

A Hybrid Evolutionary Algorithm for Multiobjective Anatomy Based Dose Optimization in HDR Brachytherapy

M Lahanas¹, D Baltas^{1,2}, and N Zamboglou^{1,2}

¹Department of Medical Physics & Engineering, Strahlenklinik, Klinikum Offenbach, 63069 Offenbach, Germany.

²Department of Electrical and Computer Engineering, National Technical University of Athens, 15773 Zografou, Athens, Greece.

Corresponding author:

Michael Lahanas, Ph.D.
Dept. of Medical Physics & Engineering
Strahlenklinik
Klinikum Offenbach
Starckenburgring 66
63069 Offenbach am Main, Germany
Tel.: +49 – 69 – 8405 –3235
Fax. : +49 – 69 – 8405 –4481 or –864480
E-mail: mlahanas@gmx.de

ABSTRACT

Multiple objectives must be considered in anatomy based dose optimization for high dose rate (HDR) brachytherapy and a large number of parameters must be optimized to satisfy often competing objectives. For objectives expressed solely in terms of dose variances, deterministic gradient-based algorithms can be applied and a weighted sum approach is able to produce a representative set of non-dominated solutions. As the number of objectives increases, or non-convex objectives are used, local minima can be present and deterministic or stochastic algorithms such as simulated annealing either cannot be used or are not efficient. In this case we employ a modified hybrid version of the multiobjective optimization algorithm *NSGA-II*. This, in combination with the deterministic optimization algorithm produces a representative sample of the Pareto set. This algorithm can be used with any kind of objectives, including non-convex, and does not require artificial importance factors. A representation of the trade-off surface can be obtained with more than 1000 non-dominated solutions in 2-5 minutes. An analysis of the solutions provides information of the possibilities available using these objectives. Simple decision making tools allow the selection of a solution that provides a best fit for the clinical goals. We show an example with a prostate implant and compare results obtained by variance and DVH based objectives.

1. INTRODUCTION

In modern interstitial HDR brachytherapy, the dose distribution is evaluated with respect to the irradiated normal tissue (NT), the planning target volume (PTV). Additionally, for organs at risk (OAR) within or adjacent to the PTV the dose must be smaller than a critical dose value D_{crit} . In practice, it is difficult if not impossible to meet all these objectives simultaneously. Some of the algorithms only take into account the catheter geometry and ignore the geometry of the PTV and OARs. Doses in adjacent OARs are reduced, by trial and error, by manually changing the dwell times of dwell positions in their neighborhood. In the last few years the anatomy of the PTV, and in some cases of the surrounding tissue including OARs, have been considered using various objective functions (Lahanas *et al* 1999, Lessard and Pouliot 2001, Milickovic *et al* 2002). These algorithms now are described as *anatomy based*. Single objective optimization algorithms transform the multiobjective problem into a single objective problem by combining the objectives into a single objective, e.g. by using a weighted sum of all objectives. The weights used, known also as *importance factors*, are considered as a measure of the significance of each objective in the optimization process. Even if the solution obtained is a global optimal solution, *i.e.* the best possible for the aggregate single objective function, it is possible that by using other importance factors, other better solutions can be obtained. With this approach all the available information on the alternative possible solutions is lost. Until now these factors have been obtained empirically although the method of their choice has never been discussed.

One goal of a treatment planning system should be the ability to have a fast response to assist the clinician in obtaining good plans. Also it should provide all the information on the possibilities given the objectives of the treatment. In order to explore the feasible region of the solution space with respect to each objective, different values for the importance factors in the aggregate objective function must be given, and then the results assessed. Furthermore, the appropriate values of these importance factors differ from case to case. This implies that for any new clinical case a large amount of effort is necessary for their determination. While current optimization methods are single objective, the dose optimization problem is a true multiobjective problem and therefore multiobjective optimization methods should be used. Most realistic problems require the simultaneous optimization of many objectives. It is unlikely that all objectives are optimal for a single set of parameters. If this is so, then in principle there exist an infinite number of solutions. A multiobjective algorithm

does not provide a single solution but a representative set of all possible solutions. Out of these representative solutions a single final solution has to be selected.

For variance-based objectives, gradient-based deterministic algorithms due to their efficiency allow the construction of the so-called Pareto or trade-off front. This contains all the information on the degree of the competition between the objectives and which is necessary for the treatment planner to select the solution which best fulfills the dose optimization goals. For convex objective functions according to the conditions of the Kuhn-Tucker theorem (Miettinen 1999), a global Pareto optimum can be obtained and the entire Pareto front is accessible using the weighted sum method. A problem of this method, as found in most of the conventional dose optimization algorithms, is that it cannot provide solutions in possible non-convex parts of the Pareto trade-off front because a convex weighted sum of objectives converges only to the convex parts of the Pareto front. Additionally it can fail if local minima exist.

For DVH based objectives, deterministic algorithms cannot be used. We modified for this reason the genetic multiobjective optimization algorithm *NSGA-II* combining it with a deterministic algorithm increasing its performance. This algorithm is not limited to convex objective spaces and does not require importance factors. Additionally, we use an external population for the accumulation of non-dominated solutions. We obtain 1000-2000 non-dominated solutions in 2-5 minutes. Comparisons of the optimization results were made using the Nucletron treatment planning system PLATO BPS 13.7.

2. METHODS

2.1. Multiobjective optimization.

The multiobjective optimization (*MO*) problem (also called multicriteria optimization or vector optimization) (Miettinen 1999) can be defined as the problem of determining the following. "A vector of decision variables which satisfies constraints and optimizes a vector function whose elements represent M objective functions. These functions form a mathematical description of performance criteria that are usually in conflict with each other. Hence, the term optimize means finding a solution which would give the values of all objective functions acceptable to the treatment planner."

We call decision variables x_j , $j=1,2,\dots,N$ for which values are to be chosen in an optimization problem. In order to know how "good" a certain solution is we need to have some criteria for evaluation. These criteria are expressed as computable functions $f_1(\mathbf{x}),\dots,f_M(\mathbf{x})$ of the decision variables, which are called *objective functions*. These form a vector function \mathbf{f} . In general, some of these will be in conflict with others, and some will have to be minimized while others are maximized. The multiobjective optimization problem can be now defined as the problem to find the vector $\mathbf{x}=(x_1,x_2,\dots,x_N)$, *i.e.* solution which optimize the vector function \mathbf{f} , subject to the following conditions:

$$\begin{aligned} J \text{ inequalities } g_j(\mathbf{x}) &\geq 0, j = 1, 2, \dots, J \\ K \text{ equality constraints } h_k(\mathbf{x}) &= 0, k = 1, 2, \dots, K. \end{aligned} \quad (1)$$

The constraints define the feasible region \mathbf{X} and any point \mathbf{x} in \mathbf{X} defines a feasible solution. The vector function $\mathbf{f}(\mathbf{x})$ is a function that maps the set \mathbf{X} in the set \mathbf{F} that represents all possible values of the objective functions. Normally we never have a situation in which all the $f_i(\mathbf{x})$ values have a optimum in \mathbf{X} at a common point \mathbf{x} . We therefore have to establish certain criteria to determine what would be considered a "optimal" solution. One interpretation of the term optimum in multiobjective optimization is the Pareto optimum.

A solution \mathbf{x}_1 dominates a solution \mathbf{x}_2 if the two following conditions are true:

- 1) \mathbf{x}_1 is no worse than \mathbf{x}_2 in all objectives, *i.e.* $f_j(\mathbf{x}_1) \leq f_j(\mathbf{x}_2) \forall j=1,2,\dots,M$
- 2) \mathbf{x}_1 is strictly better than \mathbf{x}_2 in at least one objective, *i.e.* $f_j(\mathbf{x}_1) < f_j(\mathbf{x}_2)$ for at least one $j \in \{1,2,\dots,M\}$

We assume, without loss of generality, that this is a minimization problem. \mathbf{x}_1 is said to be non-dominated by \mathbf{x}_2 or \mathbf{x}_1 is *non-inferior* to \mathbf{x}_2 and \mathbf{x}_2 is dominated by \mathbf{x}_1 . Among a set of solutions P , the non-dominated set of solutions P' are those that are not dominated by any other member of the set P . When the set P is the entire feasible search space then the set P' is called the *global Pareto optimal set*. If for every member \mathbf{x} of a set P there exists no solution in the neighborhood of \mathbf{x} then the solutions of P form a *local Pareto optimal set*. The image of the Pareto optimal set is called the *Pareto front*.

2.2 Weighted sum approaches

The weighted sum method converts the multiobjective problem of minimizing the vector \mathbf{f} into a scalar problem by forming a weighted sum of all objectives. The weighting coefficients do not necessarily correspond directly in a linear way to the importance of the objectives. The

problem is to minimize $f(\mathbf{x}) = \sum_{m=1}^M w_m f_m(\mathbf{x})$ subject to the conditions (Eq. 1). If \mathbf{x}^* is a Pareto

optimal solution then there exists a weight vector $\mathbf{w} = (w_1, w_2, \dots, w_M)$, such that \mathbf{x}^* is a solution of the multiobjective convex optimization problem. For two objectives the weighted

sum is given by $y = w_1 f_1(\mathbf{x}) + w_2 f_2(\mathbf{x})$, i.e. $f_2(x) = -\frac{w_1}{w_2} f_1(x) + \frac{y}{w_2}$. The minimization of the

weighted sum can be interpreted as finding the value of y for which the line with slope $-w_1/w_2$ just touches the boundary of \mathbf{F} as it proceeds outwards from the origin. It is therefore not possible to obtain solutions on non-convex parts of the Pareto front with this approach.

A representative convex part of the Pareto set can be sampled by running a single objective optimization algorithm each time with a different vector of importance factors. We call a set of importance factors vectors *normalized and uniformly distributed* if each importance factor of each objective takes one of the following values: $[l/k, k = 0, 1, \dots, k]$,

where k is the sampling parameter and $\sum_{j=1}^M w_j = 1, w_j \geq 0, \forall j$ for each vector \mathbf{w} . For M

objectives we have $\binom{M+k-1}{M-1}$ vectors. Each optimization run produces a new and distinct

nondominated point. The produced solutions will not necessary be uniformly distributed on the Pareto front (Das and Dennis 1997).

Single objective optimization algorithms provide in the ideal case only one Pareto optimal solution in one optimization run. Multiobjective evolutionary algorithms provide a representative set of non-dominated solutions and are ideally suited for true multiobjective optimization.

2.3 Objective functions used in HDR brachytherapy dose optimization.

The aim of HDR brachytherapy dose optimization is to cover the PTV with at least some specified dose value, and to protect the OARs and the surrounding normal tissue with dose values above some specific level. The objectives functions which are used have a common characteristic in that they penalize dose values above and or below these limits. These objectives can be expressed as combinations of objective functions f_L and f_H of the form:

$$f_L(\mathbf{x}) = \frac{1}{N} \prod_{i=1}^N \Theta(D_L - d_i(\mathbf{x}))(D_L - d_i(\mathbf{x}))^\alpha \quad (2)$$

$$f_H(\mathbf{x}) = \frac{1}{N} \prod_{i=1}^N \Theta(d_i(\mathbf{x}) - D_H)(d_i(\mathbf{x}) - D_H)^\alpha \quad (3)$$

where N is the number of sampling points, D_L and D_H are low and high dose limits, $d_i(\mathbf{x})$ is

$$\begin{cases} 1 & x > 0 \\ 1/2 & x = 0 \\ 0 & x < 0 \end{cases}$$

the dose of the i^{th} sampling point, $\Theta(x) = \begin{cases} 1/2 & x = 0 \\ 0 & x < 0 \end{cases}$ and α a parameter.

$$\begin{cases} 1 & x > 0 \\ 1/2 & x = 0 \\ 0 & x < 0 \end{cases}$$

We obtain the objectives functions proposed for the inverse dose optimization problem for prostate implants in Lessard and Pouliot 2001 by using $\alpha=1$.

For $\alpha=0$ we obtain a set of objectives which we call DVH based objectives. The advantage of using these objectives is that the results in terms of objective values are more intuitive to understand than dose variances or objectives given by Lessard and Pouliot. We use the following set of objectives:

$$f_L^{PTV} = \frac{1}{N_{PTV}} \prod_{i=1}^{N_{PTV}} \Theta(D_L^V - d_i) \text{ as the fraction of PTV with } D < D_L^V$$

$$f_H^{PTV} = \frac{1}{N_{PTV}} \prod_{i=1}^{N_{PTV}} \Theta(d_i - D_H^V) \text{ as the fraction of PTV with } D > D_H^V$$

$$f_{OAR}^j = \frac{1}{N_{OAR}^j} \prod_{i=1}^{N_{OAR}^j} \Theta(d_i - D_{crit}^j) \text{ as the fraction for the } j^{\text{th}} \text{ OAR with } D > D_{crit}^j.$$

$$f_{NT} = \frac{1}{N_{NT}} \prod_{j=1}^{N_{NT}} D_j^2 \text{ as the average squared dose in the surrounding normal tissue.}$$

D_L^V is the prescription dose, or lower dose limit and D_H^V is the high dose limit in the PTV.

f_{NT} is the mean quadratic dose in the surrounding normal tissue. N_{PTV} , N_{NT} and N_{OAR}^j are the number of sampling points in the PTV, normal tissue and in the j^{th} OAR. The objectives f_L^{PTV} , f_H^{PTV} and f_{OAR}^j are proportional to the DVH values at the corresponding dose values. As a fraction they are normalized in the range [0, 1]. The dose values are expressed as fractions of the prescription dose. If the sources are limited inside the PTV then the range of the values of f_{NT} is of comparable magnitude to the other objectives and the objectives are therefore approximately normalized.

The use of the square of the dose value ensures that high dose values are more likely to be avoided in the NT than are the more uniformly distributed moderate dose values. For D_H^V we use a value of 1.5 times the prescription dose. The dose is not normalized on any particular point or set of points.

In Milickovic *et al* 2002 we have shown results using gradient-based optimization algorithms for objectives that correspond to the case $\alpha=2$. We called these objectives variance-based objectives. We have chosen as normalization the average dose on the PTV surface. This dose value D_{ref} is set to be equal to the prescribed dose. The isodose of an optimal dose distribution of the prescription dose coincides with the PTV surface. To obtain

such a dose distribution the dose variance f_S of the sampling points (dose points) uniformly distributed on the PTV surface should be as small as possible. If the source dwell positions are inside the PTV then the use of an additional objective for the surrounding NT is not necessary. The avoidance of excessive high dose values inside the PTV is controlled by the dose distribution variance f_V inside the PTV. Normalized variances are used:

$$f_S = \frac{1}{N_S} \sum_{i=1}^{N_S} \frac{(d_i^S - m_S)^2}{m_S^2}, \quad f_V = \frac{1}{N_V} \sum_{i=1}^{N_V} \frac{(d_i^V - m_V)^2}{m_V^2}$$

Where m_S and m_V are the average dose values on the PTV surface and within the PTV respectively, and N_S , N_V the corresponding numbers of sampling points. The objective space of (f_S, f_V) is convex and gradient-based algorithms converge to the global Pareto front. If OARs are to be considered then an additional objective is included for each OAR:

$$f_{OAR} = \frac{1}{N_{OAR}} \sum_{i=1}^{N_{OAR}} \frac{\Theta(d_i^{OAR} - D_c^{OAR} m_S)(d_i^{OAR} - D_c^{OAR} m_S)^2}{(D_c^{OAR} m_S)^2}$$

Where N_{OAR_i} is the number of sampling points in the OAR and D_c^{OAR} is the corresponding critical dose as a fraction of the prescription dose or reference dose which is the average dose on the PTV surface.

2.4 The non-dominated sorting genetic algorithm (NSGA-II)

Multiobjective genetic algorithms have been used in external beam radiotherapy (Haas *et al* 1998) and in brachytherapy (Lahanas *et al* 1999). Now more powerful algorithms are available. True multiobjective evolutionary algorithms, MOEAs do not require importance factors, as reported in Wu and Zhu 2001.

Deb *et al* 2001 proposed the multiobjective genetic algorithm *NSGA-II* that combines elitism and a mechanism to distribute the solutions as much as possible over the entire local Pareto set. Elitism is a method that guarantees that the best ever found solution would always survive the evolutionary process. For multiobjective optimization elitism requires that some portion of the non-dominated solutions will survive.

We have modified this algorithm by including an initialization by a gradient-based optimization algorithm and an archiving option. We describe some important aspects of this algorithm and the modifications that we applied.

2.4.1 Original algorithm

In contrast to single objective genetic algorithms where each individual is assigned a fitness value based on a single objective value, which determines the quality of this individual, there is no such single value for a true multiobjective optimization algorithm. It is therefore natural to use the dominance relation to classify an individual of the population. *NSGA-II* classifies a population P in mutually exclusive equivalent classes P_i i.e. $P_i \cap P_j = 0 \quad \forall i \neq j$ and

$P = \bigcup_{i=1}^k P_i$. These classes are called *fronts*. The number of classes varies from generation to generation and the members in each class are equivalent. That is, it cannot be stated which individual is better. This classification which is called *non-dominated sorting* is implemented as follows.

All non-dominated individuals are classified into one category and assigned a dummy fitness value or rank. These classified individuals are therefore ignored and from the remaining members of the population the non-dominated individuals are selected for forming the next layer. This process continues until all members are classified. Individuals of the first

layer have the highest fitness while members of the last layer are assigned the smallest fitness. All Individuals from the first layer produce more copies in the next generation and therefore the population evolves towards the Pareto front.

Finite populations tend to converge to a small part of the Pareto front. This phenomenon is called *genetic drift* and requires special mechanisms to prevent this occurring. For each member of the population an estimation of the density of solutions surrounding it is calculated using the so-called *crowding distance* measure: the population is sorted for each objective function value in their ascending order of magnitude. The solutions with the smallest and largest value are assigned a very large distance estimate to guarantee that they will be selected in the next generation. This helps to significantly increase the extent of the population on the Pareto front. All other solutions are assigned a distance value equal to the absolute difference in the function values of two adjacent solutions. The overall crowding distance is the sum of the individual distance values to each objective.

The elitism is used by combining the population of children Q_t and the parent population P_t at generation t together. A non-dominated sorting is applied and a new population is formed. A population of children Q_{t+1} from P_{t+1} is formed using a binary crowded tournament selection, crossover and mutation.

2.4.2 Modifications: Hybridization using supported solutions and archiving

Comparisons of *NSGA-II* with other genetic algorithms show that it generates efficiently solutions close to the global Pareto front only for problems with a small number of parameters. A statistical analysis shows that a very large number of generations are required to approach the global Pareto front. The population initially moves fast towards the global Pareto front but slows down finally and approaches the Pareto front only asymptotically.

If a small number of individuals are moved on the global Pareto front using an efficient optimization algorithm, then they efficiently attract the other individuals in their neighborhood. These so-called *supporting solutions* (Gandibleaux *et al* and Milickovic *et al* 2001) improve very significantly the performance of evolutionary elitist algorithms and enable a fast accumulation of solutions close to the global Pareto front. We combine the *NSGA-II* algorithm with a deterministic algorithm. The idea is to improve the diversity along the front and to guide the population closer and faster towards the global Pareto front. It is important to include if possible solutions that define the extent of the Pareto front, *i.e.* the optimal value of each individual objective.

Even by applying constraints for more than three objectives the size of the Pareto front can be large and a very large population would be necessary to cover as uniformly as possible the desired part of the trade-off front. We therefore employ to *NSGA-II* an archiving method in which an external population is filled with all the non-dominated solutions found. At each generation, each solution of the genetic population is compared with all the individuals of the external population. If a solution for the genetic pool dominates solutions for the external population and if it is non-dominated by them then it is added to the external pool and all the dominated solutions are removed. If the population has reached its maximum size, as set by the user, then the non-dominated solution replaces a solution that has the highest number of individuals in its neighborhood. This mechanism prevents accumulation of solutions in a part of the objective space and allows a more uniform coverage of the Pareto front by the archived solutions. The non-domination tests usually require much less time than the function evaluation and the non-dominated sorting. Therefore a very large number of non-dominated solutions can be acquired quickly. The accumulation of non-dominated solutions is shown as an example in Fig. 1 as a function of the generation of the genetic algorithm.

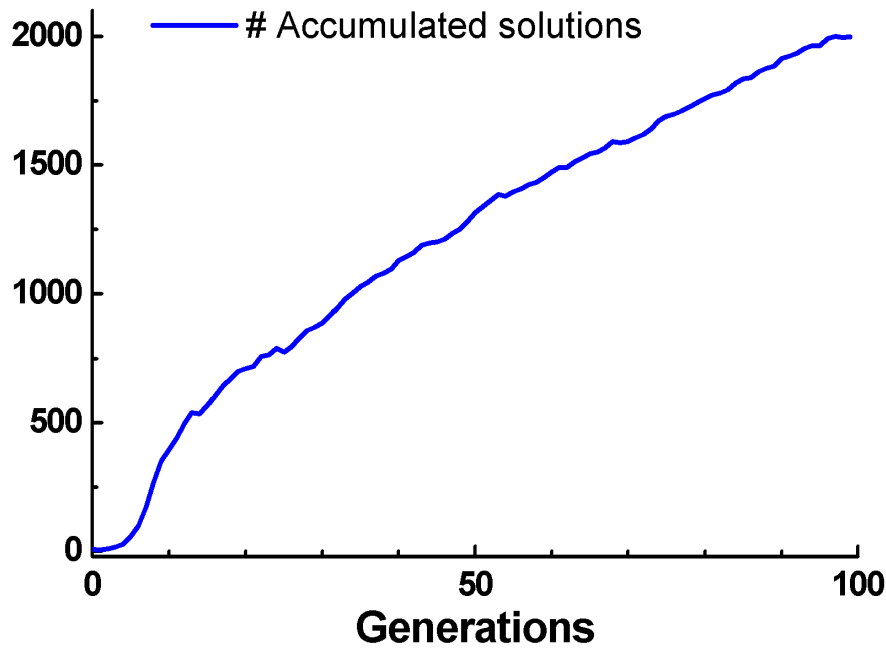


Figure 1. Example of the number of accumulated archived non-dominated solutions as a function of the generation number.

The average linkage method can be applied if necessary to reduce the number of solutions without modifying their distribution. In this case all solutions are considered initially as a single cluster. The distance of two clusters is given as the average distance between pairs of individuals across the two clusters. Two clusters with minimal distance are merged into a larger cluster. This procedure is repeated until the number of clusters reaches a user specified size. Finally, a point is selected from each cluster with minimal average distance to all other points inside the cluster.

We call this modified version of *NSGA-II*, as the modifications are essential for the performance of the algorithm for large-dimensional and multi-objective problems, the non-dominated sorting archived and supported genetic algorithm or *NSASGA*.

For the *NSASGA* algorithm a population of strings is formed, each string storing a set of dwell weights and a floating-point double precision representation is used. The crossover and mutation probability P_c and P_m must be limited in the respective ranges 0.7-1.0 and 0.001-0.1. We use a population size of 300. Tests obtained empirically have shown that the optimization results do not change significantly after 100 generations. This is because the population is initialized with a deterministic algorithm. Otherwise it would require 1000 or more generations to produce similar results.

We use an adaptive mutation and crossover operator (Deb and Agrawal 1995). Two parents with variables x_i^1 and x_i^2 produce two new variables y_i^1 and y_i^2 using a probability distribution $P(\beta)$:

$$P(\beta) = \left\{ \begin{array}{l} 0.5(\eta+1)\beta \\ 0.5(\eta+1)\frac{1}{\beta^{\eta+1}} \end{array} \right\}, \quad \beta = \left| \frac{y_i^1 - y_i^2}{x_i^1 - x_i^2} \right| \quad (4)$$

The parameter η determines the distance of the children from their parents. The probability distribution produces near-parent solutions with a larger probability than solutions far from the parents.

For DVH based objectives a large variety of solutions is possible. If the Pareto front is very large then for a small population, some important parts of the Pareto surface may not be occupied. It is possible to use constraints which allows to restrict the population into parts of the objective space which are of particular interest. Such a constraint when used for the DVH

based objectives is $f_L^{PTV} < 0.7$, i.e. only solutions with a PTV coverage of at least 30% are considered.

2.5. Decision-making tools.

For single objective optimization methods there is only one optimization result and the only decision necessary is whether or not to accept the solution. For multiobjective optimization decision-making tools are necessary to filter a single solution from a Pareto front that matches at best the goals of the treatment planner.

A utility function is a model of the decision maker's preference that maps a set of objective functions into a value of its utility. The goal is to maximize the utility. One such utility function is the Conformal Index *COIN* (Baltas *et al* 1998) that is now extended for the inclusion of OARs. It is additionally a measure of implant quality and dose specification in brachytherapy. *COIN* takes into account patient anatomy, both of the tumor, NT and OARs. *COIN* for a specific dose value D is defined as:

$$COIN = c_1 \cdot c_2 \prod_{i=1}^{N_{OAR}} \left(1 - \frac{V_{OAR}^i(D > D_{crit}^i)}{V_{OAR}^i} \right) \quad (5)$$

$$c_1 = \frac{PTV_D}{PTV} \quad (6)$$

$$c_2 = \frac{PTV_D}{V_D} \quad (7)$$

The coefficient c_1 is the fraction of the PTV (PTV_D) with dose values of at least D . The coefficient c_2 is the fraction of the calculated (body) volume with dose values of at least D (V_D) that is covered by the PTV. It is also a measure of how much NT outside the PTV is covered by D . V_{CS}^i is the volume of the i^{th} OAR and $V_{CS}^i(D > D_{crit}^i)$ is the volume of the OAR that receives a dose that exceeds the critical dose D_{crit}^i . The product in Eq. 5 covers all N_{OAR} OARs.

In the case where an OAR receives a dose D above the critical value defined for that structure, the conformity index will be reduced by a fraction that is proportional to the volume that exceeds this limit. We describe the dependence of *COIN* on the choice of the reference dose value as the *COIN distribution*. If D is chosen to be the reference dose D_{ref} then the ideal situation is $COIN = c_1 = c_2 = 1$. *COIN* assumes in this form that the PTV, the OARs and the surrounding normal tissue are of the same importance.

In the case that a treatment planner considers one objective to be more important than the others, a display table of a list of values for all solutions of the objectives, *COIN*, DVHs for all OARs, the NT and the PTV of each solution is provided. Additionally, the extreme dose values are also provided. The entire table for every such quantity can then be sorted and solutions can be selected and highlighted by the treatment planner. Constraints can then be applied such as to show only solutions with a PTV coverage $100c_1$ larger than a specified value. The PTV coverage is the percentage of the PTV that receives at least 100% of the prescription dose. This reduces the number of solutions and in this way the treatment planner understands the available possibilities. The DVHs of all selected solutions can be displayed and compared.

Another decision-making tool is the display of projections of the Pareto front onto a pair of selected objectives. For M objectives the number of such projections is $\binom{M}{2}$. The position of

selected solutions can be seen in these projections this helps to identify their position in the multidimensional Pareto front. This permits quantitative understanding of the degree of correlation between the objectives and of the possibilities provided by the non-dominated set. The Pareto front provides information such as: How much can an objective be optimized and how this modifies the other objectives values? What is the range of values for each objective?

3. RESULTS

We present results of *NSASGA* for a 83 cm³ prostate implant with 15 catheters and 125 active source dwell positions. We considered as OARs the urethra and the rectum for which the corresponding critical dose values were set to be 1.25 and 0.75 times the prescription dose of 9.0 Gy. This is a special treatment case with a high dose per fraction so that the protection of urethra from an overdose due to a high prescription dose is important. The objective values have been obtained from 500 sampling points quasi-randomly distributed inside the PTV and OARs. For the surrounding NT 2000 sampling points were used. Sampling points inside the catheters were excluded in order to avoid fluctuations of the dose variances. For the sampling points on the PTV surface we used uniformly distributed sampling points uniformly distributed on the triangulated PTV surface with a surface density of 3 points/cm² (Lahanas *et al* 2000).

We used two sets of sampling points, a set of 2000-3000 points for the calculation of the objectives values during the optimization. The DVHs presented here and derived parameters such as COIN for all solutions to be used for the decision making process were obtained from a set of 20000 sampling points in each object. The DVHs obtained from the PLATO BPS system were obtained from 30000 sampling points.

We used dose calculation point look-up tables (LUT) in which the kernel values for each dose calculation point and source dwell position pair were calculated and stored in a table in a preprocessing step. If we ignore the preprocessing time then the calculation time for the dose distributions is independent of the form of the dosimetric kernel. We used dosimetric kernels obtained by Monte Carlo simulation (Angelopoulos *et al* 1991, Sakelliou *et al* 1992, Karaikos *et al* 1998 and 1999).

The dose distribution around a cylindrical source is not isotropic, due to the attenuation of the photons in the active source material, the encapsulation material, the source drive cable, etc. Due to the cylindrical rotational symmetry the dose rate in a uniform isotropic medium is a function only of r and θ . The orientation of each source is determined from the catheter geometry and at each source dwell position a vector is calculated which is parallel to the cylindrical source axis and in opposite direction to the source drive cable.

The calculations were performed using a 933 MHz Intel III Windows NT computer with 512 MB RAM. The optimization time depended mainly on the number of possible dwell positions. The calculation of the statistics, and of DVHs for up to 1000 solutions requires 2 min. A speedup by a factor of two is possible with available standard 2 GHz PCs.

For the optimization with variance based objectives we used the limited memory Broyden-Fletcher-Goldberg-Shanno (*L-BFGS*) algorithm (Liu and Nocedal 1989). We used as optimization parameters the square root of the dwell times. This avoids completely solutions with negative dwell times and constraints for positive dwell times are not necessary. Numerical experiments based on 100000 optimization runs, using up to hundred different sets of importance factors provide evidence for global convergence, *i.e.* the results does not depend on the starting point. Additional comparisons with simulated annealing suggests that the solutions are global optimal.

We study the influence of the initialization of *NSASGA* with *L-BFGS*. We compare the Pareto fronts obtained by *NSASGA* for variance and DVH based objectives with the Pareto sets obtained by *L-BFGS* or simulated annealing.

We compare the dose optimization results obtained with variance based and DVH based objectives. Additionally we include a comparison with the optimization result obtained using Nucletron Plato BPS 13.7 system.

3.1 Variance-based objectives

We compare the Pareto fronts from *L-BFGS* and *NSASGA* for variance based objectives with and without initialization. For the comparison we did not use any constraint for f_L^{PTV} . For the initialization with *L-BFGS* we use a sampling parameter $k=4$, i.e. we initialize 35 members of the population with solutions from *L-BFGS* using variance based objectives. The two-dimensional projections of the non-dominated sets are shown in Fig. 2. We include the Pareto set obtained with $k=19$ (816 solutions) for *L-BFGS*. The optimization time for *L-BFGS* was 234 s. The optimization time for *NSASGA* with and without initialization was 134 s and 136 s respectively. The corresponding numbers of obtained non-dominated solutions were 1996 and 1511.




Figure 2. Two-dimensional projections of non-dominated solutions obtained by *NSASGA* with and without initialization. Included are the solutions found by *L-BFGS*.

3.2 DVH based objectives

We compare the Pareto front for DVH based objectives obtained with *NSASGA* with and without initialization by *L-BFGS*. We compare the *NSASGA* Pareto front additional with the Pareto front obtained with fast simulated annealing *FSA* (Szu and Hartley 1987). For *NSASGA* we use the same initialization as for the previous test with variance-based objectives ($k=4$). The optimization time with and without initialization was 219 s and 231 s respectively. The number of accumulated solution was 1332 and 2000 respectively.

For *FSA* 330 solutions ($k=7$) have been generated with 100000 iterations per solution. Additional the algorithm was initialized with solutions obtained by *L-BFGS* using DVH based objectives which were approximated by:

$$f_L^{PTV} = \frac{1}{N_{PTV}} \prod_{i=1}^{N_{PTV}} \Theta(D_L^V - d_i)(D_L^V - d_i)^\alpha$$

$$f_H^{PTV} = \frac{1}{N_{PTV}} \prod_{i=1}^{N_{PTV}} \Theta(d_i - D_H^V)(d_i - D_H^V)^\alpha$$

$$f_{OAR}^j = \frac{1}{N_{OAR}} \prod_{i=1}^{N_{OAR}} \Theta(d_i - D_{crit}^j)(d_i - D_{crit}^j)^\alpha$$

with $\alpha=0.05$. The optimization time for *FSA* was 26 hours. In Fig. 3 we show three two-dimensional projections of the non-dominated solutions from *NSASGA*, with and without initialization, and the result from *FSA*. The non-uniform distribution of solutions obtained by *FSA* is an indication that local minima exist and degenerate cases, *i.e.* for the same set of importance factors different objective values can be obtained although their weighted sum is constant. For the variance-based objectives we do not observe degenerate cases and the Pareto fronts obtained by *L-BFGS* and *FSA* are identical. If we consider the results of *FSA* as global optimal then *NSASGA* provides a fairly good representation of the global optimal Pareto front.




Figure 3. Three two dimensional projections of non-dominated solutions obtained by NSASGA with and without initialization. Included are solutions found by FSA.

3.3 Comparison of solutions obtained with variance and DVH based objectives

3.3.1 Variance-based objectives

The maximum COIN value obtained from the 816 *L-BFGS* solutions with variance-based objectives was 0.51 and the PTV coverage was 86%. In this case 15% of the urethra and 1% of the rectum receive a dose value above their corresponding D_{crit} value. If we set a constraint of at least 90% coverage for the PTV then the maximum COIN is 0.43 and the PTV coverage for this solution is 92.2%. We then have a 24% overdose in the urethra and a 4% overdose in the rectum.

A treatment planner considered this as the best solution. The reason is that there is a strong trade-off between the coverage of the PTV and the overdose to the urethra volume. From a PTV coverage of 92% - 94% the overdose to the urethra volume increases rapidly from 24% to 66%. The histograms for the DVHs of the PTV and the urethra for all solutions and for the selected solution are shown in Fig. 4. We include the best solution found by a treatment planner with the planning system PLATO BPS 13.7.

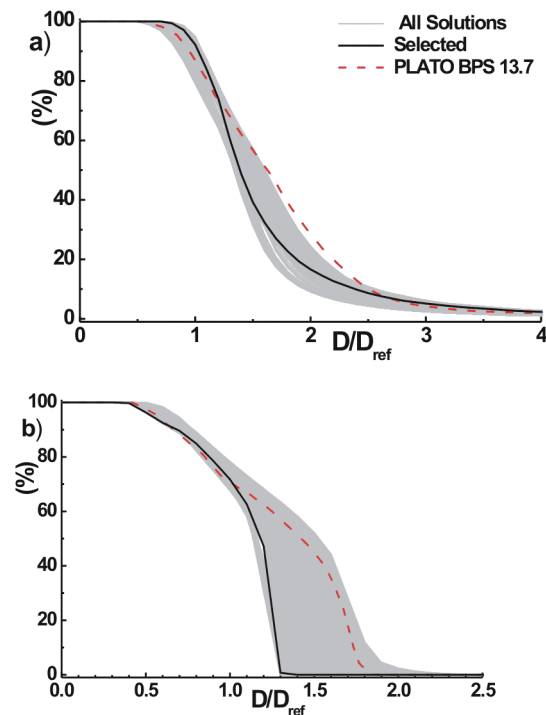


Figure 4. DVHs for a) the PTV, for b) the urethra of all non-dominated solutions obtained by *L-BFGS* for the prostate implant. The solution selected by a treatment planner is shown. Included is the best solution found using PLATO BPS 13.7.

3.3.3 DVH based objectives

We applied a DVH based optimization with *NSASGA*. 1679 solutions were accumulated in 227 sec. The maximum COIN we obtained was 0.766 and the PTV coverage for this solution was 86.5%. The overdose volume to the urethra and rectum was 3%. If we set a minimum of 90% coverage then we obtain the best solution for the urethra with a 21% overdose and with a 6% overdose for the rectum. The PTV coverage for this solution is 92.4%.

The optimization results are very similar to the results obtained with *L-BFGS* and show that the critical dose in the urethra sets a rigid constraint. In Fig. 5 we see the tradeoff between the three objectives a) PTV coverage b) the percentage of volume of urethra with overdose and c) the percentage of volume of rectum with overdose. Two-dimensional projections of Fig. 5 are shown in Fig 6. For a PTV coverage larger than 85% there is a rapid increase of the overdose volume in the urethra up to 70%, see Fig. 6a. If the critical dose for the urethra is 1.25 times the prescription dose then only with a PTV coverage of around 80% can we completely protect the urethra. In Figure 6b we see that the rectum overdose volume increases with the PTV coverage. At 100% PTV coverage solutions with a rectum overdose between 8% and 12% can be found.

The DVH of the PTV, the urethra of all solutions of *NSASGA* is shown in Fig. 7. The solution selected by the treatment planner using the decision-making tools is marked in this figure.




Figure 5. Trade-off between three objectives: PTV coverage, percentage of volume of urethra overdose and rectum overdose. The three two dimensional projections are shown. These show the trade-off between two objectives.

Comparing the *NSASGA* and *L-BFGS* results with those obtained using the Nucletron PLATO BPS 13.7 system we found the following. For the best result that we obtained with PLATO for the urethra, the overdose was 60%, the PTV coverage 87% and the COIN was 0.25. Although an experienced treatment planner spent approximately 1 hr it was not possible to find a solution with a better protection for the urethra. The DVHs of the PTV, the urethra of the selected solution generated by *L-BFGS*, *NSASGA* in comparison to the solution that was obtained by PLATO are shown in Fig. 8.




Figure 6. Tradeoff between a) PTV coverage and volume of urethra with overdose, b) PTV coverage and volume of rectum with overdose (see also Fig. 5).




Figure 7. DVHs for a) the PTV and for b) the urethra of all non-dominated solutions obtained by NSASGA for the prostate implant. The solution selected by a treatment planner is shown. Included is the best solution found using PLATO BPS 13.7.

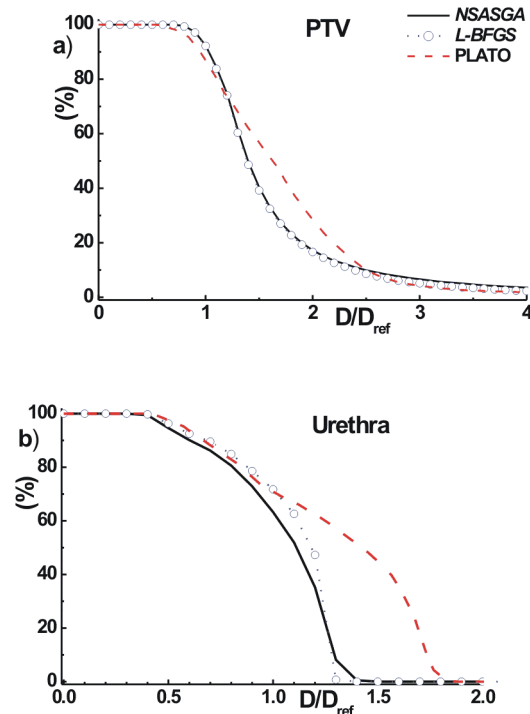


Figure 8. DVHs for a) the PTV, for b) the urethra for a solution selected from the non-dominated set produced by L-BFGS and NSASGA for the prostate implant. These solutions are compared with a solution obtained with Nucletron PLATO BPS 13.7 after 1 hour by an experienced treatment planner.

4. DISCUSSIONS AND CONCLUSIONS

In the past, emphasis has been given to optimization methods that provide a single solution. Such a solution, obtained by a weighted sum approach, is some unknown point, probably on the convex part of the global Pareto front. If for a given set of importance factors the result is not satisfactory, and if it is possible, then importance factors would be changed and by trial and error, the results would be compared with the previous solutions. As the number of objectives increases it is more difficult to guide the optimization engine to the desired result or possible result, which are not always the same result. Some of the treatment planning systems do not even provide this possibility. Therefore a treatment planner sometimes has to rescale or modify the dwell times or even manually select other source dwell positions based on information of dose distributions. This is in order to increase the PTV coverage and/or to protect OARs from overdose.

A method has been proposed by Xing *et al* 1999 in intensity modulated beam radiotherapy for the selection of a set of optimal importance factors. In this approach the problems of the importance factors have only been replaced by another set of importance factors used for the DVH based objectives that steer the optimization engine. This requires some knowledge that is not always available *a priori*. In some cases the range of the importance factors for each objective may differ significantly. In Souza *et al* 2001 a weighted sum approach with a mixed integer programming algorithm was presented and it was reported that the weights for the five objectives used should be in the range 300-500, 300-2000, 1-300, 5-100 and 10^4 - 10^7 . Even if finally, a single set of weights was used for all optimizations, the determination of an unique optimal set is difficult and does not always gives the best result or even a satisfactory result. Additionally, due to physical constraints, some ideal DVHs must be used even if it is unknown as to what can actually be achieved. Finally, if a solution is presented then the treatment planner has no insight into what alternatives solutions could have been selected.

In brachytherapy it is difficult to quantify what is meant by a optimal dose distribution, in the presence of OARs and treatment planners have their own insight as to how a dose distribution will be "optimal" for a particular type of cancer in a particular location. This means that an optimization engine should provide to the treatment planner with all possibilities that can be achieved.

Deterministic gradient-based algorithms are very powerful and a few 100 solutions can be obtained in 1-2 minutes. Using the hybrid genetic algorithm NSASGA we obtained representations of the trade-off surfaces with 1000-2000 solutions in 2-5 minutes. For DVH based objectives this large number is necessary if more than one OAR has to be considered. The objectives are directly related to the DVHs which finally are used in the decision making process. There is no need for any dose normalization. The surrounding normal tissue is considered and there is no artificial use of variances that are employed as an approximation to the optimal dose distribution. The DVH based objectives are less restrictive than the variance based objectives. The results when compared with a weighted sum based global optimization FSA algorithm are similar but FSA requires for the same number of solutions many hours or even days rather than 2-5 minutes!

The dose normalization applied with the variance-based objectives protects the surrounding normal tissue as long as there are no source dwell positions outside the PTV. If the source dwell positions inside the PTV are distributed such that large parts of the PTV are not left empty, then if we can ignore the OARs we then obtain the best COIN values. The solutions obtained by *L-BFGS* are global Pareto optimal. The normalization limits the maximum coverage of the PTV with the prescription dose in the case when the reference isodose surface due to a poor distribution of source dwell positions is such that it cannot have the shape of the PTV surface. The dose distributions with the DVH based objectives allow a greater coverage since it does not require any dose normalization.

The complexity of the Pareto front increases rapidly with the number of objectives. This is a problem not only for multiobjective optimization methods but is a general problem if one has to consider many OARs. We have transformed the dose optimization problem into a decision-making problem. Although this approach assumes an ability of the treatment planner to understand the possibilities offered from the large number of solutions we believe that this is not a serious problem and that a treatment planner can filter out using the simple decision making tools an appropriate solution in 2-3 minutes. The planner obtains first information of the possibilities from the range of DVH values for all structures of interest and from other quantities such as maximum dose, COIN etc. If none of the objectives is preferred then COIN can be used for the selection of a good solution. However this is not always true as for example shown for the prostate implant case. This demonstrated that a simple utility function cannot be used and that additional information such as provided by the non-dominated solutions is valuable and necessary for the decision.

NSASGA provides in approximately five minutes or less, a large number of non-dominated solutions. An additional 2-3 minutes are required to select a solution that satisfies in the best possible way the planners' objectives. Comparison with other implants made but not presented here show that the result is almost better and in the worst case as good as the solution provided by PLATO BPS, which in difficult cases requires a manual intervention by the treatment planner which in some cases may take 1-3 hours and sometimes even without a satisfactory result.

NSASGA does not require any importance factors, is not restricted to a special class of implant sites or type and do not need any special adjustment from case to case. *NSASGA* using the DVH based objectives provide solutions that cannot be obtained by the deterministic variance based objective algorithms such as *L-BFGS* due to the restrictions of the dose distribution or due to the problems of obtaining solutions in non-convex spaces.

As the "evolution" of evolutionary algorithms is continuing, their performance due to for example new and better niching mechanisms, or due to inclusion of local search operators will allow us to obtain faster and better results. We believe that this finally will make this class of algorithms the choice for multiobjective dose optimization in brachytherapy. First results of the application of *NSASGA* for the inverse planning problem in brachytherapy, where not only the source dwell times but also the optimal position and number of catheters has to be found, are encouraging.

ACKNOWLEDGEMENTS

This investigation was supported by a European Commission Grant (IST-1999-10618, Project: MITTUG).

REFERENCES

- Angelopoulos S, Perris S, Sakellariou K, Sakelliou L, Sarigiannis K and Zarris G 1991 Accurate Monte Carlo calculations of the combined attenuation and build-up factors, for energies (20-1500 keV) and distances (0-10 cm) relevant in brachytherapy *Phys. Med. Biol.* **36** 763-78
- Baltas D, Kolotas C, Geramani K, Mould R F, Ioannidis G, Keckidi M and Zamboglou N 1998 A Conformal Index (COIN) to evaluate implant quality and dose specifications in brachytherapy *Int. J. Radiation Oncology Biol. Phys* **40** 512-24
- Das I and Dennis J 1997 A Closer Look at Drawbacks of Minimizing Weighted Sums of Objectives for Pareto Set Generation in Multicriteria Optimization Problem *Structural Optimization* **14** 63-9
- Deb K and Agrawal R B 1995 Simulated binary crossover for continuous search space *Complex Systems* **9** 115-48
- Deb K, Agrawal S, Pratap A and Meyarivan T 2000 A fast and elitist multiobjective genetic algorithm: NSGA-II. Technical Report 20001, Indian Institute of Technology, Kanpur, Kanpur Genetic Algorithms Laboratory (KanGAL).
- Gandibleaux X, Morita H and Kato N 2001 The Supported Solutions Used as a Genetic Information in Population Heuristic, in Proceedings of the first international conference, EMO 2001, Zurich, Switzerland, edited by E. Zitzler, K. Deb, L. Thiele, C. A. Coello Coello, D. Corne, Lecture Notes in Computer Science Vol. 1993, Springer 429-42
- Haas O C L, Burnham K J and Mills J A 1998 Optimization of beam orientation in radiotherapy using planar geometry *Phys. Med. Biol.* **43** 2179-93
- D'Souza W D, Meyer R R, Thomadsen B R and Ferris M C 2001 An iterative sequential mixed-integer approach to automated prostate brachytherapy treatment plan optimization *Phys. Med. Biol.* **46** 297-322
- Karaiskos P, Angelopoulos A, Sakelliou L, Sandilos P, Antypas C, Vlachos L and Koutsouveli E 1998 Monte Carlo and TLD dosimetry of an ^{192}Ir high dose rate brachytherapy source *Med. Phys.* **25** 1975-84
- Karaiskos P, Angelopoulos A, Baras P, Sakelliou L, Sandilos P, Dardoufas K and Vlachos L 1999 A Monte Carlo investigation of the dosimetric characteristics of the VariSource ^{192}Ir high dose rate brachytherapy source *Med. Phys.* **26** 1498-502
- Lahanas M, Baltas D and Zamboglou N 1999 Anatomy-based three-dimensional dose optimization in brachytherapy using multiobjective genetic algorithms *Med. Phys.* **26** 1904-18
- Lahanas M, Baltas D, Giannouli S, Milickovic N, and Zamboglou N 2000 Generation of uniformly distributed dose points for anatomy-based three-dimensional dose optimization methods in brachytherapy *Med. Phys.* **27** 1034-46
- Lessard E and Pouliot S 2001 Inverse planning and anatomy-based dose optimization in HDR-brachytherapy of the prostate using fast simulated annealing and dedicated objective functions *Med. Phys.* **28** 773-9
- Liu D C and Nocedal J 1989 On the limited memory BFGS method for large scale optimization *Mathematical Programming* **45** 503-28
- Miettinen K M 1999 *Nonlinear Multiobjective Optimization*, (Boston, MA: Kluwer).

Milickovic N, Lahanas M, Baltas D and Zamboglou N 2001 Comparison of Evolutionary and Deterministic Multiobjective Algorithms for Dose Optimization in Brachytherapy, in Proceedings of the first international conference, EMO 2001, Zurich, Switzerland, edited by E. Zitzler, K. Deb, L. Thiele, C. A. Coello Coello, D. Corne, Lecture Notes in Computer Science Vol. 1993, Springer 167-80

Milickovic N, Lahanas M, Papagiannopoulou M, Zamboglou N and Baltas D 2002 Multiobjective anatomy-based dose optimization for HDR-brachytherapy with constraint free deterministic algorithms *Phys. Med. Biol.* **47** 2263-80

Sakelliou L, Sakellariou K, Sarigiannis K, Angelopoulos A, Perris A and Zarris G 1992 Dose rate distributions around ^{60}Co , ^{137}Cs , ^{198}Au , ^{192}Ir , ^{241}Am , ^{125}I (models 6702 and 6711) brachytherapy sources and the nuclide $^{99}\text{Tc}^m$ *Phys. Med. Biol.* **37** 1859-72

Szu H and Hartley R 1987 Fast Simulated Annealing *Phys. Lett. A* **122** 157-62

Wu X and Zhu Y 2001 An optimization method for importance factors and beam weights based on genetic algorithms for radiotherapy treatment planning *Phys. Med. Biol.* **46** 1085-99

Xing L, Li J G, Donaldson S, Le Q T and Boyer A L 1999 Optimization of importance factors in inverse planning *Phys Med. Biol.* **44** 2525-36