta)$ is the distance from the CTV to the TV in direction (ϕ, θ) . The dosimetric margin distribution (DMD) is the distribution of $D(\phi, \theta)$ over all directions (ϕ, θ) . (a) Primary prostate structures. (b) Nodal prostate structures surrounding primary PTV.

work “acceptable coverage” signifies that the CTV minimum dose exceeds the planned PTV minimum dose, which is assumed to be the dose that the physician approves as being suitable for treatment.) The effect of the TV extending beyond the PTV is to provide an additional volume around the CTV that can absorb setup errors. This volume is typically not accounted for when selecting the CTV-to-PTV margin. For this reason, margin-based treatment planning TVs can be larger than those actually required to absorb the specified percentage of setup errors.

Previous studies^{5,6} of simultaneous integrated boost (SIB) intensity modulated radiation therapy (IMRT) prostate plans have shown that the TV can be substantially larger than the PTV. In these cases, conventional margin-based planning produces dose distributions that absorb a larger percentage of setup errors than is actually required. Evaluation of CTV coverage after the fact (i.e., after plan optimization) would

expose this effect, but clinically this is rarely performed. In these SIB IMRT plans, the potential therefore exists to shrink the PTV and TV while still maintaining the desired level of CTV coverage. Shrinking the PTV and TV is potentially desirable because it can reduce the dose to normal tissue.

The goal of this work is to demonstrate an iterative planning method—referred to as coverage-based treatment planning—that is designed to produce treatment plans that meet a specified coverage criterion. At each iteration of this method the DMD is computed. From the DMD one calculates CTV coverage (i.e., the probability that the CTV will remain within the TV). The CTV-to-PTV margin is then iteratively adjusted until the specified CTV coverage is achieved. The advantage of this iterative approach is that it automatically compensates for the dosimetric margin around the CTV, shrinking the PTV and TV as much as possible while still meeting the CTV coverage criterion. This work applies the proposed method to the 27 SIB IMRT prostate plans studied previously^{5,6} and, for a 90% CTV coverage criterion, evaluates the degree to which the PTV and TV can be reduced.

The method of coverage-based treatment planning proposed here represents a step toward a probabilistic treatment planning (PTP) algorithm which can generate dose distributions that closely approximate a specified level of coverage while accounting for geometric uncertainties. Probabilistic treatment planning can be roughly defined as optimization methods which incorporate probabilistic models of geometric errors into intensity modulated radiation therapy (IMRT) inverse planning. Geometric errors include patient setup errors. Potentially they could also include interfraction and intrafraction organ motions and target and organ at risk (OAR) delineation errors. Alternative approaches to PTP may be summarized as follows.

The approach proposed here requires that target coverage be achieved for a specified percentage of errors. Other published approaches to PTP pursue a different approach. They either optimize a probability weighted dose distribution (PWDD) or use a probability weighted objective function (PWOFF) to obtain a robust dose distribution. For random (per-fraction) setup uncertainties,⁴ a common way of obtaining a PWDD is to convolve the fluence with the error distribution prior to dose calculation. For conventional fractionation schemes, this convolution method⁷ (CM) provides an acceptable approximation to the cumulative blurring effect of many setup errors. The optimizer tries to “undo” the fluence convolution by deconvolving the uncertainty distribution from the prescribed dose distribution. The end result is the addition of “horns” to the per-fraction (preconvolution) fluence profile.⁸ The horns counteract the effect of random errors, albeit imperfectly. Application of PWDD-type techniques to PTP has been investigated by Li and Xing,⁹ Birkner *et al.*,¹⁰ and Trofimov *et al.*¹¹

McShan *et al.*¹² generalized the CM approach by using a probabilistically weighted finite sum of sample dose distributions, rather than a convolution. The advantage of this technique is that errors are no longer restricted to rigid body translations. Unkelbach *et al.*^{13–16} also developed a PWDD-

type approach to PTP based on optimization of the dose expectation (first moment). They noted that optimizing the average dose can lead to unsatisfactory doses in some error scenarios. They avoided this problem by including a dose variance term, in addition to the expectation, in the objective function. The variance term dampened the more extreme dose distributions.

A number of authors have explored PWOFF approaches, including Lof *et al.*,¹⁷ Baum *et al.*,¹⁸ Sir *et al.*,¹⁹ Yang *et al.*,²⁰ and Witte *et al.*²¹ Lof *et al.* proposed a comprehensive adaptive control framework for fractionated radiotherapy, that incorporated random setup uncertainties, and demonstrated the occurrence of horns in the resulting dose distributions. This work was pursued by Reh binder *et al.*²² who used fluence modulation to compensate for random errors and couch corrections for systematic errors. Baum *et al.*¹⁸ and Sir *et al.*¹⁹ explored objective functions in which each voxel is weighted by a coverage probability, i.e., a probability that the voxel will be included in a target or OAR structure. In contrast, Yang *et al.*²⁰ and Witte *et al.*²¹ employed objective functions that are probabilistically weighted functions of biological metrics: equivalent uniform dose (EUD), tumor control probability (TCP), and normal tissue complication probability (NTCP).

In simple cases, Unkelbach and Oelfke²³ showed that the above two PTP approaches—i.e., the PWDD and PWOFF approaches—are mathematically related. More general optimization strategies which fall into neither of these solution categories are represented by Chu *et al.*²⁴ and Olafsson and Wright.²⁵ In general, there is a consensus that the above PTP approaches can compensate for the effect of random errors. However, it is not clear that they provide good solutions for systematic errors. Optimizing an average dose distribution or objective function does not prevent the occurrence of unacceptable dose distributions in some cases. In effect, it does not constrain the percentiles of the dose distributions that can be delivered to patients. It is for this reason that the current work adopts a percentile constraint on target coverage: adequate CTV coverage is required not just on average, but for 90% of setup errors.

II. METHODS AND MATERIALS

II.A. Prostate IMRT plans

This work describes a virtual treatment planning study utilizing 27 anonymized prostate IMRT plans taken from a Virginia Commonwealth University (VCU), Department of Radiation Oncology clinical database. These plans included 5 mm CTV-to-PTV margins, referred to in this work as the “default” margins. The present work varies the CTV-to-PTV margins to obtain “optimal” CTV-to-PTV margins, i.e., margins that give the desired coverage value Q . Coverage Q is defined to be the percentage of systematic setup errors that can be absorbed by a given plan. (See Sec. II C for assumptions regarding setup errors.)

Structure delineation followed a protocol used by our institution. Normal structures included the small intestines, rectum, and bladder. The rectum was contoured to the junc-

tion of the sigmoid colon. The plans include two clinical target volumes—the prostate CTV (CTV) and nodal CTV (CTVNodes)—with corresponding PTV and PTV Nodes. To derive the CTV, the prostate was expanded by 5 mm in three dimensions, except posteriorly where there was zero expansion, to create the CTVProstate. The seminal vesicles CTV (CTVSV) was contoured from the prostate- seminal vesicles interface superiorly for 20 mm. The CTVProstate and CTVSV were combined to form the CTV. (These CTV structures remained unchanged throughout the margin iteration procedure described in Secs. II D and II E.)

In the original (default) plans, the PTV was derived by expanding the CTV by 5 mm in three dimensions. CTVNodes was derived by expanding the external and internal iliac vessels of 6 mm superiorly and inferiorly and 10 mm radially, excluding small bowel, bladder, rectum, and the PTV. CTVNodes was expanded by 5 mm in three dimensions, excluding the PTV, to derive PTVNodes.

The NonPTVSmallBowel, NonPTVRectum, and NonPTVBladder were derived by excluding the portion of these OARs which overlapped either the PTV or PTVNodes. The periprostatic rectum (PPRectum) is that length of rectum spanning from the superior to inferior aspect of the PTV. A NonPTVPPRectum was delineated as well following the same method for determining the other NonPTV contours. The PosteriorRectalWall comprised the posterior first/ third of the PPRectum.

Initial computed tomograph (CT) image sets for these patients had 1 mm voxel size in the x (left-right) and y (anterior-posterior) directions, and 3 mm in the z (inferior-superior) direction. The margin iteration described below required margins to be adjusted by submillimeter increments, in order to make small adjustments to CTV coverage. The z voxel size of 3 mm resulted in granular behavior of the PTV and TV, leading to discontinuous jumps in CTV coverage. To reduce the impact of this problem, the voxel size was reduced to 1 mm in the z -direction, and contours of all structures were interpolated onto the finer 1 mm³ CT grid prior to inverse planning and margin iteration.

II.B. IMRT inverse planning

IMRT inverse planning was performed with the VCU-IMRT optimization algorithm²⁶ interfaced to the Pinnacle treatment planning system (Philips Medical Systems, Fitchburg WI). Plans were optimized using deliverable-based optimization, which includes the leaf motion constraints and the effects of radiation leakage from the MLC in the optimization.^{27,28} Dynamic multileaf collimator (DMLC) delivery was presumed during optimization. Intensity modulation was incorporated into Pinnacle³ using a transmission compensator matrix. Pinnacle’s adaptive convolution algorithm, which has been shown to adequately account for tissue heterogeneities,²⁹ was used for all dose computations. Planned treatments consisted of 28 fractions of simultaneous integrated boost (SIB) IMRT. Beam configuration consisted of seven equispaced 18 MV beams.

The prescription doses were 65 Gy to the 98% of the PTV and 50.4 Gy to 95% of the PTVNodes. (These dose levels were chosen since, in our institution, patients receive an initial 9 Gy of high dose rate brachytherapy.) At each margin iteration, the following dose-volume criteria were used for inverse planning. The PTV had dual criteria with a lower limit (LL) of 65.5 Gy to 99.9% of its volume with a penalty (relative importance) of 50, and an upper limit (UL) of 72 Gy to no more than 0.01% of its volume with a penalty of 30. The PTVNodes had a LL of 50.8 Gy to 98% of its volume with a penalty of 45; and a dual UL of no more than 40% and 5% of its volume receiving 55 and 65 Gy, with penalties of 20 and 40, respectively.

OAR dose criteria were consistent with TD5/5 doses in the published literature.^{30,31} The NonPTVRectum was limited to no more than 40% and 1% of its volume receiving 40 and 65 Gy, respectively, with penalties of 25 and 30, respectively. The NonPTVPPRectum was limited to no more than 40% and 5% of its volume receiving 45 and 65 Gy, respectively, with penalties of 30 and 40, respectively. The NonPTVBladder was limited to no more than 50% and 1% of its volume receiving 45 and 60 Gy, respectively, with penalties of 20 and 30, respectively. The NonPTVSmallBowel was limited to no more than 30% and 1% of its volume receiving 40 and 50 Gy, respectively, with penalties of 30 and 30, respectively. The PostRectalWall was limited to no more than 10% of its volume receiving 45 Gy with a penalty of 25.

II.C. Setup errors

This work considered random and systematic interfraction setup errors. It did not consider intrafraction motion or structure delineation errors. Errors were assumed to be translational and Gaussian, with systematic (Σ) and random (σ) standard deviations (SDs) of 2 mm along all axes.

[SDs appropriate to the institution and treatment site should be employed. SDs of 2 mm approximate current levels of interfraction setup errors at our institution (unpublished data) when prostate patients are localized using implanted markers. They are consistent with published data from other institutions.^{32,33} The method proposed in this work can be used with any values of σ and Σ .]

Random and systematic errors were handled differently. The 28 treatment fractions were assumed to be sufficient for random errors to produce blurring of the dose distribution, which was modeled using the convolution method.^{7,34} During IMRT inverse planning, fluence was convolved with a normal probability distribution function (PDF) having SD $\sigma=2$ mm. Consequently, the inverse planning process itself compensated for random interfractional variations by expanding the treated volume by an appropriate margin.

Systematic errors result in a displacement, rather than a blurring, of the dose distribution, and so cannot be handled in this way. Plans robust to systematic errors were obtained by using margin iteration to ensure that the treated volume encompassing the CTV was large enough to absorb the

specified fraction (i.e., 90%) of systematic errors that a patient could incur. Margin iteration is described in detail in Sec. II D below.

Note that the default 5 mm CTV-to-PTV margins correspond exactly to the systematic component of the van Herk margin formula.³ In this work, because random errors are handled via fluence convolution, the CTV-to-PTV margin is required to absorb systematic errors only. For this case, the van Herk margin formula requires a margin of $M=2.5 \Sigma=5$ mm.

II.D. Coverage-based treatment planning

Coverage-based treatment planning (CBTP) is achieved by means of margin iteration, i.e., iteratively adjusting the size of the PTV and TV to achieve target coverage that is as close as possible to the specified value. In this study, the TV surface is the PTV minimum dose isodose surface, so the size of the TV is determined by the CTV-to-PTV margin. Assuming that the PTV minimum dose is achieved somewhere on the PTV surface, the TV boundary encompasses the PTV, necessarily touching the PTV at one or more points on its surface, but generally extending beyond the PTV in other directions (see Fig. 1).

Because the TV touches the PTV surface, the CTV-to-PTV margin determines the minimum size of the TV. Consequently, the CTV-to-PTV margin is used to adjust the overall size of the TV, and therefore the dosimetric margin distribution. As the CTV-to-PTV margin grows and shrinks, the TV will grow and shrink with it, although not necessarily in direct proportion to the CTV-to-PTV margin.

Margin iteration is illustrated at a high level in Fig. 2. At each iteration one selects a value for the CTV-to-PTV margin and expands the CTV and CTVNodes by this margin to obtain the PTV and PTVNodes. (In our implementation the iterative loop was controlled by a TCL script. At the start of each iteration this script deleted the pre-existing PTVs and regenerated new PTVs using the relevant CTV-to-PTV margin.) Standard IMRT inverse planning is then performed to obtain the dose distribution corresponding to the CTV-to-PTV margin.

Having obtained the dose distribution, one finds the PTV minimum dose and the TV, from which the DMD and coverage Q are computed. If the calculated CTV coverage is within tolerance of the desired value, the current plan is optimal and the iteration stops. If not, one deletes the current PTVs, adjusts the CTV-to-PTV margin, creates new PTVs, and repeats the above steps.

II.E. Margin iteration implementation details

The implementation started with a lower bound $M_{\min}=2$ mm and an upper bound $M_{\max}=5$ mm for the CTV-to-PTV margin. Other starting values could be used, provided the resultant coverage values $Q_{\min}$ and $Q_{\max}$ bracket the specified coverage (90%). The value $M_{\min}=2$ mm was found to give $Q_{\min}<90\%$ for all plans considered here. The van Herk margin formula ensures that $M_{\max}=5$ mm gives $Q_{\max}\geq 90\%$.

FIG. 2. Flow diagram of margin iteration. The CTV-to-PTV margin is iteratively adjusted to achieve primary prostate PTV coverage that is close to the specified coverage criterion.

At each iteration, IMRT optimization was performed with a midpoint CTV-to-PTV margin $M_{\text{mid}} = (M_{\text{min}} + M_{\text{max}})/2$, and the corresponding coverage value Q_{mid} was calculated. If Q_{mid} was within 1% of Q_{opt} , M_{mid} was deemed to be optimal and the iteration stopped. Otherwise, M_{max} was set equal to M_{mid} if $Q_{\text{mid}} > Q_{\text{opt}}$ or M_{min} was set equal to M_{mid} if $Q_{\text{mid}} < Q_{\text{opt}}$. In this way the gap between M_{min} and M_{max} was progressively halved, while ensuring that the optimal CTV-to-PTV margin M_{opt} remained sandwiched between the two: $M_{\text{min}} \leq M_{\text{opt}} \leq M_{\text{max}}$. Up to five iterations were permitted. By the fifth iteration, the gap between M_{min} and M_{max} was less than 0.1 mm. At that point, if neither Q_{min} nor Q_{max} was within 1% of Q_{opt} , M_{opt} was set equal to M_{max} .

When expanding the primary and nodal CTVs to obtain the primary and nodal PTVs, care was taken to ensure that no gaps appeared between the primary and nodal PTVs. In a given iteration, the nodal CTV excludes the primary PTV. If one subsequently shrinks the primary PTV, but uses the same nodal CTV to obtain the nodal PTV, it is possible for a gap to occur between the primary and nodal PTVs. To avoid this problem, it is necessary to start with a nodal CTV that extends to the surface of the primary CTV, expand to obtain the PTVs, and then exclude the primary PTV from both the nodal CTV and nodal PTV.

II.F. DMD and coverage calculation

This section briefly describes the method of DMD and coverage calculation. The reader is referred to Gordon and Siebers⁵ for details. The dosimetric margin $D(\phi, \theta)$ is the distance that the CTV can move in direction (ϕ, θ) —where ϕ is elevation and θ azimuth—before it intersects the TV. Figures 1(a) and 1(b) illustrate the dosimetric margins around the primary and nodal CTVs. The nodal TV encloses both the primary and nodal PTVs.

Relevant structures were exported from the treatment planning system in the form of triangular meshes. A computational geometry algorithm was then used to calculate collision distances between the meshes. Collision distances (i.e.,

dosimetric margins) were calculated for 10° increments of ϕ from -90° to 90° , and 10° increments of θ from -180° to $+180^\circ$. Allowing for degeneracy of θ at $\phi = -90^\circ$ and $\phi = +90^\circ$, this gave dosimetric margins in 614 evenly spaced directions covering 4π sr. These 614 measurements constituted the dosimetric margin distribution (DMD).

Because systematic errors are assumed to be Gaussian with identical SD along all axes, systematic displacements d occur isotropically with the following three-dimensional normal cumulative distribution function (CDF) having standard deviation Σ ,

$$\Pr[d \leq r] = \text{normcdf3D}(r, \Sigma) = \text{erf}(r/\sqrt{2}\Sigma) - (2/\sqrt{\pi}) \cdot (r/\sqrt{2}\Sigma) \cdot \exp(-r^2/2\Sigma^2), \quad (1)$$

where $\text{erf}()$ is the error function. Coverage Q is the percentage of errors that are absorbed by the dosimetric margin. It follows that Q is given by the formula

$$Q = \int \int d\phi d\theta \text{normcdf3D}(D(\phi, \theta), \Sigma), \quad (2)$$

where the double integral is in reality a normalized sum over the 614 values of $D(\phi, \theta)$. Our description—here and in Sec. II D—assumes that the PTV minimum dose occurs on the surface of the PTV. This was found to be the case for the 27 plans analyzed here. The proposed methods of margin iteration and DMD calculation can be adapted to the case where the PTV minimum dose occurs at an interior point, although details are not given here.

II.G. Plan evaluation using percentile DVHs

Plans were evaluated in terms of (a) the reductions in margins and treated volumes produced by coverage-based planning and (b) the resultant DVHs. Conventional DVHs summarize the *planned* dose distribution, i.e., the dose distribution produced by IMRT inverse planning. In this work, the planned dose distribution included the effect of random setup errors (via fluence convolution) but excluded the effect

of systematic setup errors. In order to include the effect of systematic errors, this work utilized *percentile DVHs*, which were generated as follows.

After margin iteration, the effect of systematic errors on each optimized plan was simulated by shifting the dose distribution (relative to the anatomy) by 50 randomly sampled systematic offsets. (The component of each systematic offset along each 3D axis was a random value sampled from a normal distribution with mean zero and SD $\Sigma=2$ mm. Offsets were translations. Rotations were not modeled.) The doses resulting from these systematic shifts are referred to as *realized* doses to distinguish them from planned doses. Each realized dose distribution corresponds to the dose distribution that would be realized in the patient for the modeled systematic error.

Conventionally, D_v denotes the planned dose received by fractional volume v of a structure and V_d denotes the fractional volume receiving planned dose d . The (planned) DVH is a plot of v versus D_v . For systematic offset i , $D_{v,i}$ denotes the realized dose received by fractional volume v of a structure and $V_{d,i}$ denotes the fractional volume receiving realized dose d . For each systematic offset the realized DVH is a plot of v versus $D_{v,i}$.

The quantity $D_{v,p}$ denotes the p th percentile of the random variable $D_{v,i}$. For conciseness, it is convenient to interpret $D_{v,p}$ differently for target and OAR structures. For target structures (i.e., CTVs), $D_{v,p}$ is the value such that p percent of realized doses $D_{v,i}$ fall above $D_{v,p}$. For OAR, $D_{v,p}$ is the value such that p percent of realized doses $D_{v,i}$ fall below $D_{v,p}$. These definitions are consistent with general expectations regarding target and OAR doses: for targets one wants to be sure that realized doses fall above some threshold most of the time, whereas for OAR one wants to be sure that realized doses fall below some threshold most of the time.

A percentile DVH is a plot of v versus $D_{v,p}$. In this work, percentile DVHs were obtained by calculating $D_{v,p}$ for values of v from 0 to 1 in steps of 0.01. For each value of v , $D_{v,p}$ was calculated by ordering the 50 samples $D_{v,i}$ and selecting (or interpolating) the p th percentile. In the following, the percentile p usually matches the coverage criterion Q . If a plan is generated that provides target coverage for $Q=90\%$ of systematic errors, it makes sense to quantify the performance of that plan in terms of doses $D_{v,90}$, volumes $V_{d,90}$, and DVHs that are valid for 90% of errors. Additional plots of CTV $D_{98,95}$ and $D_{98,99}$ are given below (Fig. 8) to illustrate the effect of different p -values.

The fractional volume $FV_{d,p}$ is equal to $V_{d,p}$ divided by the volume of the relevant target structure. It is analogous to the ICRU conformity index, which is the ratio of the TV to the PTV volume.

III. RESULTS

III.A. Margins and treated volumes

Standard IMRT inverse planning and margin iteration was performed for all 27 patients, using a DMD-based coverage criterion of $Q=90\%$. The results of interest are the reduction in margins and treated volumes that can be achieved by

FIG. 3. Plot of margin versus patient number. M_{opt} ($\blacktriangle$) is the optimal CTV-to-PTV margin, DM_{opt} ($\square$) is the average dosimetric margin corresponding to M_{opt} , and $DM_{5\text{ mm}}$ ($\blacklozenge$) is the average dosimetric margin corresponding to default 5 mm CTV-to-PTV margins. Gray lines are average values.

meeting, rather than exceeding, the desired coverage criterion. Also of interest are the effect of margin reductions on target and OAR dose distributions. This section describes reductions in margins and treated volumes.

Figure 3 shows the optimal CTV-to-PTV margin M_{opt} (i.e., the margin produced by margin iteration) for each of the 27 patients. It also shows the average dosimetric margins DM_{opt} and $DM_{5\text{ mm}}$ corresponding to optimal and 5 mm margins, respectively. Across the 27 prostate patients, the optimal CTV-to-PTV margin for s remains close to the average value of 2.8 mm. This represents a reduction of about 40% relative to the default 5 mm CTV-to-PTV margin. For a CTV-to-PTV margin of 5 mm the average dosimetric margin is 8.8 mm. For optimal CTV-to-PTV margins it is 6.2 mm, a reduction of 30%. For both $M=5\text{ mm}$ and $M=M_{\text{opt}}$, the difference between the CTV-to-PTV margin and the average dosimetric margin reflects the imperfect conformance of the planned dose distribution to the PTV, as in Fig. 1.

Figure 4 plots the average dosimetric margin versus the CTV-to-PTV margin for 27 patients. The superimposed line is the least-squares fit. Some of the per-patient plots in Fig. 4 show that the dosimetric margin increases nonmonotonically with CTV-to-PTV margin. This behavior is discussed below. The overall significance of Fig. 4 is to show a linear relationship between the dosimetric margin and the CTV-to-PTV margin. This type of relationship could potentially be used to correct existing margin formulas to account for imperfect conformance of the planned dose distribution to the PTV. (This type of correction would be irrelevant in a mature implementation of probabilistic planning, where implicit margins are added to the planned dose distribution by the inverse planning algorithm.)

Figure 5 plots the coverage values achieved for the primary (prostate) CTV with default 5 mm CTV-to-PTV margins ($Q_{5\text{ mm}}$), and with CTV-to-PTV margins that are opti-

FIG. 4. Plots of average dosimetric margin versus CTV-to-PTV margin for 27 patients. The superimposed line is the least-squares fit.

mized to meet a coverage criterion of 90% (Q_{opt}). With 5 mm margins, the coverage values are all close to 100%. Recall that 5 mm is the value given by the van Herk margin formula to achieve $Q=90\%$. Conventional practice is to set the CTV-to-PTV margin equal to the van Herk margin formula, in the expectation that this will produce target coverage of 90%. In fact, due to imperfect conformance of the TV to the PTV, setting the CTV-to-PTV margin equal to 5 mm produces coverage values close to 100% for these patients. Figure 5 illustrates that there is room to shrink the CTV-to-PTV margin below 5 mm, while still achieving $Q \geq 90\%$.

Ideally, all of the optimized coverage values Q_{opt} would be exactly equal to 0.90. However, the finite size of CT vox-

FIG. 5. Coverage values achieved for the primary CTV with default 5 mm CTV-to-PTV margins ($Q_{5\text{ mm}}$, $\blacklozenge$), and with CTV-to-PTV margins that are optimized to meet a coverage criterion of 90% (Q_{opt} , $\blacksquare$).

FIG. 6. Percent change in planned primary PTV minimum dose plotted against percent change in primary treated volume.

els causes granularity in the coverage values, e.g., coverage can jump from 85% to 95% when the margin increases by only 0.1 mm. (See Sec. IV below.) Because of this granularity, it is generally not possible in our study to find a value of the CTV-to-PTV margin that gives coverage of exactly 0.90. Figure 5 shows that achievable coverage values of range from 0.89 to 0.96, with an average of 0.92.

Figure 6 shows the percent change in the planned primary PTV minimum dose plotted against the percent change in the primary treated volume, when the CTV-to-PTV margin is reduced from 5 mm to M_{opt} . Intuitively, reduction in the CTV-to-PTV margin should reduce the dosimetric margin and the treated volume and leave the planned minimum PTV dose either unchanged or slightly increased due to the reduced impact of OAR constraints. Figure 3 shows that the dosimetric margin is indeed reduced for all patients. The percent change in DM ranges from -43% to -18% . However, Fig. 6 shows that TV changes range from -54% to $+45\%$ (positive values are counterintuitive), while changes in the PTV minimum dose range from -6% to $+9\%$ (negative values are counterintuitive). The counterintuitive results are due to the fact that the PTV D_{min} is not constant over the plans—see Sec. IV below.

An alternative way of quantifying the reductions in irradiated tissue is to look at the volume of tissue treated to the prostate prescription dose of 65 Gy and the nodal prescription dose of 50.4 Gy. In the remainder of this section, reductions in irradiated volumes are quantified not in terms of absolute volumes, but rather in terms of fractional volumes FV_d receiving dose d . FV_d is equal to the volume of tissue receiving dose d divided by the volume of the embedded target structure, i.e., CTV or CTVNodes. Figure 7 plots the fractional volumes receiving prescription doses of 65 Gy (primary) and 50.4 Gy (nodal), as the CTV-to-PTV margin is varied via margin iteration. The data in this figure are averaged over the 27 patients.

FIG. 7. Fractional volumes receiving prescription doses of 65 Gy (primary) and 50.4 Gy (nodal). CTV $FV_{65,90}$ is the 90th percentile of the CTV volume irradiated to 65 Gy, divided by the CTV volume. (See Sec. II for an explanation of the $V_{d,p}$ and $FV_{d,p}$ notation.) Total FV_{65} is the total volume irradiated to 65 Gy, divided by the CTV volume. CTVNodes $FV_{50.4,90}$ is the 90th percentile of the CTVNodes volume irradiated to 50.4 Gy, divided by the CTVNodes volume. Total $FV_{50.4}$ is the total volume irradiated to 50.4 Gy, divided by the CTVNodes volume.

Specifically, CTV $FV_{65,90}$ is the 90th percentile of the CTV volume that is irradiated to ≥ 65 Gy, divided by the CTV volume. As one would expect, this quantity is close to 1, indicating that almost all of the CTV was irradiated to 65 Gy. Total FV_{65} is the total volume of tissue irradiated to ≥ 65 Gy, divided by the CTV volume. As the CTV-to-PTV margin was reduced from 5 to 2 mm, this quantity decreased from 2.4 to 1.8. The average CTV volume is 105 cm^3 . In reducing the CTV-to-PTV margin from 5 mm to M_{opt} , the total volume irradiated to 65 Gy was reduced on average by

19.3% or about 48 cm^3 . The total volume of non-CTV tissue irradiated to 65 Gy was correspondingly reduced by 32.3%.

Similarly, CTVNodes $FV_{50.4,90}$ is the 90th percentile of the CTVNodes volume irradiated to ≥ 50.4 Gy, divided by the CTVNodes volume. Total $FV_{50.4}$ is the total volume irradiated to ≥ 50.4 Gy, divided by the CTVNodes volume. As the CTV-to-PTV margin was reduced from 5 to 2 mm, this quantity decreased from 3.7 to 3.4. The average CTVNodes volume is 427 cm^3 . In reducing the CTV-to-PTV margin from 5 mm to M_{opt} , the total volume irradiated to 50.4 Gy was reduced on average by 2.8% or about 45 cm^3 . The total volume of non-CTVNodes tissue irradiated to 50.4 Gy was correspondingly reduced by 3.8%.

Based on Fig. 7, reducing the CTV-to-PTV margin from 5 mm to M_{opt} reduces the volumes of non-target tissue irradiated at 65 Gy and 50.4 Gy. For 65 Gy the reduction is large (32%). For 50.4 Gy the reduction is small (4%). The reason for the difference in percent reduction between 65 and 50.4 Gy is unclear, but is perhaps because the IMRT optimization criteria (given above) permit a significant percentage of the PTVNodes to receive 65 Gy, without penalty. It appears that margin reductions can significantly reduce the volume of PTVNodes receiving 65 Gy, but have little effect on the volume of tissue outside PTVNodes that receives 50.4 Gy.

Table I gives the percent reduction in the total FV_{65} on a per patient basis. Reductions range from a minimum of 8.7% to a maximum of 28.6%, with an average of 19.3%. Table II gives the percent reduction in the total $FV_{50.4}$. Reductions range from a minimum of 2.4% to a maximum of 11.7%, with an average of 2.8%.

Table III lists the reduction in region of interest (ROI) volumes irradiated to various dose levels when the CTV-to-PTV margin is reduced from 5 mm to M_{opt} . Values in Table III are percentages of the total volume of the corresponding ROI. The dose levels included in Table III are those featured in the IMRT inverse planning criteria (see above). Note,

TABLE I. Table of fractional volumes treated to 65 Gy for default 5 mm CTV-to-PTV margins ($M_5 \text{ mm}$) and optimal CTV-to-PTV margins (M_{opt}). The total FV_{65} is the total volume of tissue treated to 65 Gy, divided by the volume of the prostate CTV. Data are presented for each of the 27 prostate plans. Δ is the percent decrease in FV_{65} in going from 5 mm to optimal CTV-to-PTV margins.

Pat No.	Total FV_{65} ($M_5 \text{ mm}$)	Total FV_{65} (M_{opt})	Δ (%)	Pat No.	Total FV_{65} ($M_5 \text{ mm}$)	Total FV_{65} (M_{opt})	Δ (%)
1	2.307	1.767	23.4	15	2.553	1.925	24.6
2	2.540	1.813	28.6	16	3.180	2.443	23.2
3	2.445	2.039	16.6	17	2.414	1.894	21.5
4	3.159	2.445	22.6	18	2.437	2.119	13.1
5	2.545	2.179	14.4	19	2.512	2.011	20.0
6	2.167	1.771	18.3	20	2.371	1.899	19.9
7	2.547	2.054	19.4	21	3.013	2.718	9.8
8	2.270	1.857	18.2	22	2.411	2.113	12.4
9	2.344	1.983	15.4	23	2.107	1.668	20.8
10	1.952	1.579	19.1	24	2.410	1.824	24.3
11	2.886	2.277	21.1	25	2.557	2.088	18.3
12	2.173	1.608	26.0	27	2.449	1.996	18.5
13	2.480	1.966	20.7	27	2.330	2.127	8.7
14	2.860	2.185	23.6				

TABLE II. Table of fractional volumes treated to 50.5 Gy for default 5 mm CTV-to-PTV margins ($M_{5\text{ mm}}$) and optimal CTV-to-PTV margins (M_{opt}). The total $FV_{50.4}$ is the total volume of tissue treated to 50.4 Gy, divided by the volume of the nodal CTV. Data are presented for each of the 27 prostate plans. Δ is the percent decrease in $FV_{50.4}$ in going from 5 mm to optimal CTV-to-PTV margins.

Pat No.	Total $FV_{50.4}$ ($M_{5\text{ mm}}$)	Total $FV_{50.4}$ (M_{opt})	Δ (%)	Pat No.	Total $FV_{50.4}$ ($M_{5\text{ mm}}$)	Total $FV_{50.4}$ (M_{opt})	Δ (%)
1	3.637	3.329	8.5	15	3.680	3.530	4.1
2	3.161	2.940	7.0	16	4.772	4.412	7.5
3	3.252	3.032	6.7	17	3.762	3.500	7.0
4	3.274	3.137	4.2	18	4.408	4.244	3.7
5	3.793	3.704	2.4	19	3.913	3.620	7.5
6	3.182	2.809	11.7	20	4.064	3.726	8.3
7	3.302	3.155	4.5	21	3.654	3.438	5.9
8	3.508	3.328	5.1	22	4.025	3.763	6.5
9	3.468	3.304	4.7	23	3.262	3.114	4.5
10	3.992	3.697	7.4	24	3.732	3.447	7.7
11	4.022	3.710	7.7	25	4.007	3.786	5.5
12	3.903	3.532	9.5	27	3.490	3.294	5.6
13	3.540	3.367	4.9	27	4.033	3.880	3.8
14	3.283	3.131	4.6				

however, that for critical structures, the IMRT planning criteria apply to the non-PTV parts of the ROIs. In contrast, the values in Table III are percentages of the whole ROI volume, not the non-PTV part. The values in Table III cannot therefore be directly compared with the target values in the IMRT planning criteria.

Table III shows that this method of reducing CTV-to-PTV margins to achieve 90% primary CTV coverage has little effect on OAR volumes irradiated by doses less than 60 Gy. However, it reduces volumes irradiated by higher doses. Notably, it reduces bladder $V_{60,90}$ by 5.6% on average (reductions for individuals vary from 1.7% to 10.6%), and periprostatic rectum $V_{65,90}$ by 4.9% on average (reductions for individuals vary from 0.9% to 12.3%).

Reductions in high doses to OAR are directly attributable to reductions in the volume of the PTV, and therefore in the overlap volume between the PTV and OAR. These reductions are unlikely to be as evident if one were, e.g., to consider doses to the non-PTV parts of the OAR. In that case one would see an increase in the volume of the non-PTV portion of the OAR, without much change in the overall dose

distribution given to the non-PTV portion. Evaluating dose on the whole OAR makes it clear that the volume of OAR being irradiated to a high dose has decreased.

III.B. Doses and DVHs

Figure 8 shows plots of planned PTV D_{98} and realized CTV $D_{98,p}$ versus the CTV-to-PTV margin. Plotted doses are averaged over the 27 patients and normalized to the 65 Gy prescription dose. CTV $D_{98,p}$ is plotted for $p=90\%$, 95% , and 99% , with $p=90\%$ being the value that is most relevant, since it matches the 90% coverage criterion.

The IMRT optimizer achieves a more or less constant value for the planned PTV D_{98} regardless of the CTV-to-PTV margin. The relationship between the realized CTV $D_{98,p}$ and the planned PTV D_{98} is determined by the size of the CTV-to-PTV margin. If the margin is too small to absorb systematic errors, CTV $D_{98,p}$ will fall short of PTV D_{98} . If the margin is too large, CTV $D_{98,p}$ will exceed PTV D_{98} . If the margin is well matched to systematic errors, CTV $D_{98,p}$

TABLE III. Reductions in the irradiated volumes when CTV-to-PTV margins are reduced from 5 mm to M_{opt} . Numbers in the table are differences in percent volumes $V_{d,90}$, i.e., the percent of the ROI volume irradiated to dose d . The “avg” columns gives the average reduction in $V_{d,90}$, while the “range” columns give the range of the reductions, taken over the 27 plans. Negative values represent increases in $V_{d,90}$ when CTV-to-PTV margins are reduced from 5 mm to M_{opt} . Results are for the whole ROI, rather than the non-PTV part of the ROI. (PP Rectum=periprostatic rectum, P rectal wall=posterior rectal wall).

ROI+volume	Avg (%)	Range (%)	ROI+volume	Avg (%)	Range (%)
CTVNodes $V_{55,90}$	0.1	-4.9: +6.9	Bladder $V_{45,90}$	3.4	+1.1: +7.1
CTVNodes $V_{65,90}$	0.0	-3.1: +1.8	Bladder $V_{60,90}$	5.6	+1.7: +10.6
Rectum $V_{40,90}$	0.2	-5.3: +5.3	Small Bowel $V_{40,90}$	0.9	-0.7: +7.3
Rectum $V_{65,90}$	2.5	+0.7: +5.6	Small Bowel $V_{50,90}$	1.7	-0.1: +11.3
PP rectum $V_{45,90}$	2.1	-1.0: +8.0	P rectal wall $V_{45,90}$	1.5	-2.7: +10.7
PP rectum $V_{65,90}$	4.9	+0.9: +12.3

FIG. 8. Plots of planned PTV D_{98} and realized CTV D_{98} versus the CTV-to-PTV margin for the primary target. Doses are averaged over the 27 patients and normalized to the 65 Gy prescription dose. The value of CTV $D_{98,p}$ is such that, in the presence of systematic errors, realized CTV D_{98} will exceed CTV $D_{98,p}$ for percentage of errors. CTV $D_{98,p}$ is plotted here for $p = 90\%, 95\%, 99\%$.

will equal PTV D_{98} . In this last case, the PTV DVH dose is truly a representative of the expected CTV DVH dose values when setup errors are accounted for.

Figure 8 shows that CTV $D_{98,90}$ is equal to PTV D_{98} at a CTV-to-PTV margin of around 2.5 mm. Based on Fig. 3, it was concluded above that the optimal CTV-to-PTV margin is 2.8 mm. These figures are consistent to within the level of accuracy of this study.

The curves for CTV $D_{98,95}$ and $D_{98,99}$ illustrate the tradeoff between the coverage criterion and the optimal CTV-to-PTV margin. In each case, the optimal CTV-to-PTV margin is the one which ensures that realized CTV D_{98} matches planned PTV D_{98} for the specified percentage of errors. The higher the required coverage probability, the greater the required margin.

Figure 9 shows percentile DVHs ($p=90\%$) for a single patient for two cases: optimal CTV-to-PTV margin (thick lines) and 5 mm CTV-to-PTV margin (thin lines). These DVHs are characteristic of patients exhibiting a large difference between the optimal and 5 mm DVHs. (The term “large difference” signifies 5%–10% difference in the volumes of one or more OAR receiving high (>50 Gy) doses). Based on visual inspection, 5 of the 27 patients showed similarly large DVH differences. Of the remaining patients, 15 showed moderate differences (1%–5% difference in the volumes of one or more OAR receiving high doses) and 7 showed small differences ($<1\%$ difference in the volumes of one or more OAR receiving high doses) between optimal and 5 mm DVHs.

The DVH results showed that, at least for some patients, use of optimal margins ($Q=90\%$) produces a significant reduction in higher doses to organs at risk. Based on the data,

FIG. 9. Percentile DVHs ($p=90\%$) for CTV-to-PTV margin M equal to M_{opt} (thick lines) and 5 mm (thin lines). This figure is representative of patients exhibiting large differences between the two sets of DVHs. (PP rectum = periprostatic rectum, PR wall = posterior rectal wall).

it appears possible that reductions in this type tend to be largest for patients who are hard up against the OAR DVH constraints.

IV. DISCUSSION

CBTP as proposed in this paper utilizes two important concepts. First, CBTP evaluates the coverage probability realized by a treatment plan instead of hoping that a coverage level is achieved through specification of a nominal CTV-to-PTV margin. In this study, the CBTP concept was found to reduce irradiated volumes. It is possible, however, that there may be cases where a margin formula results in lower coverage probabilities than are desired, requiring larger (not smaller) treatment volumes. Importantly, without explicit evaluation of the coverage probability for a given plan, the realized coverage probability will be unknown. To avoid gross undercoverage or overcoverage, it is logical that reporting of expected coverage probability should become standard treatment planning practice.

Second, the relationship between the requested CTV-to-PTV margin and the realized TV can be used to adjust the coverage probability. While in this paper an automatic iterative procedure was used, manual procedures could be utilized to yield equivalent clinical benefits. The methods developed, with some minor extensions as noted below, are applicable to arbitrary treatment sites, planning systems, and treatment techniques.

IV.A. General assumptions and limitations

This work is based on several assumptions. In the case of random setup errors, the work as presented is restricted to translational interfraction errors. Random setup errors are further assumed to be normally distributed with the same

standard deviation along each axis, and their effect is modeled using the convolution method. The work generalizes to other types of *random* errors—e.g., intrafractional motion, anisotropic, and rotational errors—provided that the effect of these errors can also be modeled using the convolution method or otherwise be incorporated into the dose distribution.

Systematic setup errors are likewise assumed to be translational and normally distributed, with the same standard deviation along each axis. These assumptions are embodied in Eq. (1). The work can be generalized to other types of *systematic* errors—e.g., anisotropic and rotational errors—by replacing Eq. (1) with an analogous equation derived from an alternative systematic error model.

Importantly, as implemented here the proposed method assumes shift invariance of the dose distribution. That is, it assumes the effect of systematic errors is to shift (i.e., translate or rotate) the entire dose distribution, without changing its isodose contours. This assumption, which is embodied in Eq. (2), has been shown to be reasonable for prostate treatments.^{6,35} However, shift invariance may not hold for treatment sites where there are tissue heterogeneities. In these cases, the proposed method may be inaccurate. Instead, recalculating dose for each systematic error would be required.

This work used a target coverage criterion of 90%. Specifically, after compensating for random setup errors via the convolution method, plans were required to be capable of absorbing 90% of systematic setup errors. The default 5 mm CTV-to-PTV margins that were employed in this study corresponded exactly with this criterion: the SD of systematic errors was $\Sigma = 2$ mm, and the corresponding systematic component of the van Herk margin formula, which is predicated on 90% coverage, is $M = 2.5 \Sigma = 5$ mm.

The proposed method generalizes to other coverage criteria. Suppose, for example, one were to adopt a 95% coverage criterion. In that case the proposed method would need to be evaluated by comparing the optimal margins that are produced by margin iteration with CTV-to-PTV margins that are consistent with 95% coverage. In the case of a 95% coverage criterion, the systematic component of the van Herk margin formula would change to $M = 2.8 \Sigma = 5.6$ mm. (The reader is referred to van Herk *et al.*³ for the underlying mathematics.) This work gives results only for a 90% coverage criterion. However, it is anticipated that the same qualitative results—reduction in CTV-to-PTV margins and high doses to normal tissues—would also be obtained for other coverage criteria. (See Fig. 8 for expected coverage-based planning results).

IV.B. Specific assumptions and limitations

Section III A noted that CTV coverage experienced discrete jumps when the CTV-to-PTV margin was changed only slightly. For example, coverage jumped from 85% to 95% when the margin increased by only 0.1 mm. This was caused by the finite CT grid size (1 mm³), in combination with the way that DMDs were calculated.

In our implementation DMDs were calculated by first converting isodose surfaces to ROIs, then exporting ROIs as triangular meshes. It appears that in the process of converting isodose surfaces to ROIs, Pinnacle modifies the isodose surface so that it follows CT voxel boundaries. This ensures that as the CTV-to-PTV margin is adjusted in submillimeter increments, voxels will either be wholly included or wholly excluded from the TV. As whole voxels are brought inside the TV coverage experiences significant jumps.

Because of this effect, it was not always possible to find a value of M for which Q was within 1% of Q_{opt} . In our implementation M_{opt} was set equal to the smallest value of M for which $Q \geq 0.99 Q_{\text{opt}}$ (where Q_{opt} was 0.9). This resulted in “optimized” coverage values ranging from 0.89 to 0.96 (Fig. 5). It may be possible to consistently achieve coverage values that are very close to the Q_{opt} by using a better method of DMD calculation, e.g., one that avoids the conversion of isodose surfaces to ROIs.

Figure 3 shows that when the CTV-to-PTV margin is reduced from 5 mm to M_{opt} , the average dosimetric margin is reduced for all patients. The percent change in DM ranges from -43% to -18% . However, Fig. 6 shows that for two patients the percent change in the TV is positive. Also, for a number of patients the percent change in the planned PTV minimum dose is negative. Intuitively, reduction in the CTV-to-PTV margin should reduce the treated volume and leave the PTV minimum dose either unchanged or slightly increased due to the reduced impact of OAR constraints.

These counterintuitive results can be explained as follows. Motivated by the van Herk margin formula³ and prior work,^{5,6} this work focused on the PTV minimum dose as the primary dose metric relating to coverage and correspondingly defined the TV to be the volume enclosed by the PTV minimum dose isodose surface. In contrast, IMRT inverse planning criteria were defined in terms of the primary PTV D_{98} (and nodal PTV D_{95}). Figure 8 shows that the value of the PTV D_{98} remains quite stable as the CTV-to-PTV margin is varied. However, Fig. 6 shows that PTV D_{min} varies significantly.

It appears that the counterintuitive values in Fig. 6 are due to instability of D_{min} as compared with D_{98} . Dose falls away sharply near the edge of the high dose region around the PTV. Even though IMRT inverse planning achieves values for D_{98} that are consistently close to the prescription dose, it appears that D_{min} can vary significantly due to dose fluctuations (or voxel granularity) right at the edge of the PTV. This variation produces the spurious fluctuations in D_{min} and the TV seen in Fig. 6.

This leads to the following conclusion. In retrospect, a better way to have implemented this work would have been to define the primary TV as the volume enclosed by the PTV D_{98} isodose surface, and then tune the CTV-to-PTV margin to ensure that this TV was of the correct size to provide the specified coverage. This would have ensured consistency between the margin iteration and IMRT inverse planning—both would have used PTV D_{98} as the critical dose metric—and would likely have produced more stable treated volumes. This provides guidance for future work.

The variations in $D_{\min}$ explain the nonmonotonicity observed in some plots of dosimetric margin versus CTV-to-PTV margin in Fig. 4. Variations in $D_{\min}$ will have the effect of spuriously expanding the treated volume (i.e., volume inside the $D_{\min}$ isodose surface). Such variations could occur for some values of the CTV-to-PTV margin, but not for others. This in turn will lead to erratic behavior of some of the plots in Fig. 4.

Finally, in our implementation the same CTV-to-PTV margin was used to expand both the primary and nodal CTVs. Specifically, CTV-to-PTV margin was adjusted to achieve the desired coverage of the primary CTV, without regard to coverage of the nodal CTV. Results in Gordon *et al.*⁶ showed that, for this set of patients, use of the same margin for the primary and nodal PTVs will produce very similar coverage of the primary and nodal CTVs. This observation was used here to justify CTV-to-PTV margin selection based on primary CTV coverage only.

However, it is conceivable that even with the same margin, primary and nodal CTVs could experience different levels of coverage due to the presence of different dosimetric margins around these structures. Consequently, a more complete implementation would optimize the margin so that the desired coverage is provided not only for the primary CTV but also for the nodal CTV. This would result in a trivial change to the margin iteration procedure described above, but would likely entail longer computing time for the margin iteration because the primary and nodal CTV-to-PTV margins would need to be separately varied.

IV.C. Applicability to other treatment sites and techniques

The proposed method of coverage-based treatment planning has been evaluated here for seven-field SIB IMRT prostate treatments, but is generally applicable to other treatment sites and treatment techniques. The principal advantage of the method is to exploit situations where the dosimetric margin around a CTV is larger than the CTV-to-PTV margin. In these cases, it is possible to reduce the size of the CTV-to-PTV margin while still maintaining the desired level of CTV coverage.

At this time it is not clear what treatment sites and techniques will produce dosimetric margins that are large relative to the CTV-to-PTV margin. The dosimetric margin may depend on whether the plan is 3D conformal or IMRT. It may also depend on factors such as the beam energy, number of treatment angles, beam arrangement (which may be constrained by the treatment site), the choice of IMRT optimization criteria, and the specific treatment technique (e.g., simultaneous integrated boost). More research is necessary to quantify the dosimetric margins for other treatment sites and techniques.

In particular, it is not clear whether CBTP would be of use for ungated lung treatments where breathing motion leads to an elongated PTV. In this case, the applicability of CBTP would depend on how well the planned dose distribution conforms to the anisotropic PTV. It is possible that CBTP

will be useful in the case of adaptive treatment. For instance, where a dosimetric margin has been “created” by random setup errors in previously delivered fractions, CBTP may be able to reduce the CTV-to-PTV margin used for the remaining fractions.

IV.D. Implications of anisotropic dosimetric margins

In current IMRT planning, a common practice is to use an isotropic CTV-to-PTV margin. This creates the incorrect perception that all points on the CTV surface are equally likely to be underdosed due to a setup error. However, in cases where the TV does not conform perfectly to the PTV, the dosimetric margin around the CTV is anisotropic, implying that certain points on the CTV surface are more likely than others to be underdosed. The proposed CBTP method does not alter this fact. Instead it recognizes it and, where possible, uses it to create a more optimal treated volume.

Typically, current planning does not discriminate between subparts of a single PTV. All subparts are implicitly assumed to be equally susceptible to disease progression, and so one attempts to deliver the prescription dose uniformly across the PTV. In this scenario, it makes no difference where the CTV is underdosed. Instead, one is concerned only with the overall probability of CTV underdosing. In future, if biological data show that certain subparts of a PTV are more sensitive to underdosing than others, the proposed CBTP method (together with conventional IMRT planning) may need to be re-evaluated. In that case, additional constraints might need to be placed on the margin optimization so that the more susceptible parts of a PTV are guaranteed a high probability of coverage.

A further subtlety associated with anisotropic margins has to do with biological fractionation effects. Due to random setup errors across a multifraction treatment course, CTV regions for which the dosimetric margin is small could conceivably experience a higher degree of dose variability—i.e., a higher degree of variability in the fractionated doses that contribute to the total cumulative dose—than regions for which the dosimetric margin is large. Depending on the biological properties of the CTV tissue—e.g., the values of the parameters in the linear-quadratic model—this variability could affect the tumor control probability. This issue remains a subject for future research. It was side stepped in this work by adopting the frequently-used assumption that the effect of random errors can be modeled using the convolution method. The convolution method effectively assumes that dose variability (due to random errors) is negligible.

V. CONCLUSIONS

A previous study by Gordon and Siebers⁵ showed that, for the set of 27 SIB IMRT prostate plans used as the starting point for this study, the average dosimetric margin is about twice as large as the CTV-to-PTV margin. In these plans, if one sizes the PTV conventionally the TV ends up being substantially larger than it needs to be in order for the plans to

absorb the specified percentage (90%) of setup errors. Conversely, one can shrink the PTV and still achieve the specified level of CTV coverage.

This work demonstrates an iterative method of plan optimization that converges to the CTV-to-PTV margin that yields the desired coverage probability. In this method, one iteratively adjusts the CTV-to-PTV margin until the treated volume matches the specified level of coverage. When this method was applied to the 27 prostate cases, the CTV-to-PTV margin was reduced from the conventional value of 5 mm (i.e., the value given by the systematic component of the van Herk margin formula) to an average value of 2.8 mm.

This reduction in PTV size produced a corresponding decrease in the volume of normal tissue receiving high dose. The total volume of tissue receiving ≥ 65 Gy was reduced on average by 19.3% or about 48 cm³. Individual reductions varied from 8.7% to 28.6%. The volume of bladder receiving ≥ 60 Gy was reduced on average by 5.6% (reductions for individuals varied from 1.7% to 10.6%), and the volume of periprostatic rectum receiving ≥ 65 Gy was reduced on average by 4.9% (reductions for individuals varied from 0.9% to 12.3%).

The general principles of CBTP are applicable to arbitrary treatment sites and delivery techniques. However, the benefits in each specific case are to be determined. Importantly, observed deviations between coverage implied by CTV-to-PTV margins and coverage achieved by a given treatment plan imply the need to perform coverage probability analysis on a per-plan basis.

^{a)} Author to whom correspondence should be addressed; Electronic mail: jgordon@mcvh-vcu.edu

¹ "Prescribing, recording and reporting photon beam therapy," ICRU Report No. 50, 1994 (unpublished).

² "Prescribing, recording and reporting photon beam therapy," ICRU Report 62, 2000 (unpublished).

³ M. van Herk, P. Remeijer, C. Rasch, and J. V. Lebesque, "The probability of correct target dosage: Dose-population histograms for deriving treatment margins in radiotherapy," *Int. J. Radiat. Oncol., Biol., Phys.* **47**, 1121–35 (2000).

⁴ M. van Herk, "Errors and margins in radiotherapy," *Semin. Radiat. Oncol.* **14**, 52–64 (2004).

⁵ J. J. Gordon and J. V. Siebers, "Evaluation of dosimetric margins in prostate IMRT treatment plans," *Med. Phys.* **35**, 569–575 (2008).

⁶ J. J. Gordon, A. J. Crimaldi, M. Hagan, J. Moore, and J. V. Siebers, "Evaluation of clinical margins via simulation of patient setup errors in prostate IMRT treatment plans," *Med. Phys.* **34**, 202–214 (2006).

⁷ W. A. Beckham, P. J. Keall, and J. V. Siebers, "A fluence-convolution method to calculate radiation therapy dose distributions that incorporate random set-up error," *Phys. Med. Biol.* **47**, 3465–3473 (2002).

⁸ T. C. Chan, T. Bortfeld, and J. N. Tsitsiklis, "A robust approach to IMRT optimization," *Phys. Med. Biol.* **51**, 2567–2583 (2006).

⁹ J. G. Li and L. Xing, "Inverse planning incorporating organ motion," *Med. Phys.* **27**, 1573–1578 (2000).

¹⁰ M. Birkner, D. Yan, M. Alber, J. Liang, and F. Nusslin, "Adapting inverse planning to patient and organ geometrical variation: Algorithm and implementation," *Med. Phys.* **30**, 2822–2831 (2003).

¹¹ A. Trofimov *et al.*, "Temporo-spatial IMRT optimization: Concepts, implementation and initial results," *Phys. Med. Biol.* **50**, 2779–2798 (2005).

¹² D. L. McShan, M. L. Kessler, K. Vineberg, and B. A. Fraass, "Inverse plan optimization accounting for random geometric uncertainties with a

multiple instance geometry approximation (MIGA)," *Med. Phys.* **33**, 1510–1521 (2006).

¹³ J. Unkelbach and U. Oelfke, "Inclusion of organ movements in IMRT treatment planning via inverse planning based on probability distributions," *Phys. Med. Biol.* **49**, 4005–4029 (2004).

¹⁴ J. Unkelbach and U. Oelfke, "Incorporating organ movements in inverse planning: Assessing dose uncertainties by bayesian inference," *Phys. Med. Biol.* **50**, 121–139 (2005).

¹⁵ J. Unkelbach and U. Oelfke, "Incorporating organ movements in IMRT treatment planning for prostate cancer: Minimizing uncertainties in the inverse planning process," *Med. Phys.* **32**, 2471–2483 (2005).

¹⁶ D. Maleike, J. Unkelbach, and U. Oelfke, "Simulation and visualization of dose uncertainties due to interfractional organ motion," *Phys. Med. Biol.* **51**, 2237–2252 (2006).

¹⁷ J. Lof, B. K. Lind, and A. Brahme, "An adaptive control algorithm for optimization of intensity modulated radiotherapy considering uncertainties in beam profiles, patient set-up and internal organ motion," *Phys. Med. Biol.* **43**, 1605–1628 (1998).

¹⁸ C. Baum, M. Alber, M. Birkner, and F. Nusslin, "Robust treatment planning for intensity modulated radiotherapy of prostate cancer based on coverage probabilities," *Radiother. Oncol.* **78**, 27–35 (2006).

¹⁹ M. Y. Sir, S. M. Pollock, M. A. Epelman, K. L. Lam, and R. K. Ten Haken, "Ideal spatial radiotherapy dose distributions subject to positional uncertainties," *Phys. Med. Biol.* **51**, 6329–6347 (2006).

²⁰ J. Yang *et al.*, "A new method of incorporating systematic uncertainties in intensity-modulated radiotherapy optimization," *Med. Phys.* **32**, 2567–2579 (2005).

²¹ M. G. Witte, J. van der Geer, C. Schneider, J. V. Lebesque, M. Alber, and M. van Herk, "IMRT optimization including random and systematic geometric errors based on the expectation of TCP and NTCP," *Med. Phys.* **34**, 3544–3555 (2007).

²² H. Rehbringer, C. Forsgren, and J. Lof, "Adaptive radiation therapy for compensation of errors in patient setup and treatment delivery," *Med. Phys.* **31**, 3363–3371 (2004).

²³ J. Unkelbach and U. Oelfke, "Relating two techniques for handling uncertainties in IMRT optimization," *Phys. Med. Biol.* **51**, N423–N427 (2006).

²⁴ M. Chu, Y. Zinchenko, S. G. Henderson, and M. B. Sharpe, "Robust optimization for intensity modulated radiation therapy treatment planning under uncertainty," *Phys. Med. Biol.* **50**, 5463–5477 (2005).

²⁵ A. Olafsson and S. J. Wright, "Efficient schemes for robust IMRT treatment planning," *Phys. Med. Biol.* **51**, 5621–5642 (2006).

²⁶ Q. Wu and R. Mohan, "Algorithms and functionality of an intensity modulated radiotherapy optimization system," *Med. Phys.* **27**, 701–711 (2000).

²⁷ J. V. Siebers, M. Lauterbach, P. J. Keall, and R. Mohan, "Incorporating multi-leaf collimator leaf sequencing into iterative IMRT optimization," *Med. Phys.* **29**, 952–959 (2002).

²⁸ N. Dogan *et al.*, "Improving IMRT dose accuracy via deliverable monte carlo optimization for the treatment of head and neck cancer patients," *Med. Phys.* **33**, 4033–4043 (2006).

²⁹ I. B. Mihaylov, F. A. Lerma, M. Fatyga, and J. V. Siebers, "Quantification of the impact of MLC modeling and tissue heterogeneities on dynamic IMRT dose calculations," *Med. Phys.* **34**, 1244–1252 (2007).

³⁰ B. Emami *et al.*, "Tolerance of normal tissue to therapeutic irradiation," *Int. J. Radiat. Oncol., Biol., Phys.* **21**, 109–122 (1991).

³¹ P. Rubin and G. W. Casarett, "Clinical radiation pathology as applied to curative radiotherapy," *Cancer* **22**, 767–778 (1968).

³² D. Litznerberg *et al.*, "Daily prostate targeting using implanted radiopaque markers," *Int. J. Radiat. Oncol., Biol., Phys.* **52**, 699–703 (2002).

³³ M. W. Lometti *et al.*, "On-line prostate localization using radiopaque markers and an electronic portal imaging device: Analysis of seed migration and interfraction gland rotation," *Int. J. Radiat. Oncol., Biol., Phys.* **69**, S367 (2007).

³⁴ J. J. Gordon and J. V. Siebers, "Convolution method and CTV-to-PTV margins for finite fractions and small systematic errors," *Phys. Med. Biol.* **52**, 1967–1990 (2007).

³⁵ M. van Herk, P. Remeijer, and J. V. Lebesque, "Inclusion of geometric uncertainties in treatment plan evaluation," *Int. J. Radiat. Oncol., Biol., Phys.* **52**, 1407–1422 (2002).