

Optimisation of Irradiation Directions in IMRT Treatment

Lenrick A. Johnston
Department of Operational Research
University of Auckland
New Zealand
rick@semrits.co.nz

Abstract

Inverse planning has the ability to dramatically improve treatment outcomes for cancer patients undergoing radiation treatment. One of the major challenges in inverse planning is the development of effective tools to support the planning process. Not only must a large amount of patient data be collected and analysed, but a significant number of output variables can be changed and has to be determined to optimise treatment. These variables include irradiation directions, intensity profiles and beam parameters.

We present a prototype decision tool for planners which optimises irradiation direction. This decision tool also incorporates intensity profile optimisation, to create an integrated approach to planning. Optimisation of irradiation directions is a computationally difficult problem. A number of strategies are explored which will produce good treatment plans within a realistic time frame. These strategies incorporate heuristic optimisation techniques, aiming to minimise both tumour underdosing and overdosing of healthy organs simultaneously. Initial results indicate that significant improvements can be made in primary tumour treatment over current approaches.

1 Background

Cancer is currently one of the leading causes of death in 1st world industrialised nations. Three main methods of treatment exist: surgery, chemotherapy, and radiotherapy. While surgery and chemotherapy are relatively specialised forms of treatment, neither are as dependent as radiotherapy on computer-based technologies. Radiotherapy focuses a number of radiation beams from external sources onto the cancerous tissue. The aim is to cause minimal damage to healthy organs, while guaranteeing a lethal dose to the tumour.

Radiotherapy is a particularly successful treatment for localised tumours, contained within a certain region of the body. The procedure is non-invasive, has relatively few side-effects compared to either chemotherapy or surgery, and is quick and painless to perform.

However radiotherapy does present technical difficulties for those planning treatment. A large number of factors – including the intensity and frequency of the

radiation beam, the placement of the beam, spreading and diffraction effects, critical organ location – must be considered in the planning of radiotherapy treatment.

1.1 Intensity Modulated Radiotherapy (IMRT) Planning

Traditional radiotherapy planning uses a forward approach, i.e. directions and intensities are chosen by the planner, and the resulting dose distribution calculated. If it is unsatisfactory the process is repeated until a "good" dose distribution is obtained. With modern radiotherapy equipment, especially IMRT (intensity modulated radiation therapy) these techniques can hardly take into account the complex interactions between all planning parameters.

Inverse Planning is a 'backward' strategy to treatment planning, which uses a target dose as a starting point and attempts to find intensities that will produce that dose. Inverse Planning frames the problem clearly in terms of optimisation: it specifies the dosage which the tumour should receive, the dosage which healthy organs should not exceed, and finds intensities that match these doses optimally. This approach has the advantage of avoiding trial and error; it specifies the desired physical output as a starting point. The fact that Inverse Planning makes no prior assumptions suggests that it has the potential to produce significantly improved treatment plans.

1.2 Project Objective

This project investigates the optimisation of gantry angles in the Intensity Modulated Radiotherapy Planning (IMRT) problem. In particular, we will examine possible methods to find very good treatment plans by varying the angle and number of the radiation beams in relation to the patient. This research builds on previous work done by Burjony (2001) and Hamacher and Küfer (2002). The main difficulty lies in the large number of factors that must be considered: the position of tumour and healthy organs, optimal angle, intensity profile, shape and position of the radiation beam, and the number of beams to be used.

The aim of this project will be to produce treatment plans which are as close as possible to the desired 'ideal' dose bounds set by planners. These bounds specify both an ideal lower bound for the tumour (to guarantee the destruction of the cancerous tissue), and an ideal upper bound on healthy organs. While satisfying these bounds simultaneously is generally not possible, the goal of this project will be to find the optimal dose of radiation and gantry angles which will compromise most effectively between these two opposing needs.

2 A Complete Model

Previous models of the problem have not captured gantry angle as a variable in the optimisation problem. The existing formulation is as a Linear Program, which assumes that gantry angles are fixed by planners before optimisation is performed. Given a number of critical organs (those considered important by planners) and selected input angles, the existing model optimises the dosage given to each organ and outputs the intensity profile required to achieve this dosage.

However, since the *directions* of irradiation are changed by planners to manually produce better solutions, it seems reasonable to incorporate this variable into the optimisation problem. Indeed, changes in gantry position can have a significant effect both on tumour coverage and healthy organ damage. As Figure 1 shows, the gantry head is able to turn around the patient to produce radiation from almost any position.



Figure 1. Gantry as it turns around the patient.

2.1 Dose Distribution and Discretisation

One of the challenges in current research is to incorporate dose distribution information as the gantry angle changes. It is currently established practise to assume that the relationship between emitted radiation and received dosage is linear: if twice as much radiation is emitted, twice as much radiation will be absorbed (Hamacher & Küfer, 2002).

However the dose distribution does not change linearly as the angle is changed, being related not only to the radiation intensity but also to the angle. The full model is therefore a high-dimensional multi-criteria non-linear optimisation problem which is difficult to solve with current optimisation technologies. Non-linear methods such as those presented by Tervo & Kolmonen (1999) presently use a combination of greedy heuristics, careful initialisation and ‘geometrical intuition’ to obtain reasonable results within an acceptable timeframe. One of the key difficulties is that the non-linear problem is multi-extremal, with many local optima. These problems cannot be solved by traditional methods of convex programming; at best, successive local optima can be explored for a global optimum.

This project considers an alternative approach which aims to overcome the difficulties associated with non-linear optimisation. In principle it is possible to consider each beam angle combination separately, solving each linear program (LP) in turn to find a solution. Angle discretisation means that the dose distribution is no longer related to both angle and intensity simultaneously. This discretisation decouples the problem allowing a single variable to be optimized. This approach creates a large number of LPs, each representing a particular set of gantry positions.

In reality, because step and shoot radiation therapy machines only have a limited number of possible operating positions, discretisation of gantry angle is an appropriate method of modelling the problem. An optimal beam placement of 31.4159 degrees will not yield a different solution in reality from 31 or even 32 degrees if the machine does not have this tolerance. Auckland Hospital currently uses tolerances of 1° in their radiotherapy planning.

In addition, a number of positions cannot be replicated by the machine due to physical impedances. This is easily modelled in a discretised formulation of the problem: positions which are physically impossible can easily be discarded.

2.2 Constraints

The goal of treatment planning is to successfully treat the target tumour to the bounds specified by medical specialists, while simultaneously preventing radiation damage to healthy organs. In order to achieve this, upper bounds are specified for healthy organs, and lower bounds are specified for the cancerous volume.

In most problems these two types of constraints are not achievable simultaneously. Either the radiation in the healthy organs must be allowed to go over set limits, or the tumour must be given a reduced dosage. In order to capture this possibility, a vector T is introduced which measures the deviation from ideal dosage. This relative deviation vector measures how much the constraint is violated.

The goal of radiotherapy planning then becomes to minimise these violations T , rather than focussing on the radiation beam intensities themselves. This model incorporates the varying ability for healthy organs to tolerate radiation, meaning that a better idea of the true impact of a radiotherapy plan can be obtained.

2.3 Model Formulation

As outlined above, the problem of optimising beam intensities for all possible angles is split into a number of sub-problems. Each of these problems represent a specific possible set of beam angles which can be expressed in terms of a single LP.

Indices

k = body part: $1 \dots K$

m = volume element: $1 \dots M$

i = voxel: $1 \dots M_k$

j = bixel: $1 \dots N$

h = gantry position: $1 \dots H$

λ = set of angular positions which the collimator supports: $1^\circ, 2^\circ, \dots, 360^\circ$

$S(\lambda, H)$ = the set of independent, non-recurring positions that the gantry ($1, \dots, H$) can be in. This is related to the size of the set λ : 1° resolution there

are $|S(1, H)| = \binom{360}{H}$ different combinations of angles possible.

Parameters

K = number of body parts involved in the radiation treatment plan.

$k=1$ for the tumour and $k=2 \dots K$ for healthy organs

H = number of gantry angles used in treatment of a single patient. H is set initially by planners.

B = number of beam elements, or *bixels*, per gantry position ($1 \dots H$). Each beam is discretised into a number of point sources. B is constant across all gantry positions; each beam has the same number of beam elements. This model is appropriate for multi-leaf collimators, which are able to change the radiation intensity across the beam width.

$N = H \times B$: the total number of beam elements (bixels) for all gantry positions.

Each gantry position ($1 \dots H$) has B bixels, giving a total of $H \times B$ bixels.

$M = M_1 + M_2 + \dots + M_K$: number of volume elements – called *voxels*. The body is discretised into M volume elements of equal size for the purposes of optimisation. M_k is the number of voxels in organ k .

D = dose vector, the amount of radiation absorbed by each voxel in the body is D_m .

P_{ij} = represents the contribution of bixel j to the total dose of radiation absorbed by voxel i under unit intensity.

L_1 = the minimum amount of radiation which each voxel of the cancer must absorb, in order to create a dose lethal to the tumour.

U = a vector of size $(K-1)$ representing the maximum dose U_k of radiation that each voxel of healthy organ k can safely absorb.

μ_k = weights of importance for each body part. (The sum of these weights is 1.)

$\Theta_s = (\theta(s)_1, \dots, \theta(s)_H)$ the angles of the gantry for a particular combination S

Decision variables

T_k = relative deviations in doses from those specified (U, L)

x_{hn} = radiation intensity in each bixel of the beam. For each head h , there are N bixels which can each emit differing amounts of radiation.

Model

- 1) $\min F(x, T)_{\Theta_s} = \sum_{k=1}^K \mu_k T_k$
- 2) $D_1 = P_1(\Theta_s) x \geq (L_1 - T_1) e$
- 3) $D_k = P_k(\Theta_s) x \leq (U_k + T_k)$ for $k=2, \dots, K$
- 4) All $x, T \geq 0$

Explanation

- 1) Minimise the weighted sum of the dose deviations in each of the organs
- 2) The dose in the tumour must be larger than $L_1 - T_1$
- 3) The dose in healthy organ k must be less than $U_k + T_k$

In this model, each LP will have a dose distribution expressed by $D = P(\Theta_s)x$, where matrix $P(\Theta_s)$ is given by entries that are $p(i,j)(\Theta_s)$ – the dose in voxel j resulting from radiation of unit intensity in bixel i given a set of angles Θ_s . This means that the dose arriving in voxel j from bixel i with unit intensity will depend on the angles chosen.

This decoupling of the problem allows an enumerative approach to optimisation; a number of LPs with different beam angles are solved. Optimality is ensured (within the limitations of the discretisation) since the entire solution set has been explored.

2.4 Comments

One possible method to solve this problem is to fully explore the solution space for every possible combination of gantry positions S . While a fully enumerative formulation will always guarantee optimality, it is an undesirable method of exploring the solution space. Problems with a small number of gantry positions can be solved reasonably quickly, however more than two points of insertion creates a huge number of LPs to be solved.

For example, the number of LPs to be solved is 360 using 1 gantry position, 64620 with 2 positions and over 7.7 million with three positions. Given that Auckland Hospital uses between 4 and 11 head placements in a typical treatment plan, a fully iterative technique is simply not possible. In the following section we discuss more intelligent ways of producing solutions to the Inverse Planning problem.

3 Summary of Optimisation Techniques

This section provides a summary of the techniques used to optimise the problem expressed in section 2.3. Rather than a complete mathematical description, a summary of method has been presented to give an indication of the variation in approaches trialled.

3.1 Approaches to optimisation

There are two possible methods of approaching the problem of gantry position: as part of the original problem definition, or as a separate optimisation process distinct from the selection of intensity profiles.

Perhaps the most intuitive approach is to integrate gantry angle into the problem definition, as it captures all aspects of the Inverse Planning problem in one mathematical model. The advantage of integrated models is that they guarantee that the optimal solution of the Inverse Planning problem is within the problem definition. Two models were investigated:

1. **A Mixed integer formulation**, with the use of Boolean variables y_h to turn ‘on’ and ‘off’ the various gantry angles. A constraint $\sum_{h=1}^{360} y_h \leq H$ is used to control the number of directions and additional constraints $x_{hn} \leq y_h$ guarantee that radiation is only emitted from selected directions.
2. **A local search using steepest descent**. Local search heuristics are used in a wide variety of applications and are very successful in quickly reaching good solutions to large optimisation problems.

However, two-phase methods are also an attractive way of modeling the problem, since they consider angle optimisation separately from intensity optimisation – thus avoiding difficulties with large scale MIPs. In a two-phase approach, a simplified version of the complete model is solved to select good beam directions. Angles are then passed to the original LP formulation and solved. Two approaches are examined:

1. **Set covering**. This approach aims to ‘cover’ the tumour while minimising damage to healthy organs.
2. **Selection using LP relaxation**. This technique uses linear programming techniques to produce a number of angles which provide a lower bound on the optimal solution. These angles are then investigated in turn using heuristic techniques.

3.2 Mixed Integer Techniques

On initial analysis, the most favourable approach to the discretised IMRT problem is to express gantry positions as part of the optimisation process, and simply solve the problem. This approach has the advantage of being relatively simple to formulate, incorporating all aspects of the model previously developed.

The methodology was to use Boolean variables to limit the number of irradiation directions in the solution. This resulted in a mixed integer linear program (MIP) – one which incorporated both continuous variables (irradiation intensities) and Boolean variables (whether a head position is used or not) in its formulation.

This formulation expresses the entire problem as a single mixed integer programming problem, and is therefore guaranteed to find an optimal solution. However

the solution process requires the use of branch and bound techniques, because of the Boolean variables used in the formulation. This technique is very time consuming for large problems and often cannot generate an optimal solution within a reasonable time period.

3.3 Local Search

While the mixed integer approach expresses the entire model as a single programming problem, it had a number of drawbacks. The method was slow to improve on its initial solution, and in more difficult cases failed to find a feasible integer solution. In order to compare the performance of the mixed integer model, other methods were developed which avoided the use of branch and bound strategies.

Local search techniques are used in a wide variety of optimisation problems and were employed here in an attempt to find good solutions in a minimal time frame. Local search methods aim to alter each gantry position in turn, choosing the best angle for each head before considering another. In this way, the positions around the starting point are successively explored for better combinations of angles. Local search is particularly applicable in this problem, where there is a large number of feasible solutions and no established relationship between the resulting dose distributions.

Local search techniques are simple to formulate, but do not guarantee an optimal solution to the Inverse Planning problem.

3.4 Set Covering Strategies

The first approach taken towards two-phase optimisation was an attempt to replicate the current (human-driven) planning process. The initial goals were to imitate a planner's decision strategies with the aim of producing a number of good plans, which could then be evaluated and compared.

Set covering is a simple and intuitive model for the Inverse Planning problem, as it provides an easy method for integrating all possible angles into the formulation. Set covering techniques aim to cover the tumour with the required lethal dose, while minimising damage to healthy organs. This aim also implicitly means that the lowest possible radiation dosages will be used to 'cover' the tumour.

This model is a simplification of the complete model expressed in Section 2.3. The tumour is irradiated with required intensity, rather than optimising the deviation from dose bounds in conjunction with healthy organs. This means that – unlike an LP formulation - a lethal dose must *always* be applied to the tumour. This does not necessarily represent the true planning process, where the tumour is more likely to be 'under-covered' in order to limit treatment side-effects and maintain patient quality of life.

While this model is not guaranteed to provide optimal angles to the master LP problem, it can be used to obtain good directions with limited computational resources.

3.5 LP Relaxation Model

The final approach taken was to investigate the possibility of using information from the mixed integer (MIP) method to generate solutions. The MIP strategy was very effective at generating good solutions, but did so unreliably and with few improvements in objective. In particular, the *LP relaxation* of the mixed integer method was used to generate solutions in this model.

The LP relaxation model incorporates every possible gantry position in the original formulation. The output from this problem represents the optimal solution if a large

number of insertion points could be made from every possible position in the set λ . This formulation is used as part of the mixed integer method to provide a ‘lower bound’ on the solution. In most cases, between ten and forty distinct gantry positions are generated as part of the solution process.

The LP relaxation produces solutions which nearly always meet all constraints, with no damage above allowable limits. Unfortunately, using a large number of insertion points is generally not possible due to limited treatment times. Between three and seven gantry positions are standard, given the time restrictions on patient treatment. However the LP relaxation is valuable as a first approximation, indicating angles which could possibly produce good solutions.

For this reason, a heuristic was designed to explore the solutions obtained by the LP relaxation model. The assumption made is that angles featuring in the LP relaxation will also be ‘good’ angles, when a smaller number of insertion points are allowed.

4 Summary of Results

4.1 Angle Optimisation vs. Planner Logic

This project produced a number of methods for angle optimisation of the Inverse Planning problem. Each of the methods produced some results which were inconsistent with current treatment strategies, and conventional planner logic.

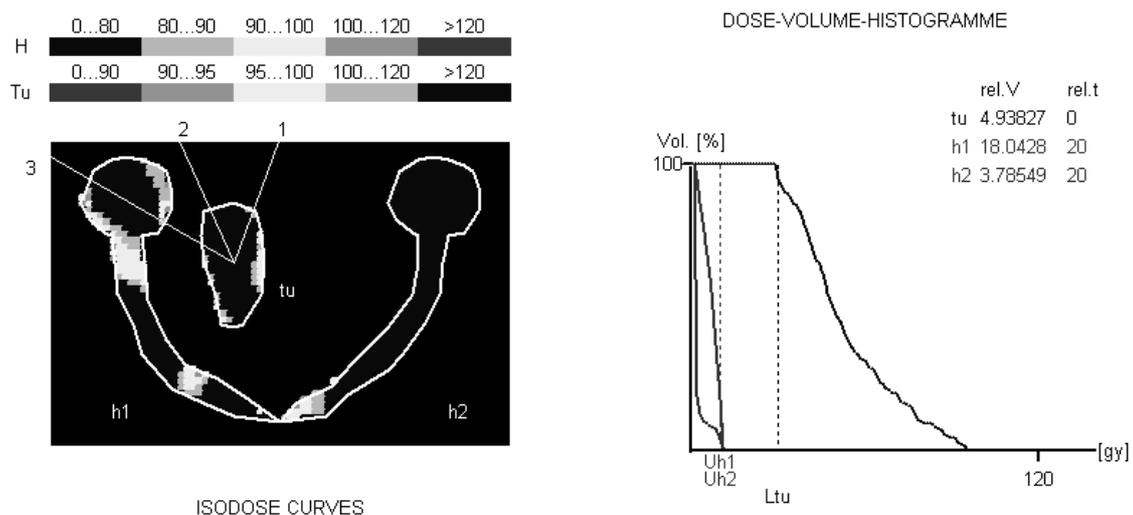


Figure 2. Solution indicating directions for treatment and a dose-volume histogram.

Figure 2 shows a treatment strategy which has completely satisfied all constraints for the problem. However, the angles for treatment are counter-intuitive; the third beam passes directly through the eyeball and nerve (top left). Damage is avoided by altering radiation intensities across the beam, to avoid irradiating the eyeball.

These results suggest that there are a number of improvements which can be made to the current design of radiotherapy plans, in order to maximise patient welfare. Optimisation software has the potential to be a powerful tool for radiotherapy planners because it considers counter-intuitive strategies as well as those traditionally used.

4.2 Objective Quality

The following points were noted with regard to the solution quality obtained by the various methods:

- The mixed integer method descended to solutions with good objectives most quickly. However in most cases the initial solution found did not improve, despite very long run-times. In most cases heuristic methods performed as well or better for longer run-times (see Figure 3).
- Of the heuristic methods trialled, the local search formulation converged consistently to the best solutions. However the gap between the local search and other methods was typically very small, within around 10%. This means there is very little practical difference between the various heuristics in terms of quality (disregarding time).

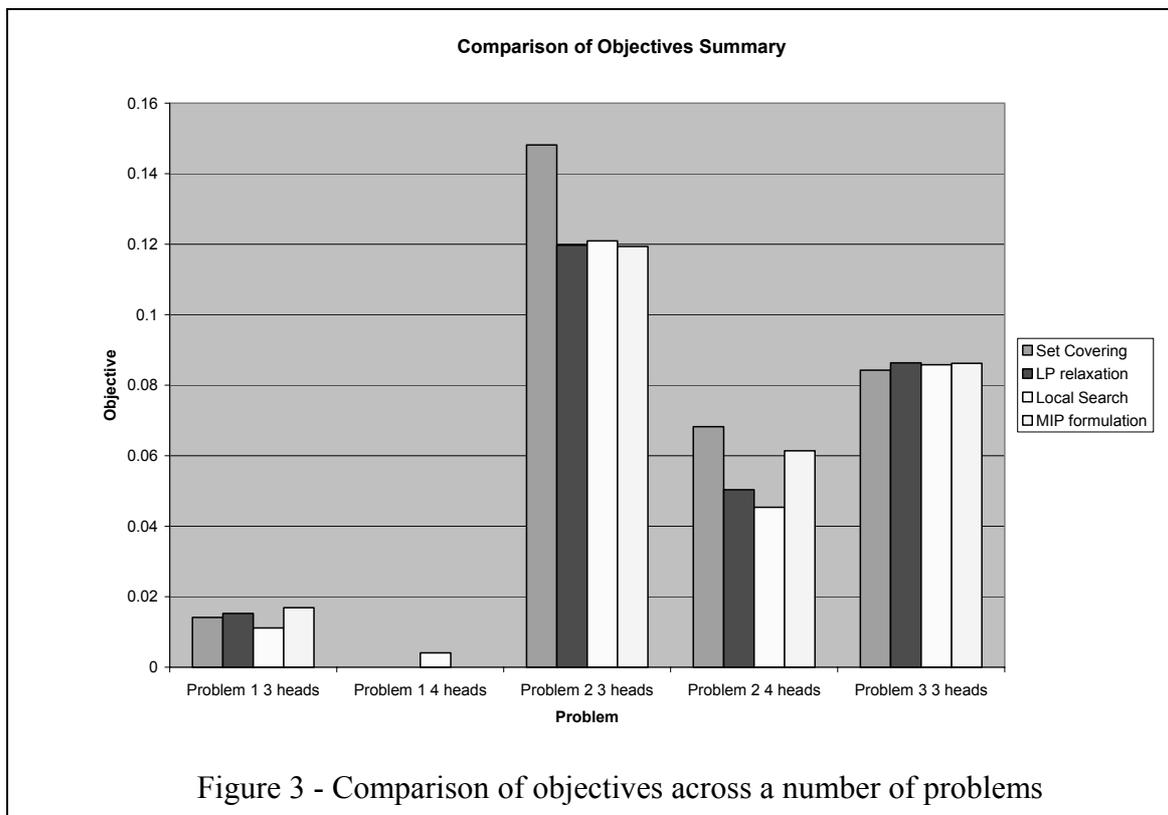


Figure 3 - Comparison of objectives across a number of problems

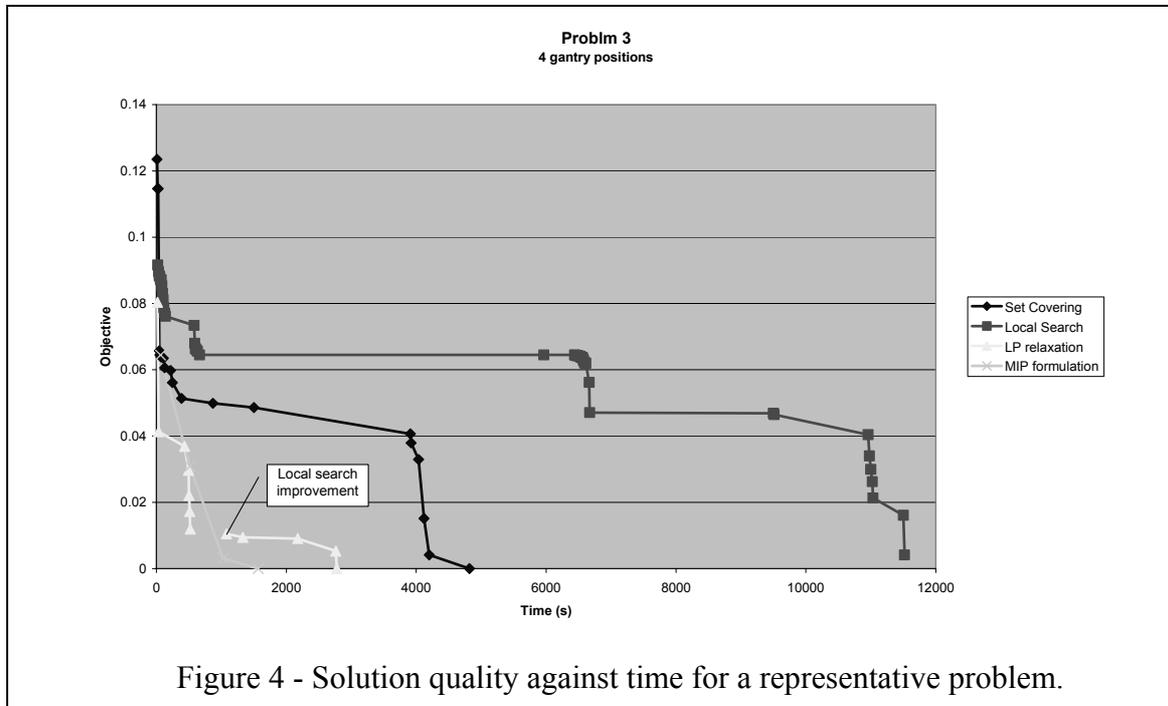
4.3 Speed of Descent

While the mixed integer formulation provided both superior objective quality and speed for small runtimes (100-600 seconds), this is largely due to the fact that this method ran on a much quicker solver (CPLEX). The performance of CPLEX is hundreds of times better than the non-commercial solvers which the other methods were trialled on.

This suggests that – while the heuristics trialled could potentially descend to solutions much more quickly on commercial LP packages – there are few improvements that could be made to the MIP method.

Attention then turns to the other methods considered, with the mixed integer model providing a point of reference. Figure 4 is a representative sample of the data collected, indicating that:

- The local search method was considerably slower in improving the quality of the solution, although it did produce consistently good solutions
- Set covering descended to a better solution more quickly than the local search method in most cases, but did not always generate solutions of the same quality
- Of the heuristic approaches, the LP relaxation method provided the fastest decrease in objective with performance similar to the MIP formulation.



4.4 Conclusion

All of the optimisation methods discussed in this project represent significant improvements over ‘intuitive’, or design reuse strategies in their placement of beams. While optimality was seldom declared in more difficult problems, very good plans were quickly obtained. Optimisation of gantry position provides dramatically improved treatment outcomes from intensity optimisation alone, giving planners another method of juggling between dosage to cancerous tumours and dosage to healthy organs.

While a mathematical model can never take into account all possible factors in the modelling of a problem, the model proposed by this project does capture the essential factors in planning: tumours, healthy organs and radiation. In this way, optimisation allows planners to focus on the more important aspects of Inverse Planning, avoiding the currently used ‘trial-and-error’ approach.

5 References

- Burjony, M., *Balanced Solution of the Inverse Problem in Cancer Radiation Therapy*, University of Auckland, 2001
- Hamacher, H.W. & Küfer, K.H., *Inverse Radiation Therapy Planning – a multiple objective optimisation approach*, Discrete Applied Mathematics, Vol 118, pp.145-161, 2002
- Tervo, P. & Kolmonen, P., *A model for the control of multileaf collimator and related inverse planning algorithm*, University of Kuopio, 1999