Characterisation and evaluation of a novel transmission detector for intra-fraction monitoring of radiotherapy.

DANIEL JOHNSON

Submitted in accordance with the requirements for the degree of

Doctor of Philosophy

THE UNIVERSITY OF LEEDS

School of Medicine

UNIVERSITY OF LEEDS

November 2018

For Mum and Dad

The candidate confirms that the work submitted is their own, except where work which has formed part of jointly authored publications has been included. The contribution of the candidate and the other authors to this work has been explicitly indicated below. The candidate confirms that appropriate credit has been given within the thesis where reference has been made to the work of others.

The work in Chapter 4 of the thesis has appeared in publication as follows:

A simple model for predicting the signal for a head-mounted transmission chamber system, allowing imrt in-vivo dosimetry without pretreatment linac time, Journal of Applied Clinical Medical Physics 15(4), 270279.
Johnson, D., Weston, S. J., Cosgrove, V. P. & Thwaites, D. I. (2014)

The author was responsible for the collecting and analysing of the data and writing of the text. Weston, S. J., Cosgrove, V. P. & Thwaites, D. I.
contributed to discussion on the data and analysis as well as comments on the manuscript.

Conference Abstracts

Initial evaluation of the Device for Advanced Verification of IMRT Deliveries as an in-vivo dosimeter
D. Johnson, V. Cosgrove, S. Weston, D.I. Thwaites
Proc 11th Biennial ESTRO Physics and Technology conference, London, May 2011, Radiother Oncol 99 suppl 1 (2011), 363

Modelling the response of a LINAC-head mounted multi-wire transmission detector for prediction of signals for IMRT plan verification
D. Johnson, V. Cosgrove, S. Weston, D.I. Thwaites
Proc. ESTRO 31, Barcelona, May 2012, Radiother Oncol 103, suppl 1, 2012, 332

Comparison of dose calculations performed by OMP, DIAMOND and RadCalc for 11 clinical VMAT plans
D. Johnson, V. Cosgrove, S. Weston, Sm.M Dawoud, D.I. Thwaites
Proc ESTRO 31, Barcelona, May 2012, Radiother Oncol 103, suppl 1, 2012, 516

In-vivo verification of IMRT delivery using a transmission detector, not requiring pre-treatment time on a linac
D. Johnson, V. Cosgrove, S. Weston, D.I. Thwaites
Proc. ESTRO 2nd forum, Geneva April 2013, Radiother Oncol 106, suppl 2, 436

Direct in vivo verification of IMRT delivery using a 2-D transmission detector, without requiring reference to pretreatment measurements or additional pre-treatment linac time
Johnson D, Cosgrove V, Weston S, Thwaites DI
Proc Int Conf Medical Physics, Brighton Sept 2013, Med Phys Int,

The use of a novel transmission detector to perform in vivo measurements, without the need for any preparation time on the linac
D. Johnson, V. Cosgrove, S. Weston, D.I. Thwaites
Proc Eng Phys Sci Med, Perth, Nov 2013, Australas. Phys Eng Sci Med (2014) 37, 214

Evaluation of transmission detector model using Monte Carlo simulation of VMAT delivery
D. Johnson, V. Cosgrove, S. Weston, D.I. Thwaites
Proc ESTRO 35, Torino, Apr-May 2016, Radiother Oncol 119 (2016), suppl 1, 695

Acknowledgements

x

Que Será, Será.

Doris Day, 1956

Abstract

The goal of this thesis was to characterise a novel transmission detector in the context of signal prediction. This was to eliminate the need to collect a baseline signal for the device before treatment. This not only saves time, but, by independently generating the baseline signal, the process is less prone to missing errors.

A simple analytical algorithm was designed and was found to be capable of detecting gross errors, however, it was shown not to be accurate enough to detect MLC position errors that could have a clinical effect on the delivery. MU check software was commissioned, however the fluence distribution it produced lacked the complexity for accurate signal prediction. A Monte Carlo model of a linac was built and validated then used to demonstrate that the detector could be modelled as two slabs of Perspex; the signal being proportional to the dose measured in the air between them. Two Monte Carlo models were then made using different systems, these were both evaluated by comparing predicted signals to measured signals for VMAT plans. Both models performed well and were capable of detecting leaf errors $\sim 1mm$; the merits of both are discussed with regard to error detection and ease of use.

Abbreviations

#	Fraction
3DCRT	3 Dimensional Conformal RadioTherapy
BED	Biologically Effective Dose
BOC	Beatson Oncology Centre
CAX	Central AXis
CM	Component Module
CNS	Central Nervous System
CPDF	Cumulative Probability Distribution Function
CT	Computed Tomography
CTV	Clinical Tumour Volume
DAVID	Device for Advanced Verification of IMRT Deliveries
DBS	Directional Bremsstrahlung Splitting
DMLC	Dynamic MultiLeaf Collimator
DTD	Dolphin Transmission Detector
DVH	Dose Volume Histogram
EGS	Electron Gamma Shower
EPID	Electronic Portal Imaging Device
EUD	Equivalent Uniform Dose
FFF	Flattening Filter Free
GTV	Gross Tumour Volume
H&N	Head and Neck
IGRT	Image Guided Radiotherapy
IMAT	Intensity Modulated Arc Therapy
IMRT	Intensity Modulated Radiotherapy
IVD	In-Vivo Dosimetry
LIDA	Leeds Institute for Data Analytics
linac	Linear Accelerator
LTH	Leeds Teaching Hospital
MARC1	Medical Advanced Research Computer 1
MC	Monte Carlo
MLC	Multi Leaf Collimator
MLCi2	Brand name for Elekta's 40 Leaf pair treatment head
MU	Monitor Unit
NSRH	North Staffordshire Royal Hospital
NTCP	Normal Tissue Complication Probability
OAR	Organ At Risk
ODI	Optical Distance Indicator
PDD	Percentage Depth Dose
PTV	Planning Treatment Volume
QA	Quality Assurance
SAD	Source to Axis Distance
SSD	Source to Surface Distance
TLD	ThermoLuminescent Dosimeter
TPS	Treatment Planning System
VMAT	VoluMetric Arc Therapy
XVI	X-Ray Volume Imaging
XVMC	X-ray Voxel Monte Carlo

Contents

---

[1]This work was done by many members of the physics team at Leeds Teaching Hospital (LTH). Though the author was involved in collecting beam data and verifying the beam model, the work was done for the clinical service and the burden shared throughout the department

List of Figures

List of Tables

Outline of thesis

The first chapter gives the reader a general introduction to radiotherapy and the requirement for accuracy in its delivery. In-vivo dosimetry's role in the QA framework and how it gives the hospital physicist the confidence that this accuracy is being achieved is discussed at the end of the chapter. Chapter 2 considers the general case for IVD and the devices that are currently available. The research question is posed and the novel work completed by the author for the purpose of this thesis is presented in Chapters 3 to 6. The chapters are arranged chronologically, a summary of each is included below. Presenting the material in this manner is designed to assist the reader in understanding how decisions regarding the direction of the work were reached. As with all research sometimes negative results are recorded. For example, the work outlined in Appendix B did not yield useful results, yet it helped inform decisions that led to the work in the subsequent chapters hence its inclusion. Chapter 7 includes a discussion on this research, how it sits with current radiotherapy practice and suggestions on future work.

- **Chapter 1** *Introduction to Radiotherapy* External beam radiotherapy is introduced. The basic radiobiology that underpins the subject is outlined and the requirement for a high level of dosimetric certainty in the delivery of treatments is demonstrated. The key components of the modern linear accelerator are outlined including the Multi-Leaf Collimator (MLC) and how it has been used to develop Intensity Modulated Radiotherapy (IMRT). While IMRT affords much more conformal dose distributions, small errors in leaf position can lead to much larger dosimetric errors than traditional 3D Conformal radiotherapy. MLC QA, pre-treatment verification and in-vivo dosimetry can be used together to give the radiotherapy physicist confidence that the linac is delivering treatments that maintain suitably high levels of dosimetric accuracy to ensure that clinical objectives are realised.

- **Chapter 2** *Literature review* In-Vivo Dosimetry (IVD) is reviewed. The case is made for IVD by considering:

  - The high level of accuracy required in advanced radiotherapy deliveries
  - The unfortunate reality that gross errors have occurred

  A synopsis of the current literature pertaining to IVD in the UK is given. Techniques for performing IVD including point detectors, electronic portal imaging and log file analysers are reviewed. Particular attention is given to transmission detectors with individual products discussed. The Device for Advanced Verification of IMRT Deliveries (DAVID) - a multi-wire transmission detector and the subject of this thesis - is introduced and the current

literature on the device is reviewed. The requirement for an algorithm to predict the DAVID signal for the purposes of efficiency and safety is demonstrated leading to the research objective:

> **Investigate the characteristics of the DAVID device with the intention of being able to predict the response to advanced deliveries in order to facilitate intra-fraction monitoring of radiotherapy that is both efficient and safe.**

- **Chapter 3** *Preliminary work for the predictive algorithms for the DAVID* The DAVID is shown to be a stable device with a response that is linear in relation to both delivered Monitor Units (MUs) and field size. This demonstrates that, in principle, it should be possible to predict the signal provided enough is known about the field. Two algorithms are developed one relying on a look-up table, the other using the linear response of the primary and secondary response to leaf separation. These are tested using basic MLC apertures with the linear-response algorithm predicting the signal more accurately than the look-up table approach.

- **Chapter 4** *An analytical model for predicting the DAVID signal* The linear-response algorithm outlined in chapter 3 is developed to include the effects of varying MU, asymmetric leaf separation, scatter, jaw position and penumbral effects. The increased complexity of the algorithm allows it to predict the response of the DAVID to clinical IMRT treatment deliveries. This work was published during the course of the thesis (Johnson et al. 2014). It is shown that the uncertainty associated with the out-of-field response is a limiting factor in the algorithm's ability to predict the DAVID response. It is postulated that a dose or fluence map of the delivery would have the potential to predict the response more accurately.

- **Chapter 5** *Monte Carlo Modelling of the DAVID* An EGSnrc-based Monte Carlo (MC) model of the Elekta MLCi2 head is developed and matched to machines used clinically at Leeds Teaching Hospitals. At the time the chapter was written, there was little published literature on MC modelling of the MLCi2 head, so this work was novel as well as being necessary to confirm an accurate model of the DAVID.

  The linac model was used to show that the collection wires had a 3% impact on the signal collected by the DAVID. This effect could, however, be accounted for using a multiplicative correction factor. Using this information a simple model of the DAVID - that did not include collection wires - was created in EGSnrc. The combination of the DAVID model and the MLCi2 model were both calibrated and used to predict the response of five clinical VMAT fields. Modelling of the DAVID in MC and the method used to calibrate a transmission detector are both novel, it is the authors intention to publish this work in the near future.

  An equivalent model of the DAVID was also constructed as a QA phantom in the Monte Carlo-based treatment planning system Monaco. This was used to predict the response of the same five VMAT fields.

Both Monte Carlo approaches were shown to be capable of accurately predicting the DAVID response to a VMAT delivery. While the EGSnrc approach was slightly more accurate, this comes at the expense of both time and effort if applied routinely in the clinic.

- **Chapter 6** *The performance of the models in detecting MLC errors using the DAVID* There is a review of the current literature regarding MLC errors and their impact on the dosimetric accuracy of VMAT deliveries. The general consensus is that leaf bank errors that either increase or decrease leaf bank separation have the most severe impact on delivered dose. Five VMAT plans were modified to include these errors with various magnitudes. The analytical model devised in Chapter 4 and the two Monte Carlo models developed in Chapter 6 were used to predict response of the original and modified plans. These predictions were compared to measured signals to give an indication of the algorithm's sensitivity to the simulated errors. The Monte Carlo approaches were shown to be almost equivalent and considerably more sensitive than the analytical approach. Using Monaco to model a transmission detector is novel as is the MLC error testing of the DAVID in this context. The author intends to publish this work on the completion of the thesis.

- **Chapter 7** Discussion of findings, conclusions and suggestions for future work.

A key contribution of the work to the field of radiotherapy physics is a robust approach to implementing IVD that is not only efficient, but sensitive to gross treatment errors and more subtle leaf errors that have been shown to have a significant impact on the dose delivered by advanced (e.g VMAT) techniques. Additionally, I developed two novel calibrated Monte Carlo models for transmission detectors which could be used to predict the expected measurements.

Chapter 1

Introduction to Radiotherapy

In 2014 359,960 people were diagnosed with cancer in the UK (CRUK 2015). Radiotherapy was used to treat approximately 50% of these cases. 40% of all cured cancers are treated with radiotherapy, 16% of all cured cancers are treated with radiotherapy alone (CRUK 2014). This chapter provides an outline of the biology underpinning radiotherapy and how the response of tissue to radiation demands a certain level of accuracy when delivering treatments. The basic principles of linear accelerators are discussed and how modern advances are enabling more conformal treatments. The chapter concludes by demonstrating that these new techniques must be delivered accurately and then looks at the tools that are used to ensure that this is the case.

1.1 Biological basis and accuracy requirements for radiotherapy

1.1.1 Underlying principles

Radiotherapy is the treatment of cancer using ionising radiation. Exposure of tissue to ionising radiation causes the production of free radicals, these are highly unstable and react almost immediately ($\sim$ milliseconds) with nearby molecules, transferring chemical damage to them. In the context of radiotherapy, the energy deposited by ionising radiation is referred to as dose and is given the unit of Gray, defined as the absorption of one joule of radiation energy per kilogram of matter (BIPM 2006). In the case of exposing tissue to ionising radiation, cells will be damaged. If the damage is sufficient it will initiate a DNA Damage Response (DDR) that can lead, through various mechanisms, to cell death.

The proportion of cells that are damaged as a function of dose is characterised by the linear quadratic model:

$$L = \alpha d + \beta d^2 \tag{1.1}$$

Where L is the total number of cell deaths α and β are constants to be discussed shortly. Using the Poisson model, the surviving fraction (S) can be considered as

the probability of no cell deaths giving (Jones et al. 2001):

$$S = e^{-L} \tag{1.2}$$

Or:

$$S = e^{-(\alpha d + \beta d^2)} \tag{1.3}$$

Radiobiological studies have shown (Joiner & Van der Kogel 2016) that each successive fraction in a multidose schedule is equally effective, so the effect of n fractions can be expressed as:

$$S^n = (e^{-(\alpha d + \beta d^2)})^n \tag{1.4}$$

Taking the natural logorithm of both sides leads to:

$$E = n(\alpha d + \beta d^2) \tag{1.5}$$

Where E is $-\ln S$ and is conventionally referred to as the 'log cell-kill' (Dale et al. 2007).

The linear quadratic model shown in Equation 1.1 was initially an empirical fit to describe chromosome damage to cells as a consequence of irradiation (Lea & Catcheside 1942). However, it has been shown to have solid bio-physical basis. Damage to the cell's DNA from the free radicals produced through exposure to ionising radiation is regarded as the principle mechanism that brings about cell death (Mayles et al. 2007). Single breaks in the DNA can normally be repaired by intra-cellular enzymes, well-separated double breaks are repaired in a similar way; double breaks close together are not easily repaired. Sufficiently damaged DNA can trigger apoptosis - programmed cell death or check point activation - where mitosis is inhibited (Dale et al. 2007, Joiner & Van der Kogel 2016). The α term in the LQ model is associated with lethal damage caused by single-strand breaks while the β term is associated with lethal damage caused by double-strand breaks (Dale et al. 2007, Jones et al. 2001).

If we take Equation 1.5 and consider the situation when the total dose D is delivered through a large number of very small fractions we see that the linear term becomes dominant and the quadratic term can, essentially, be ignored:

$$E = n\alpha d \quad = \alpha D \text{ as } \quad d \to 0 \tag{1.6}$$

Here, dose D, is the dose required to give a specific effect; in this case D is defined as the Biological Effective Dose (BED). While BED represents the physical dose required for a given effect for the case where $d \to 0$ it can be achieved through more realistic dose / fractionation regimes (Dale et al. 2007). From Equation 1.6 we see that:

$$BED = \frac{E}{\alpha} \tag{1.7}$$

Combining Equation 1.7 and Equation 1.5 it can be shown that:

$$BED = nd\Big(1 + \frac{d}{\alpha/\beta}\Big) \tag{1.8}$$

Equation 1.6 demonstrates that a fractionated treatment regime that delivers a higher total dose through a number of equally weighted fractions can have the same biological effect as a single large-dose fraction given in one go. Fractionating a treatment exploits inherent differences in the radiation response between carcinogenic and normal tissue with beneficial consequences. When considering fractionated treatments it is useful to consider the 5 Rs of radiobiology:

- Repair of sub-lethal damage - Typically much more active in normal tissue and by limiting the dose to normal tissue the repair is more likely to work.

- Re-assortment of cells within the cell cycle - Cells have varying radiosensitivities depending on where they are in the mitotic cycle. Fractionating the treatments means that more cells will be exposed to radiation when they are in their most vulnerable cell phases. This effect acts more strongly against carcinogenic cells as they typically divide at a much higher rate than normal tissue.

- Re-population - Cells that have undergone sub-lethal damage will go on to proliferate and repopulate, this is the case for both normal and carcinogenic tissues. In the case of treatments that continue for longer than the cell cycle duration, it will be necessary to deliver a higher dose in order to maintain the same level of tumour control.

- Re-oxygenation - Tumour tissues do not normally have particularly good blood supplies to all of the cells. During the course of a radiotherapy treatment the tumour shrinks, revascularisation occurs; previously hypoxic regions of the tumour can be reoxygenated. Though this may help repopulation, the presence of oxygen also increases the potency of radiation. Typically in more oxygenated areas there are more double strand breaks, this is considered to be due to the production of highly-reactive oxygen free radicals. The oxygen effect is thought to be one of the main reasons why fractionated radiotherapy has better outcomes than high single-fraction treatments (Dale et al. 2007).

- Radiosensitivity - how tissue responds to treatment depends on its radiosensitvity. Some carcinogenic tissue is damaged more effectively with fractionated treatment than others. At present there are a number of studies looking at how dose fractionated regimes can be tailored to better suit specific cases (Morrison et al. 2018, Lawless et al. 2017, Fisher & Rabinovitch 2014, Baker et al. 2016). This remains an interesting and active area of research.

1.1.2 Accuracy and precision in radiotherapy

Recent work in the rapidly-developing field of immunotherapy has demonstrated that the use of certain monoclonal antibodies can enhance the effect of radiotherapy by inhibiting tumour repair pathways or by triggering an immune response

(Farkona et al. 2016, Chalmers 2016). However, for simplicity, the classic target-cell hypothesis of radiotherapy shall be used to illustrate the key concepts of fractionated deliveries. This states that any single surviving cell has the capacity to proliferate, therefore every clonogenic cell must be killed for the treatment to be successful (Munro & Gilbert 1961, Mayles et al. 2007). Using Poisson statistics, it can be shown that, given an average number of clonogenic cells N, the probability of their being no surviving cells is:

$$TCP = e^{-N} \tag{1.9}$$

where TCP is the tumour control probability that is equal to 1 when there are no remaining cancerous cells. Now taking the Equation 1.4 - that describes the surviving fraction of cells for a fractionated treatment - it can be shown that:

$$N = -N_0 \times e^{-n(\alpha d + \beta d^2)} \tag{1.10}$$

Where N_0 is the initial number of clonogenic cells. Combining Equation 1.9 with Equation 1.10 generates Equation 1.11, describing the TCP. A plot of TCP as a function of dose using typical values is shown in Figure 1.1.

$$TCP = e^{-N_0 \times e^{-n(\alpha d + \beta d^2)}} \tag{1.11}$$

The likelihood of inflicting damage to normal tissue depends on how much of the
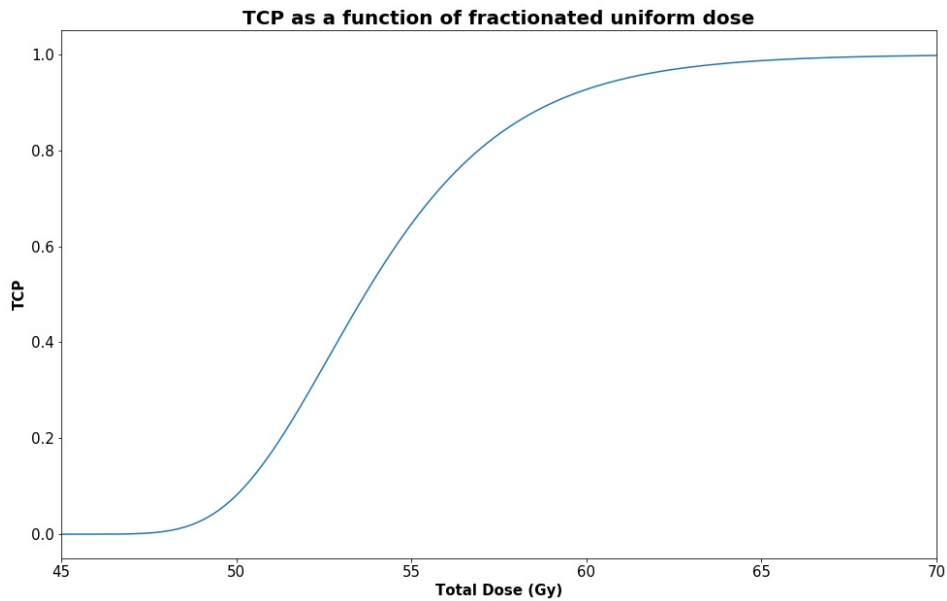
Figure 1.1: TCP curve for a 2Gy per fraction (#) treatment regime using typical values $\alpha = 0.286$, $\beta = 0.032$ and $N_0 = 10^8$ (Dale et al. 2007, Mayles et al. 2007, Brahme 1984, Joiner & Van der Kogel 2016). The gradient of the curve around the treatment dose determines the effect a dosimetric error would have on the TCP

organ is irradiated and whether the organ is more serial or parallel in character (ICRU 1999), making the tissue more susceptible to point or average dose respectively. There are a number of mathematical models describing Normal Tissue Complication Probability (NTCP) (Adamus-Górka et al. 2011), the one by Lyman

(1985) is the most widely cited (Equation 1.12)

$$P(D,V) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{t} exp\left\{-t^2/2\right\} dt \qquad (1.12)$$

Where $P(D,V)$ is the NTCP and t is the upper limit of the normally distributed probability function defined as:

$$t(D,V) = \frac{D - D_{50}(V)}{mD_50(V)} \qquad (1.13)$$

and

$$D_{50}(V) = \frac{D_{50}(1)}{V^n} \qquad (1.14)$$

Where D_{50} is the dose associated with 50% probability of a response; $D_{50}(1)$ is the tolerance dose for 50 % complications for uniform whole-organ irradiation; V is the volume fraction being irradiated by dose D; n is specific to the organ and associated with the gradient of its dose response. The Lyman response can be plotted on a three axis graph with z, x and y being V, dose and NTCP respectively. For a given V the NTCP relationship to dose is sigmoidal much like the TCP curve shown in Figure 1.1 (Lyman 1985, Adamus-Górka et al. 2011).

While the Lynman model - and ones like it - are used frequently in research, they are rarely used in the clinic. The model outlined here, for example, assumes that the OAR sub volume receives a uniform dose - this is generally not observed in practice. Furthermore, the complexity of various organs, the variety of responses they can have to regional irradiation and the limited amount of of good clinical data means that applying any mathematical model to determine NTCP will always be difficult (Joiner & Van der Kogel 2016). What can be certain is that irradiating healthy organs carries with it an associated risk, increased dose to larger volumes increases this risk. Clinical decisions and tolerances associated with specific OARs tend to be based on maximum doses, maximum volumes receiving certain doses and interpretation of the Dose Volume Histogram (DVH) derived from experience and/or national/international guidelines and/or clinical trials. With this approach in place, minor unavoidable toxicities are frequent (eg. cataracts, erythema, rectal bleeding). However, severe OAR radiation responses (eg organ failure, paralysis) are rarely seen and when they are they are normally anticipated. In this instance the patient is consented to this possibility to enable an informed decision on the course of their treatment.

Tumour control is achieved by delivering dose to the tumour, however the maximum dose delivered is limited by the risk to normal tissue. The necessary dosimetric accuracy of these deliveries is determined by the gradient of the TCP and/or NTCP curves around the dose that is being delivered. Deviation from the anticipated dose will either lead to lower TCP or higher NTCP than anticipated, in both cases resulting in poorer clinical outcome. Older (Brahme 1984) and more recent (QUANTEC 2010) studies have agreed on the gradients associated with TCP; consideration of these values by Brahme (1984) showed that the most critical loss in tumour control was found when dosimetric inaccuracies were introduced at the highest level of TCP. Given this mathematically rigorous analysis it is agreed

that delivered dose should be accurate to within 3% (relative SD). Meaning that there is a 95% chance that dosimetric changes at twice this level will be clinically observable (Thwaites 2013, Brahme 1984)

1.2 The Linear Accelerator

1.2.1 Introduction

The most common energy band and modality in worldwide use for cancer treatment is MV photon radiation as these beams can deposit dose at clinically-appropriate depths. This is now almost exclusively generated by linear accelerators. The most common type of radiotherapy treatment unit is the c-arm style linear accelerator (Figure 1.2). These are capable of generating MV radiation over the range of 4-25MeV. Throughout the rest of this thesis the term radiotherapy will be with regard to external beam radiotherapy delivered by a linac. The technical aspects associated with generating this radiation are discussed in the following section.

Figure 1.2: An Elekta Versa HD$^{\text{TM}}$

1.2.2 Generating Megavoltage (MV) radiation

To produce therapeutic electromagnetic radiation accelerator technology is used – this is where electromagnetic waves with specific properties are generated and utilised to accelerate the electrons to high speeds before crashing them in to a target, producing electromagnetic radiation through the Bremsstrahlung process. Throughout the generation of the therapeutic beam, there are a number of feedback and control mechanisms to ensure that the beam passing through the shaping

apparatus is stable, symmetric and at the right energy (Waldron 2002) giving confidence that that treatment is being delivered with appropriate accuracy. What follows is brief discussion on the key components of a linac, summarised in Figure 1.4.

- The electron gun produces electrons through thermionic emission by using a current to heat a filament. These electrons are then accelerated to speeds $\sim$0.5c using a voltage.

- At the same time as this is happening microwaves are produced by either a magnetron or klystron.

- The microwaves are directed up to the accelerating waveguide.

- Shortly after the microwaves enter the waveguide the electrons are injected.

- The electrons are accelerated along the waveguide to speeds of $\sim$ c.

- Electromagnets are used to direct the electron beam on to a high-Z target.

- On hitting the target the electrons slow down and high energy photons are produced through the Bremsstrahlung process, these pass through the linac head (Section 1.2.3) before reaching the patient.

- In practice the radio-frequency power required to accelerate electrons to such energies cannot be sustained, so linacs operate in a pulsed mode. This pulsing is controlled by the modulator that distributes high voltages and currents to the linac components.
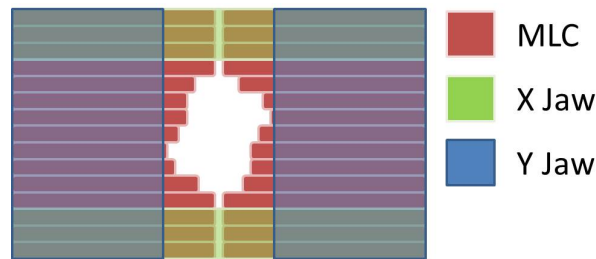
Figure 1.3: Typical MLC-jaw set up. The MLC conforms the radiation to the desired shape; the X jaws reduce transmission through the gaps between "closed" leaf pairs, as there is a gap between the closed leaves the X jaws need to attenuate the primary beam to acceptable levels before it hits the patient making them necessarily thicker. The Y jaws back up the MLC reducing inter-leaf leakage. Both X and Y jaws reduce leaf transmission.
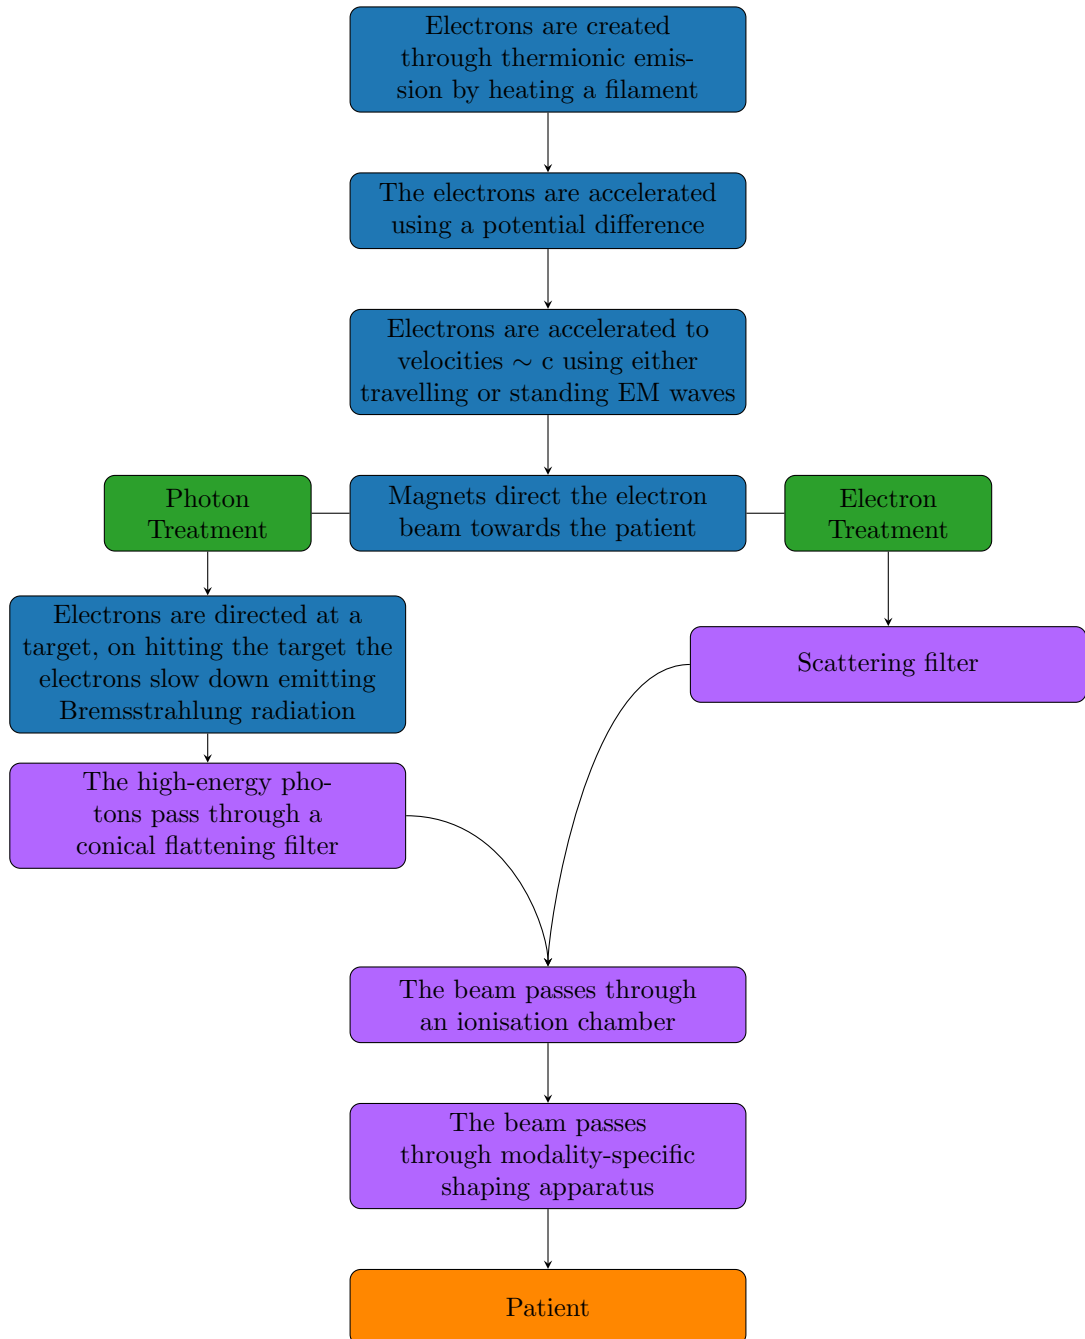
Figure 1.4: The process of generating therapeutic radiation. Although the whole unit is typically called a linac, the accelerating part, held under vacuum is shown in blue. The linac head, responsible for shaping the radiation, is shown in purple.

1.2.3 The linac head

The linac head (Figure 1.5) shapes and monitors the radiation in accordance with the radiotherapy plan and the linac's own control system. The architecture of the head is specific to both the model and vendor. Typically a modern treatment head will contain the following:

- Primary collimator – attenuates peripheral, high angle scatter radiation produced in the target. Ultimately defines the maximum size of the radiation field.

- Flattening filter – differentially attenuates the peaked radiation profile from the target so that the dose distribution is approximately flat at 10cm deep in water.

- Segmented ion chamber – checks the beam output, symmetry and flatness.

- Wedge – can be used to preferentially attenuate the radiation profile to better improve conformity.

- Mirror and filament lamp – produces an optical field that matches the radiation field to assist with patient setup and qualitative component assessment.

- Primary jaws – Set the extent of the field in the direction perpendicular to the MLC movement.

- Backup jaws - Move in the same direction as the MLCs reducing leakage (Figure 1.3).

- Multileaf collimator – shapes the beam exit aperture; adjusts the radiation field shape to better conform to tumour dimensions.
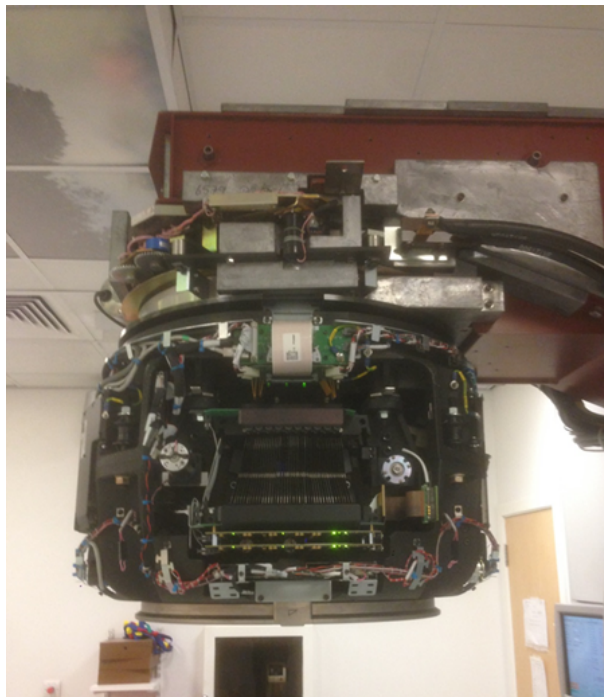
Figure 1.5: An Elekta Versa HD$^{\text{TM}}$ treatment head with the covers removed

1.2.4 The Multileaf Collimator

The Multileaf Collimator (MLC) was developed in the 1980s with Scanditronix (Scanditronix AB, Uppsala, Sweden) (Brahme 1987, 1988a), Siemens (Siemens, Concord, CA) (Boesecke et al. 1988), Varian (Varian Oncology Systems, Palo Alto) (Moeller 1989, Galvin et al. 1993) and Philips (Philips Medical Systems, Shelton) (Jordan & Williams 1991) releasing models with 32, 27, 26 and 40 leaf pairs respectively (Jordan & Williams 1994). Initially the MLC allowed the radiation beam to be conformed around target volumes (AAPM 2001), as shown in Figure 1.6, but the technology paved the way for Intensity Modulated Radiotherapy (IMRT) (Section 1.3). This approach to radiotherapy delivery led to highly conformal dose distributions and new ways to utilise IMRT remain an exciting area of research and development.
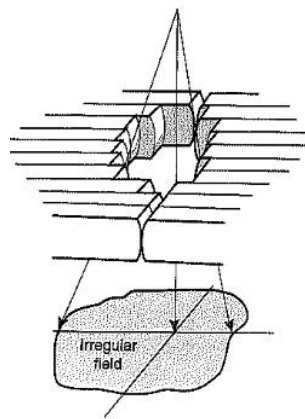
Figure 1.6: MLC used to conform the radiation field around the edge of a treatment volume. Taken from Greene & Williams (1997)
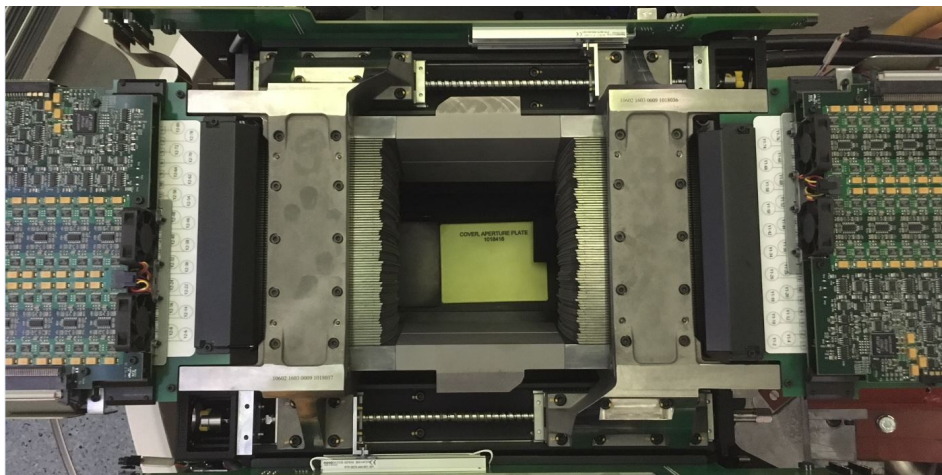
Figure 1.7: A 80-leaf-pair MLC bank from an Elekta Agility$^{\text{TM}}$ treatment head. The view is along the beam axis looking up, through the MLC aperture, to the target.

MLC banks consist of between 20 and 80 tungsten leaf pairs; each leaf can move either forwards or backwards, but is restricted to a single plane. By moving the leaf pairs, irregular apertures can be made for the radiation beam. As an

example, for the Elekta Agility MLC, leaf velocities of 65 mm s^{-1} are possible with the combined effect of leaf movement (35 mm s^{-1}) and leaf-bank movement (30 mm s^{-1}) (Thompson et al. 2014)

Radiation leakage and transmission through the leaf bank are key issues. Adjacent leaves can never be perfectly abutting, resulting in inter-leaf leakage; mechanical reasons prevent leaf pairs from being completely shut, leading to transmission; the leaves are of finite thickness so the beam can never be fully attenuated. There are several methods for combating these various issues (Huq et al. 2002, AAPM 2001).

Inter leaf leakage. Leakage between the leaves is reduced to between 1 and 5% by stepping the leaf shape (Figure 1.8). In addition to stepping the MLCs backup jaws are employed that move in the same direction as the MLCs, they are set to be as far back as the most retracted leaf (Figure 1.3). When using this method the interleaf leakage of the leaves that are not backed up by the jaws is hard to avoid. Another method employed on the new Elekta Agility head is to align the focus of the MLC divergence on a different point to the beam divergence (Thompson et al. 2014).

Elekta   Siemens   Varian
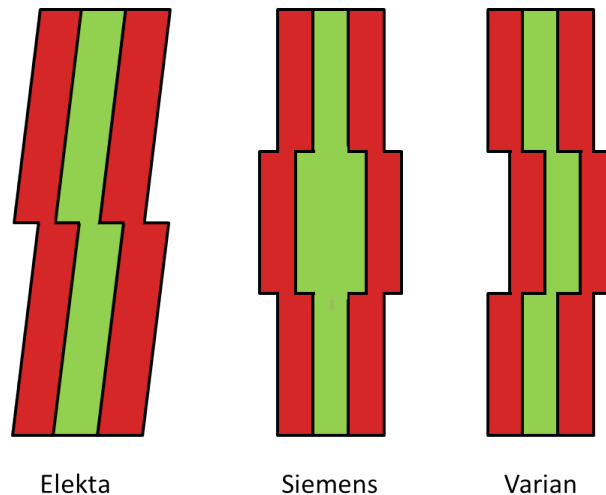
Figure 1.8: Approaches taken by different manufacturers for reducing inter-leaf leakage.

Leakage between the ends of closed leaves. To ensure a consistent penumbra over the range of leaf movement, MLC leaf ends are normally rounded; even if the leaves were to shut, the point at which they touched would offer limited shielding. This problem can be avoided if the closed leaves are not left in the centre of the field but moved to the edge so they are under a jaw. Or a jaw that moves perpendicular to the leaf movement can be used to cover the closed leaves.

Leaf transmission. Technical issues prevent the leaves being so thick that they attenuate all the radiation, a typical leaf thickness is about 7cm; this is sufficient to reduce the intensity of the primary beam to about 1%. This is still not ideal for clinical treatments but the backup jaws serve to limit the overall transmission.

Typically leaves project to between 0.25 and 1cm at the isocenter. The light projection through a shape defined by MLC will have quite jagged edges leading

one to think that the ability of the MLC to conform to the smooth edges normally found in nature would be poor, however, the effect of scatter inside the patient smooths the resulting dose distribution (Figure 1.9).
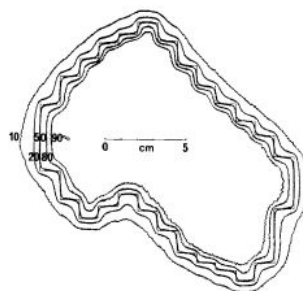
Figure 1.9: Isodose contours produced by a MLC that projects to 1cm leaf widths at the isocenter. The smooth 90% isodose suggests that 1cm leaves are adequate for accommodating for typical biological shapes. Taken from with permission Brahme (1987)

1.3 Intensity Modulated Radiotherapy (IMRT)

1.3.1 Basic Principles

The MLC was initially used to conform the treatment fields to the shape of the treatment site (Figure 1.9). However, the ability to generate small, irregular shapes meant that beams could be built up through the superposition of a number of smaller beams, or segments. The composite beams will have non-uniform fluence distributions. With suitably advanced optimisation algorithms these beams can be designed to have much higher levels of dose conformality and/or better sparing of critical structures than could ever be achieved through 3DCRT planning, see Figure 1.10 (Brahme et al. 1982, Bortfeld 2006, Webb 2003). This is expected to lead to better treatment outcomes, despite generally coming at the expense of an increased low-dose bath - though the benefits of high-dose conformality typically out weigh the risks this brings (Chang 2015). On account of the varying fluence across these beams this technique became known as Intensity Modulated RadioTherapy (IMRT).
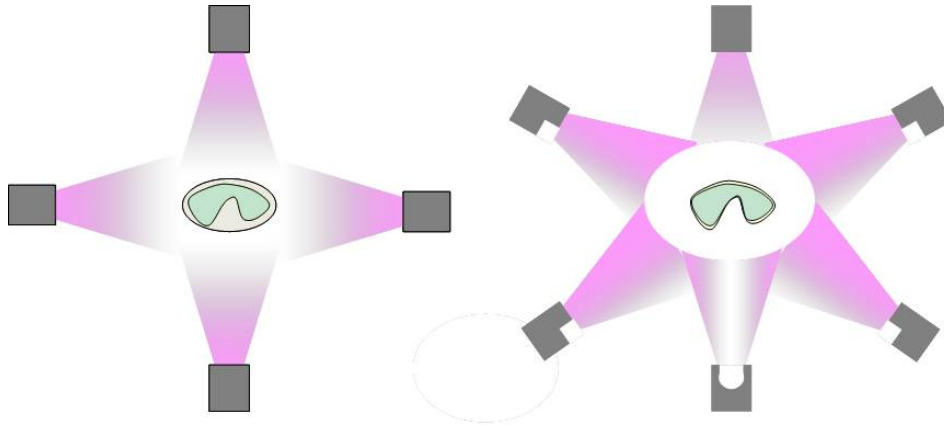
Figure 1.10: Five treatment beams treating a tumour with concave attributes (right-hand picture). The varying intensities of the beams makes the resultant dose conform in a way that would not be possible with conventional treatments (left-hand picture) (Bortfeld 2006)

Using a MLC the intensity of a beam can be modulated two ways: Dynamic modulation: the MLCs are moved while the beam is on or Step and Shoot: the radiation is turned off while the leaves are moved.

Step and shoot The intensity of the beam is modulated by splitting it up in to several segments. A segment is a small shape, defined by the MLCs, with some fraction of the beam MUs being delivered through it. After the MUs are delivered, the radiation is turned off and the MLCs are moved to define the next segment shape before the next fraction of beam MUs are delivered. Any particular segment shape is unlikely to bear much resemblance to the tumour outline but the resultant dose distribution will match the shape of the tumour.

Dynamic IMRT: The beam aperture is changed while the beam is on (Convery & Rosenbloom 1992, Mayles et al. 2007, Cho 2018). The dose is proportional to the time in-between the leading edge of the MLC passing – exposing the volume – to the trailing edge passing – blocking the volume and the dose rate during this time (Figure 1.11). Intensity Modulated Arc Therapies (IMAT)(Palma et al. 2010) are where IMRT is delivered while the gantry is rotating. There are a number of brand names associated with this, but VMAT - initially an Elekta term - has become the generic term. VMAT plans have been shown to have equivalent or superior dose distributions to static arc IMRT techniques, with the main advantage being a reduction in delivery time (Otto 2008, Palma et al. 2010, Ling et al. 2008).
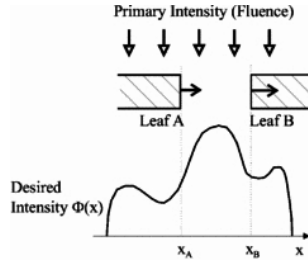
Figure 1.11: Generation of one-dimensional intensity-modulation with one leaf pair in the unidirectional 'sweep' mode (from left to right), the intensity at every point x is proportional to the difference between time $tB(x)$ when the edge of the leading (B) leaf crosses point x and starts the irradiation, and time $tA(x)$ when the trailing (A) leaf crosses over point x and stops the irradiation. Taken with permission from Bortfeld (2006).

Inverse planning 3DCRT treatments are planned by first contouring the CT scan. This involves delineating various features of the disease and organs at risk, as discussed in ICRU (1994) and ICRU (1999). In conventional forward planning, using a Treatment Planning System (TPS), beams are set at different angles determined by a mix of protocol and experience. The planner modifies the beam weighting, MLC shape and wedging in order to get the best dose distribution possible. After each amendment to the set up the plan needs to be recalculated for the new dose distribution evaluated. The plan can be visually assessed or a more quantitative assessment can be performed by analysis of the DVH. IMRT provides significantly increased degrees of freedom and enables inverse planning to be used. Inverse planning involves the plan undergoing the same contouring process, after this though the target doses or dose constraints are defined for the clinical volumes. The TPS then uses iterative methods to achieve the most suitable beam and segment choices so that the dose constraints are met. The iterative process is unlikely to achieve all of the dose objectives and compromises will have to made; what constraints are compromised is determined by clinical importance and appropriate priority can assigned by attaching certain weights to the different constraints. The TPS assesses the plan in terms of a cost function; Equation 1.15 shows the cost function for n dose constraints where D_D is the delivered dose, D_R is the required dose and $W(n)$ is the weight - or importance - of the constraint. The aim of the iterative planning process is to minimise the cost function (Webb 2003, Brahme 1988b).

$$cost = \sum_n W(n) \times [D_D(n) - D_R(n)]^2 \tag{1.15}$$

Treatment volumes will be patient specific, yet the same treatment site in different patients will, in most cases, require a similar planning paradigm. It is this assumption that leads to the development of the class solution; this is a starting point for the iteration. The use of the class solution not only speeds up the computing process but reduces the likelihood of reaching a false optimisation. It has been demonstrated that when plotting the cost function value as a function of iterations, one can sometimes observe local minima, the planning algorithm may perceive these as optimal and end the iterative process. In this way, the starting

point of the iteration affects the finishing point and starting with a good guess increases the chance of an optimal finish (Webb 2003).

The ability to increase the conformality of the treatment dose brings with it obvious advantages. With this in mind, it is now recommended that any patient with locally advanced head and neck (H&N) cancer should be treated with IMRT (SCoR 2015). Following an investment of 23 million pounds by the government, the percentage of patients in the UK receiving inverse-planned IMRT has risen from fewer than 14% in 2012 (Cancer Research 2019) to over 40% nationally with some centres delivering 100% of their radiotherapy with advanced techniques (NATCANSAT 2019).

1.4 QA and Verification of IMRT Dose delivery

IMRT provides the opportunity to deliver more conformal treatments than 3DCRT yet these advanced treatments are more susceptible to geometric errors which, in radiotherapy, translate to dosimetric uncertainty (Thwaites 2013). Furthermore, although delivery techniques may have changed, the dosimetric accuracy requirements discussed in Section 1.1.2 remain the same. It is the role of the multidisciplinary radiotherapy team to ensure the geometric uncertainties associated with radiotherapy are minimised so that IMRT can be delivered with minimal uncertainty, thus ensuring optimal clinical outcomes.

1.4.1 Patient and set-up uncertainty

The patient's position needs to match the CT scan position used for the dose calculation in order for the dose to go where it is intended. To account for variability in patient position and patient size, set up uncertainty and machine variability expansion margins are applied to treatment volumes (Van Herk 2004). However, IMRT dose distributions can be much more conformal, so whereas a 3DCRT plan might have less conformal dose distributions such that a movement of the patient would still result in the CTV being covered by the treatment dose, movement of the patient receiving a highly conformal IMRT delivery could cause the CTV to move outside the treatment-dose volume. This will result in a loss of tumour control and, if OARs are near the treatment volume, an increase in NTCP. X-ray Volume Imaging (XVI) is now standard on new linacs and Image Guided Radiotherapy (IGRT) is becoming increasingly common for radical treatments. The consequence of improved imaging is increased certainty regarding the patient position and position of the OARS and treatment volumes. This gives greater assurance for highly conformal treatments and even allows a reduction in the expansion margins applied to treatment volumes (RCR 2008b).

1.4.2 Beam shaping uncertainty

For 3DCRT deliveries, field shapes are roughly conformed with the projection of the PTV at a specific gantry angle. Tolerances on leaf and jaw position were built in to margin recipes so that the CTV was always covered by the beam. Typically tolerances of 2mm (projection at isocentre) on leaf and jaw position were recommended (AAPM 2009). IMRT deliveries typically contain small segments

and will usually involve the MLC obscuring part of the treatment volume. In the case of a 3DCRT plan, systematic errors in the position of the leaf bank - modifying the position of all the leaves for all the treatment - would affect the edge of the PTV. The same fault in a IMRT plan will affect the dose across the whole target volume, discussed more in Section 6.1. Ultimately this means that, in order to maintain dose certainty with IMRT deliveries, much higher, more stringent tolerances are placed on MLC position accuracy with figures <1mm recommended (IPEM 2018, AAPM 2009, Thwaites 2013). These tight tolerances demand a lot of the linac's leaf positioning equipment[1]. Ensuring this level of positional accuracy, leading to confidence in dose delivery is one a key role of the radiotherapy physicist. This is achieved through:

- MLC Quality Assurance (QA)

- Pre-treatment verification

- In-vivo dosimetry

MLC QA There are number of tests currently recommended to ensure that the MLC is performing optimally (IPEM 2008, AAPM 2001, IPEM 2018). The tests should be a part of a quality system ensuring that they are performed at appropriate intervals and subject to locally-derived tolerances (Thwaites et al. 2005, Leer et al. 1998, British Standards Institution 2000). These are designed to inspect the following aspects of the system controlling the leaf positions:

- Independent movement of the individual leaves

- The relationship between leaf pairs

- Movement of the leaf banks

- Relationship between the leaves and collimators (IPEM 2018).

Patient-specific pre-treatment verification Pre-treatment verification is recommended by a number of bodies (IPEM 2008, AAPM 2018). Generally the process involves a recalculation of the treatment plan on a phantom, delivery to that phantom then a comparison of the calculated and measured dose. Point doses can be measured by a chamber but increasingly, the preferred method is to use a detector array to measure and compare the dose in a plane or volume. The results from these arrays are normally subjected to gamma index analysis (See Appendix A); this approach has been shown to be capable of detecting small offsets in leaf positions (Fredh et al. 2013). Increasingly, the measured signal is processed and projected on the patient's CT scan allowing a clinical evaluation of the dose delivered by the linac. This is an improvement on an assessment based metrics derived from the more-arbitrary gamma index analysis of the distribution as it can be difficult to interpret what these mean with regard to the dose received by the patient (Schreiner et al. 2013). Although it is not a widely adopted stance, it is the opinion of the author that verification of a representative collection of treatment plans with a array-style device represents a complete test of the linac beam and and its ability to deliver a clinical dose distribution. Although sensitive to errors it lacks specificity, but appropriate implementation could reduce the

---

[1]1mm at the isocentre corresponds to about 0.3mm at the leaf edge

monthly QA burden by allowing the omission of beam tests that have little direct clinical relevance.

1.4.3 In-vivo dosimetry

A good QA program ensures that the MLC is working well. Pre-treatment verification ensures that a plan is deliverable and that the TPS-generated dose matches the dose delivered by the linac. However, these tests only imply that the dose will be delivered correctly. With the sub-millimetre demand on the accuracy of MLC necessary to maintain dose certainty to the level where clinical outcomes are not compromised, measurement of the plan as it is delivered to the patient seems prudent. Implementation of intra-fraction measurement, or In-Vivo Dosimetry (IVD) is not simple and comes at considerable cost. The next chapter looks more thoroughly at the case for IVD and reviews devices and methodologies for implementing it. This leads to the objective of this work on the characterisation and use of the DAVID transmission dosimeter.

Chapter 2

Literature review

2.1 What is in-vivo dosimetry

In vivo is the Latin for within the living. Whilst some devices and techniques do allow dose to be measured inside a patient, in radiotherapy the term generally refers to the measurement of dose received by the patient during treatment. This is in contrast to *ex vivo* or *in vitro* measurements that are made before or after the treatment with a phantom as a surrogate for the patient (Mijnheer et al. 2013). Whilst measuring the dose received by the patients seems a prudent step it is quite difficult to implement. IVD measurements must either use detectors that do not attenuate the treatment beam too much, or rely on the measurement of dose that is exiting the patient.

2.2 The case for clinical IVD

2.2.1 Radiotherapy errors

Sadly there have been a number of radiotherapy errors; some of the errors discussed in the following section appear in the Chief Medical Officer's report that concluded IVD should become routine in the UK (Donaldson 2007, Department of Health 2006)

North Staffordshire 1994

Radiotherapy treatments can be delivered isocentrically or with fixed Source to Surface Distance (SSD). In the 1980s fixed SSD deliveries were prevalent mode of delivery. At the North Staffordshire Royal Hospital (NSRH), when a patient was treated with a fixed-SSD treatment that was not a metre, a MU scaling factor was applied. In 1982 the centre acquired a new TPS that automatically applied this correction for non-isocentric fields, unfortunately this was not known to the staff. Between 1982 and 1991 all isocentric treatments delivered at NSRH had an unnecessary scaling factor applied to them resulting in the under dosing of over 1000 patients by between 5 and 35%. In cases where the sole modality of treatment was radiotherapy, and the reduction of the dose was 20%, typically bladder and prostate cases, the chance of these patients being disease free in 5 years was reduced by 50% (Ash & Bates 1994, Scottish Executive 2006).

Leeds 2004

In 2004 Leeds Teaching Hospital (LTH) used Theraplan Plus as a TPS. Theraplan lacked networking capabilities and the treatment field data was sent to the treatment set on paper where it was manually entered into the treatment database. As with all deliveries, multiple checks were performed on the database entry, but none of these detected the initial data entry error. Unfortunately the error was a failure to enter the planned use of a wedge. Elekta units use a physical wedge. The wedge attenuates the field, so more monitor units are needed to deliver an equivalent dose than would be delivered by an open field. Typically this is three times as much. In this instance, the patient was subject to the MUs for a wedged field that did not have a wedge in place, receiving 94 Gy compared to the prescribed 40 Gy (Toft 2004).

Glasgow 2006

In 2006 Lisa Norris, a patient at the Beatson Oncology Centre (BOC), received a dose 58% higher than planned (Scottish Executive 2006). The patient had been diagnosed with pineoblastoma and prescribed both chemotherapy and radiotherapy. The radiotherapy prescription was 35 Gy in 20# (1.75Gy per #) to the whole Central Nervous System (CNS) followed by 19.8 Gy in 11# (1.8Gy per #) to the tumour bed.

The BOC had just upgraded their Treatment Planning System (TPS) so that it was more integrated with the Record and Verify (R&V) system. Prior to the upgrade the paper treatment form stated the Monitor Units (MUs) for the treatment in terms of **Normalised Monitor Units**. If the TPS calculated that, to deliver 2Gy (200cGy), 50MUs were required, then the treatment form would have had on it:

$$\frac{MU}{Fraction\,Dose} = Normalised\,MU \tag{2.1}$$

$$\frac{50MU}{2Gy} = 25\,MU\,per\,100\,centiGray \tag{2.2}$$

In order to get the daily treatment MU, the treatment radiographers would need to multiply the normalised monitor units by the fractional dose (in centiGray) and divide by 100.

$$Daily\,MU = \frac{MU\,per\,100\,centiGray \times Fractional\,Dose(cGy)}{100} \tag{2.3}$$

$$Daily\,MU = \frac{25 \times 200}{100} = 50MU \tag{2.4}$$

After the upgrade this round-about process was streamlined. Paper forms were no longer used and **Normalised MU was replaced with MU**. However, on account of the complexity associated with entire CNS deliveries, the new system was not adopted and the old, paper-based system with normalised monitor units was kept in place. In the case where the error occurred, **the planner wrote the MU on the form, not the normalised MU**. When the treatment radiographer

applied the scaling factor, a larger MU was calculated and used for the treatment.

The patient was given almost 60% more dose for the treatment fields irradiating the head. When another CNS patient was planned by the same planner and the same mistake made, it was found out before treatment started; the internal investigation it prompted uncovered the first error.

Panama 2000

In a centre in Panama, the TPS calculated the MUs for a treatment as if the beam were attenuated by a complete lead sheet, not, as intended, a lead sheet with a cut out. Consequently the MU calculated by the system was much higher, as was the dose delivered. At the time the report was published in 2001, 5 people had died as a direct consequence of this error (Vatnitsky et al. 2001). This is discussed in greater detail in Appendix B.1.

Epinal 2004

At the Epinal Hospital in May 2004 the decision was made to replace the use of physical wedges with dynamic wedges for prostate treatments. A failure to understand the subtleties of the TPS user interface led some members of staff to plan with physical wedges. The number of MU being calculated and transferred to the treatment unit was, in these cases, much higher than intended for treatment with dynamic wedges. Lack of proper checking software and procedures allowed this practice to continue until new checking software, that lead to the discovery of the problem, was installed in August 2005. During this time 23 patients had been mistreated (Ash 2007, Wack et al. 2007).

2.2.2 New York 2005

At a New York Hospital in 2005 a patient was treated for H&N cancer; the first four fractions were delivered as prescribed. On reviewing the plan the doctor wanted to increase the shielding to the teeth, so the treatment was re-planned for the fifth fraction. The TPS generated new optimal fluences and was in the process of calculating the MLC positions and movement. However the computer crashed and the MLC information was not saved, but, critically the fluence information was. The plan was reopened on another computer, with the fluence information, but not the MLC data. Dose can be and was calculated with just the fluence information, the plan was reviewed and sent the the linac for delivery. The plan was delivered with the jaws in a fixed position and the MLCs retracted - their default position in the absence of MLC data in the treatment plan. 39Gy in 3 fractions was delivered to the patient's neck, from the base of skull to the larynx with life-threatening consequences (Bogdanich 2010).

Central to this treatment error was a computer failure that led to MLC data transferred to the linac which did not correspond to MLC positions that would generate the dose in the reviewed plan. Due to time pressures, the QA checks were not followed that were in place to catch errors of this nature.

2.2.3 Arguments for and against IVD

Radiotherapy errors can be avoided by

- Thorough and appropriate commissioning of radiotherapy equipment.

- A well thought out and implemented physics QA program.

- A high standard of patient care and treatment by radiographer staff.

- Appropriate use of up-to-date planning, checking and record and verify software by properly trained staff.

- Comprehensive quality system.

It is not the intention to replace any of the above measures with IVD; personal patient monitoring should be used to catch the very small fraction of cases that manage to get through the system; on these points most agree (Feldman et al. 2001). Discussion on the subject tends to focus on whether implementing an IVD system is the best use of typically limited resources.

Harrison & Morgan (2007) and Mackay & Williams (2009) question the cost and effectiveness of implementing a IVD system. While neither suggest that IVD is a bad idea they question its use for all patients and suggest that both the time and money could, in a resource-limited environment, be better spent on the development of advanced techniques to produce better patient outcomes. The author suggests that, given the very low incidence of major radiotherapy errors, it may be possible that development of advanced techniques that result in better treatment outcomes, reduced toxicity and fewer cases of recurrence in large numbers of patients could have a more positive impact on the population's health than a nation-wide implementation of of IVD that would only benefit a vanishingly-small fraction. However, this proposition is almost impossible to prove.

Harrison & Morgan (2007) also point out that a lot of the radiotherapy incidents that are used to justify IVD were brought about through transcription errors and are, in modern radiotherapy centres, almost entirely eliminated through the widespread uptake of electronic transfer of data with software to perform data-transfer checks.

Harrison & Morgan (2007), Mackay & Williams (2009), Klein et al. (2005), Munro (2007) question the effectiveness of IVD in its ability to detect clinically-relevant treatment errors, however, their discussion is largely focussed on point measurements made with a diode. The limitations of this approach are discussed in Section 2.3.1 and newer technologies, discussed later in this chapter, have shown to be better tools for IVD.

Mijnheer et al. (2013) recognise that a comprehensive QA program and additional patient-specific verification on IMRT deliveries should negate the necessity for IVD and that there is, at present, limited hard evidence to show value in IVD. However, with the increased complexity of modern IMRT treatment techniques that have higher demand on the linac's engineering, they recommend IVD be performed. Mijnheer et al. (2013), Munro (2007), MacDougall et al. (2017) advocate the implementation of IVD. Mijnheer et al. (2013) state that ideally IVD should be performed on all radiotherapy fractions. Recognising this is not possible in most radiotherapy centres, they recommend that IVD be performed, in some capacity, on all patients receiving radiotherapy with radical intent and if this is not possible, then, at the very least, the recommendations issued by IAEA (IAEA 2013) should be followed. These state that IVD should be performed when:

- Introducing new treatment protocols or irradiation techniques.

- Changing treatment equipment.

- Implementing new software, such as upgrades of planning systems, machine controlling software, network communication, or patient management systems.

- Carrying out total body or total skin irradiations.

- Delivering single fraction treatments.

- Performing treatments with curative intent where the dose is close to normal tissue tolerance.

Other international bodies to recommend the use of IVD include the World Health Organisation (WHO 2008) and the International Commission on Radiological Protection (ICRP 2001).

Although economic advantages of delivering IVD have been discussed (Williams & McKenzie 2008, IPEM 2006, Mackay & Williams 2009), for the author the key motivations for performing IVD are:

- Assurance that advanced techniques are being delivered with the accuracy necessary to be confident that planned dose distributions match delivered doses so that clinical outcomes are not compromised.

- In the wake of a number of incidents, reassurance that patients being treated with radiotherapy are safe.

2.2.4 Conclusion: IVD in the UK

The North Staffordshire, Leeds and Glasgow errors in Section 2.2.1 were discussed in the annual report by the chief medical officer released 2006. Consideration of these events and the rapidly advancing field of radiotherapy led the chief medical officer to include in the report's recommendations that: "IVD radiation checks should be made routine" (Department of Health 2006).

This was followed by Towards Safer Radiotherapy (BIR 2007), released in 2007 this document contained official recommendations by radiotherapy professionals based on the report by the chief medical officer. Among other things, it outlined the need for departments to have a quality system and what to have in them; it categorised errors and defined vernacular and systems for reporting them and recommended the use of IVD.

For centres that did not have a quality system, implementing one takes time but does not require any new equipment. Likewise for error-reporting protocol. On the other hand, implementing IVD takes considerable resources in terms of physics setting it up and radiographer training as well as the capital outlay necessary to buy the equipment and software. These were the main reasons why the reaction to the report was not entirely positive (Section 2.2.3). In response to the reaction to the IVD recommendations in Towards safer radiotherapy a report came out in 2009 (RCR 2008a). This short document stated that implementing IVD should be done taking into account the risk, potential benefit and cost associated with its deployment. Risk, in this case, could be the diversion of resources away from other

important duties. Echoing the statements made by IAEA (2013) and Mijnheer et al. (2013) it also highlighted that priority should be given to new or more-unusual treatments. Releasing this document took the pressure off radiotherapy centres allowing radiotherapy professionals to consider options and not rush into costly decisions to comply with Towards Safer Radiotherapy.

In 2014, supported by Public Health England, the Radiotherapy Board comprising of representatives from the Institute of Physics and Engineering in Medicine, Royal College of Radiologists and Society and College of Radiographers, requested a review and update of the response to Towards Safer Radiotherapy (BIR 2007). In vivo dosimetry in UK external beam radiotherapy: current and future usage was published in 2017 (MacDougall et al. 2017). In order to determine the current IVD practices and make appropriate recommendations, it conducted a nation-wide survey in the summer of 2014.

In summary, 73% of providers said that they were routinely using IVD, when compared to 37% of providers performing routine IVD in 2008 (Edwards et al. 2007). The report demonstrated that although Department of Health (2006) had provoked some debate, the response was, on the whole a more widespread uptake of routine IVD. Of the 27% not routinely performing IVD, only a small fraction said that they had no intention of doing so. Of the centres routinely performing IVD, 85% were using diodes with 15% using EPIDS. It was noted that the large percentage associated with diode usage reflects historical practice with the report commenting that diodes were not appropriate for IVD of highly modulated fields or arc deliveries.

The first recommendation made by MacDougall et al. (2017) was that: "RT providers should implement local protocols for verifying therapeutic radiation dose is delivered as prescribed." Some of the other recommendations suggested that emphasis should be placed on the development of the still-emerging EPID-based approach. However, MacDougall et al. (2017), like the other UK recommendations (Department of Health 2006, BIR 2007, RCR 2008a) and the international recommendations (IAEA 2013, ICRP 2001, WHO 2008), did not state what approach should be taken by radiotherapy providers to implement IVD. The rest of this review is concerned with looking at the devices and techniques for performing IVD in the radiotherapy clinic.

2.3   Devices for IVD

2.3.1   Point detectors

Diodes

Diodes are small semi-conductor devices that typically lie on the skin of the patient during treatment and give a real time reading of dose. In addition to these favourable properties, they are mechanically robust and have no need for an external voltage. These properties are probably why, until recently, diodes were the most commonly used in vivo dosimeter in the UK (MacDougall et al. 2017) and US (AAPM 2005). Unfortunately, the diode's response is a function of field size, dose rate, temperature and accumulated dose; as a consequence they need regular calibration. Higgins et al. (2003) perform monthly calibration; Mijnheer (2008)

and Huyskens et al. (2001) report an accuracy of 1.5% and 1.6% (1sd) respectively if the devices are well calibrated. Even with good calibration, the response of diodes depends on the angle of the beam; the response can change by as little as 2% for angles 45° but can be as high as 8% for angles above 60° (depending on the manufacturer (Higgins et al. 2003)) This makes them less than ideal for use as an in-vivo dosimeter.

Also, despite not requiring an external voltage, most diodes must be attached to a cable during acquisition meaning that it is impractical to place the device at the treatment site. Some wireless systems do exist but these have yet to have widespread uptake.

Despite these shortcomings (Huyskens et al. 2001) implemented an in-vivo program that was able to detect both major (>10%) and minor errors (<10%) in therapy delivery. They concluded that a "5% action level was appropriate" though in some cases (tangential or wedged beams) a higher 6-7% action level might be more practical.

TLDs

Thermoluminescent Devices (TLDs) are small ($\sim$ 1mm) detectors; the response has little dependence on temperature, dose rate or therapeutic energy. They do not require an external voltage and have the major advantage over diodes in that they do not require cabling during acquisition, making them useful for intra-cavity applications. Though it should be noted that the heating process required to extract the signal is time-consuming and work intensive (Izewska & Andreo 2000). The fact that TLDS do not require a wire is probably the main advantage over diodes, but this facet also highlights their biggest shortcoming: the inability to produce real-time readings (Mijnheer 2008).

2.3.2   Discussion

At the time when Towards Safer Radiotherapy (BIR 2007) was released, diodes and TLDs were the most common technology appropriate for the role of in-vivo dosimeter. International reports released around the same time (IAEA 2013) recommended them for this purpose. In the case of diodes, this led them to be the most widely-used device for performing IVD (MacDougall et al. 2017, AAPM 2005). However the point-like nature of these devices has rendered them increasingly obsolete as radiotherapy delivery techniques have advanced. While it is possible to place TLDs internally (Azorín 2004) and perform in-vivo dosimetry in the truest sense of the phrase, this resource-heavy invasive procedure comes with its own risks and is rarely performed in practice. Typically both TLDs and diodes are placed on the skin surface. The measured dose is then compared to the dose calculated at the same point in the TPS. There are a number of problems with this approach. Firstly the detectors are sensitive to back scatter and must be placed flush with the patients skin; in the case of TLDs bolus material is normally necessary. Also with wedged and IMRT fields, that can create high dose gradients across the field, the position of the detector is critical. With scope to get the set up wrong, the tolerances on point measurements tend to be quite high, often 5% and sometimes, as high as 10%. Assuming the diode is perfectly placed, a tolerance of

10% equates to a difference in SSD of 5cm. So it might be stated that as the diode is on the patient, it's sensitive to patient set up (Edwards & Mountford 2009), but it is the opinion of the author that a detector that is not sensitive to a difference in SSD of 5cm does not adequately verify the patient set up.

The most significant limitation of point detectors in the context of modern radiotherapy was brought about through the advent, and widespread uptake, of arc therapy. Typical use of diodes requires them to be placed on the surface where the beam enters. In the case of a seven-field step-and-shoot IMRT delivery, treatment staff will have to enter the room several times to re-site the diode. For a VMAT deliveries, it is not possible to get a meaningful result. It is now recommended that "diode-based point dose measurements are not appropriate for verification of dose delivered when treating with highly modulated, rotational or adaptive techniques" (MacDougall et al. 2017).

2.3.3 Electronic Portal Imaging Devices

Electronic Portal Imaging Devices (EPIDs) are MV imaging panels initially devised for verification of patient set up. The amorphous-silicon-flat-panel EPID now comes as standard with every linac install (MacDougall et al. 2017). These devices are comprised of a method to convert radiation dose deposition into a light signal, a light detector and a readout system (Van Elmpt et al. 2008). The X-ray converter is constructed from amorphous silicon which also gives the device its name. The silicon is, at an atomic level, inherently disorganised meaning that the panel, which images in the main beam, is less susceptible to radiation damage. The fast sampling time, high resolution, digital output and linear dose response that is independent of dose rate proved to be an effective MV-energy imaging device (Van Elmpt et al. 2008). A combination of the a-Si EPID's attributes and its increasing ubiquity make it a strong candidate for performing in-vivo dosimetry. EPIDs as IVD devices have a number of advantages:

- They are included as standard on all modern linacs making hardware costs low.

- They are sensitive to patient set up.

- They do not perturb the beam or affect the delivery in any way.

In many ways, EPIDs represent the ideal IVD device. However, they were not originally intended to be used as dosimeters, the fact that they are not in widespread use for IVD reflects the complexity in managing this limitation. Turning the EPID signal in to a patient dose is non-trivial, this is largely down to the tendency for amorphous silicon EPIDs to over respond to low energy scatter. There has been some success in using EPIDs to determine a point dose in a patient (Piermattei et al. 2007, Celi et al. 2016), but it is only recently that the use of the EPID to back project a 2/3D dose in the patient CT has been turned in to a commercial product (Delaby et al. 2017). The NKI solution has been implemented in the clinic and an initial report given of its use on 4337 patients. 17 serious errors were found, 9 of these would not have been discovered without EPID-based IVD (Mans et al. 2010). Wider-spread implementation is under way (Ricketts et al. 2016) and

investing resources in the development of this technology is recommended (Mac-Dougall et al. 2017). Time will tell if this iteration of the software is feasible in the typical clinic.

2.3.4 Log File analysers

Linacs record the position of the MLCs, Jaws, gantry and collimator during treatment and save them in a log file that is available to the user. It has been shown that analysis of these log files can detect delivery errors in Elekta (Tyagi et al. 2012, Fontenot 2014) and Varian (Litzenberg et al. 2002, Rangaraj et al. 2013, Vazquez-Quino et al. 2017) machines. Mobius3D (Mobius Medical Systems, Huston, TX, USA) is a commercial system that uses the beam shaping positions in the log file and plan file to calculate the dose on the patient CT using a collapsed cone algorithm (Childress et al. 2012). The dose calculated using the plan file acts as a secondary MU check (see Appendix B.1), the dose calculated using the log file is described as the delivered dose and serves as a check of the dose actually received by the patient (Fontenot 2014, Mobius3D 2019). The system has been commissioned on both Elekta (Nelson et al. 2016) and Varian (McDonald et al. 2017).

There are a number of advantages to log file analysis:

- The process can be automated and performed on each delivery saving considerable time.

- It has been suggested that this sort of analysis is more sensitive to errors than planer or point dose measurement (Rangaraj et al. 2013, Childress et al. 2015)

- The analysis can be very specific. Unlike the gamma assessment, the analysis performed by the Mobius software allows the dosimetric impact the error has made on the plan to be determined. The exact cause is also identified, reducing the time it takes to remedy a machine-based error (Childress et al. 2015).

There are two key shortcomings with log file analysis. The first is that the beam profile and energy are generally not reported; the processing software assumes that the beam matches the beam model data exactly. Although the machine should interlock if the beam energy or profile changes significantly, there is scope inside the tolerance values for the beam to be asymmetric and have slightly different energy values bringing in to question the accuracy of the "delivered dose" calculation performed by the software. Furthermore, any error that results in beam instabilities not causing an interlock will go undetected. This leads in to the second, and, for the author, the key problem with log file analysis as a form of in-vivo monitoring: the lack of independence. It is quite possible that the information in the log file might not reflect a delivery error - if a significant error was detected by the machine, then it should interlock anyway. One purpose of IVD is to detect errors when these interlocks have failed. Although machines are subject to strict protocol that dictates the accurate calibration of beam shaping apparatus (AAPM 2009, Childress et al. 2015) there have been cases where these values have been reported incorrectly and only detected via a QA measurement. (Agnew et al. 2014).

Log file analysis is currently being implemented by both Elekta and Varian users in a number of centres for treatment delivery validation (Tyagi et al. 2012, Fontenot 2014, Litzenberg et al. 2002, Rangaraj et al. 2013, Vazquez-Quino et al. 2017, Nelson et al. 2016, McDonald et al. 2017). The debate regarding its effectiveness continues (Childress et al. 2015).

2.4 Transmission detectors

2.4.1 Introduction to transmission detectors

Transmission detectors are flat-panel detectors that can be attached to the head of the linac - normally using the fixtures that hold the accessory tray or electron applicators. Transmission detectors can accurately monitor the linac output and can, in some cases, interrupt the beam in the event that the output deviates from the baseline by more than pre-determined presents (Islam et al. 2009, Pasler et al. 2017). In this way, unlike IVD devices previously discussed, transmission detectors have the potential to actually prevent a fault or error causing serious harm.

In the case of an EPID measurement, the radiation has passed through the patient allowing it to assess the patient set up. The upstream nature of transmission detectors mean that they have no way of reviewing the patient set-up. Furthermore, by necessarily obscuring the whole field, they will, to varying degrees, attenuate the beam. This can be accommodated for by applying a scaling factor to the treatment MU, however this does introduce an element of risk. By interacting with the beam prior to the patient they also act as a secondary electron source and have been shown to increase the number of contamination electrons at the surface by up to 18% (Asuni et al. 2011).

MacDougall et al. (2017) describe transmission detectors as an emerging technology with potential for use as a daily treatment monitoring tool. However, on account of the limited clinical experience or published works on the subject, the discussion was limited to little more than an acknowledgement that transmission detectors existed and could be used for IVD. What follows is a more in-depth review on specific models of transmission detector.

2.4.2 Slanted-plate transmission detector

Design

Islam et al. (2009) developed a large-area, transmission-style ionisation detector. The detector is made up of three, 22cm square, plate electrodes. Two polarising electrodes sandwich the collecting electrode. The polarising plates and inner collecting plate are 3.18mm and 1.59mm thick respectively; with a PMMA frame the total thickness of the detector is 4cm.

A particularly novel aspect of this design is that polarising plates are placed at an angle to the collecting plate (Figure 2.1). The collection efficiency is proportional to the plate separation so, in having a varying plate separation, the chamber is made spatially sensitive in one dimension. If the chamber were to be irradiated with a small square field towards the end where the plates are closer then the signal would be larger than if the plate were to be irradiated by the same field where the plates are further apart. The gradient of the polarising plates is subject to a

compromise between the sensitivity to spatial change; the plate separation varies linearly from 2 to 20mm corresponding in a signal change of approximately 0.5% per mm. The operation of large-plate like ionisation chambers is often hindered
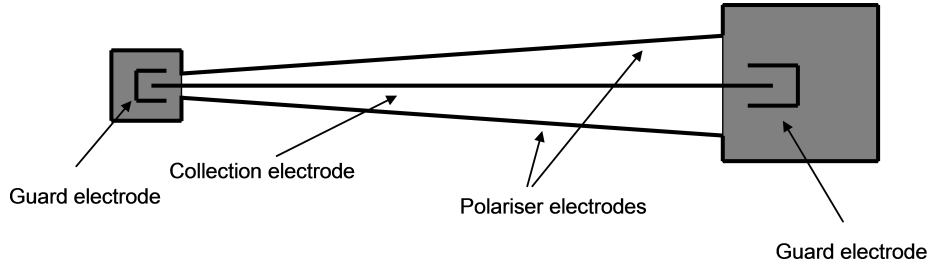
Figure 2.1: Transmission detector relying on varying sensitivity, brought about by varying proximity of the collection electrode and polarising electrode, to give a spatially sensitive signal (Islam et al. 2009)

by the fact that electrometers can saturate in normal usage due to the amount of charge collected. Electrometers that are capable of integrating the large amounts of charge generated normally do so at the expense of the resolution of small charge readings. Islam et al. (2009) have overcome this problem by using two electrometers "operating in a switching configuration"; this allows currents over the range of 0.5nA – 0.5mA to be collected accurately. The details of the system are not important here, but while one electrometer is set to collect the other is set to rest, the roles are changed when the collecting electrometer nears saturation.

Islam et al. (2009) also came up with a way of predicting the signal based on the characteristics of the detector, the MUs to be delivered and the geometry of the MLCs as shown in Equation 2.5.

$$
\begin{aligned}
S_{calc} = MU \times k \times AOF(X,Y) \times &\int_{A1} F(x,y)b(x,y)dxdy \\
&+ \int_{A-A1} T_{MLC}(x,y)F(x,y)b(x,y)dxdy \\
+ \int_{R-A} T_{Jaw}(x,y)&T_{MLC}(x,y)F(x,y)b(x,y)dxdy
\end{aligned}
\tag{2.5}
$$

Where MU is the monitor units, k is a system constant, $AOF(XY)$ is the integrated output factor for the jaw settings (this was found empirically), F dictates the fluence distribution (including penumbra) and b is the detector response. The later two integrals in Equation 2.5 refer to leakage with T corresponding to transmission through either the MLCs or the Jaw. A is the open aperture, $A1$ is the active area covered by the MLCs and R the active area covered by the jaws.

Testing and discussion

Islam et al. (2009) found the electrometer system to be "stable and reproducible" with a linear dose response in the range of 1-2000 MU. The observed dependence on dose rate was a variation of 0.2% over the range 100 – 600MU min^{-1}. An 11 segment prostate treatment was delivered and the average agreement was 0.5% with a maximum deviation of 3.5% per segment. This was repeated with a H&N treatment with an average agreement of 0.9% and maximum standard deviation of

5%. In both cases the largest discrepancies were seen in the smaller-area segments. Most of the results were lower than 1% and it was thought by (Islam et al. 2009) that the errors could be reduced if the smaller segments were combined with a larger adjacent segment. However, the difference signal from the total delivery when compared with predicted responses was found to be <1%.

Further testing was completed by artificially introducing probable errors to square fields. A 15cm^2 field was reduced by 1mm in width simulating a leaf bank error; the result was a 0.7% disagreement with the result from a true 15cm^2 field. This was repeated with a square field of 3cm^2 the disparity was seen to increase to 3%. The movement of a single leaf by 5mm in a 3cm^2 field generated a 2% difference to the signal (Islam et al. 2009). Variations in machine output were matched by equivalent variations in the device and a change in energy from 6 to 10MeV reduced the signal by 6.3 +/- 1.5% (Pasler et al. 2017).

More realistic scenarios were tested by artificially introducing leaf bank errors. 15 clinical treatment plans were used; 5 prostate, 5 H&N and 5 partial-arc breast VMAT plans, planned on Pinnacle (V.14 Philips Radiation Oncology Systems, Milpitas, CA) and delivered on an Elekta Synergy linac (MLCi2). The plans with the errors in were delivered to the device and compared to the signal generated from the delivery of the unmodified plans. 2mm and 1mm unidirectional leaf bank shifts were seen to have a limited (<1%) effect on the signal. This was deemed by the authors to be of little consequence as recalculating the modified plans on in the planning system showed that these shifts made negligible differences to the calculated dose distribution. Leaf bank shifts toward and away from the central axis resulted on average in a +/- 17.9% difference and 5.4% difference from the unmodified signal for 2mm and 1mm respectively with positive increases in dose for movement away from the axis and negative for toward (Pasler et al. 2017).

The effect of the device on the beam characteristics was generally negligible with the most significant difference being a 3% rise in surface dose (Islam et al. 2009). The reproducibility of signal generated from VMAT plans was ∼ 1% (Pasler et al. 2017).

This detector has an interesting design with a novel approach to incorporating spatial sensitivity in the direction of leaf movement. The aluminium composition makes the device opaque to visible light which means that the Optical Distance Indicator (ODI)[1] cannot be used and optical field matching cannot be performed. This could be compensated for by making the finished product easy to slide in and out of position, but this solution is far from ideal in a working environment.

Unfortunately, sensitivity in the direction of leaf travel is gained at the expense of spatial sensitivity in the direction perpendicular to the leaf movement. So, while the device is able to detect single leaf errors, it is not able to pinpoint their origin (Islam et al. 2009). Also, though the leaf-direction sensitivity is one of the key selling points of the design, it is shown that 2mm unilateral moves in the leaf bank do not generate a deviation from the baseline that is above the uncertainty in the measurement of 1% (Pasler et al. 2017). Furthermore, even if the device were sensitive to these errors, DVH analysis of unilateral shifts demonstrated that they have limited clinical impact. Leaf errors and their effects are discussed more in Section 6.1.

---

[1] Beam of light shone through a calibrated graticule so a projected image displaying the SSD is visible on the patient's surface

The method of predicting the signal (Equation 2.5) was applied to 1 prostate and 1 H&N and was seen to agree with the total signal generated by the delivery of the plans to within 1% (Islam et al. 2009). This promising result does not seem to have been followed up in subsequent work.

2.4.3 COMPASS

IBA Dosimetry (Germany,Schwarzenbruck) have created a patient-specific QA system that consists of three parts: The MatrixX detector (Godart et al. 2011, Nithiyanantham et al. 2015), The Dolphin Transmission Detector (DTD) (Thoelking et al. 2016a) and the COMPASS software. The COMPASS software can perform several functions (Thoelking et al. 2016b):

1. Import patient images, structures and DICOM plans from the TPS and calculate the dose in the patient with a collapsed cone algorithm. In this way it acts as independent check of the TPS dose calculation (See Appendix B.1).

2. Import and analyse the measurement from the MatrixX detector array for the purposes of patient-specific pre-treatment verification.

3. Import the measurements from the DTD for plan verification during the patient's treatment.

The DTD is a 2D array consisting of 1513 vented parallel-plate ionisation chambers[2] that cover an active area of 24 cm^2. This projects to a 40 cm^2 at 100cm SSD when measuring at SDD 60 cm. The centre-to-centre distances of the chambers range from 0.5cm to 1cm. The height and diameter of the individual chambers are 0.2cm and 0.32cm respectively (Thoelking et al. 2016a). With the DTD in position there is an increase in surface dose that is field size and SSD dependent. A maximum of 11% higher than without the DTD in place for a 30cm^2 field at 80cm SSD, this was however deemed clinically acceptable. The PDD below D_{MAX} was observed to vary by less than 1% with the DTD in place. Although the MU factor, to correct for the attenuation of the beam by the device, was seen to vary slightly by field size, a generic value was used for all IMRT plans assessed in the study. When subject to 2%/2mm gamma index analysis, no significant differences were seen between the IMRT distributions with and without the DTD in place (Thoelking et al. 2016a) similar results were found through Monte Carlo techniques by Nakaguchi et al. (2017).

The device was tested in more realistic clinical scenarios by delivering 18 clinical step-and-shoot IMRT and VMAT plans for various sites with the DTD in place. The previously-evaluated (Godart et al. 2011, Nithiyanantham et al. 2015) MatrixX was also used to asses the delivery at the same time. Both recorded equivalent gamma scores for all of the plans.

It would appear that the strongest attribute of the IBA system is the facility to use the DTD signal to reconstruct the delivered dose inside the patient's CT data set. This allows a three-way comparison between the delivered dose, the dose

---

[2]IBA seem to have abandoned a similar design that used diodes (Stasi et al. 2010), this is because of the effect that it was having on the treatment beam (Asuni et al. 2011, Venkataraman et al. 2009, Thoelking et al. 2016a).

calculated by the Collapsed Cone (CC) algorithm (Ahnesjö 1989) in COMPASS and the planned dose distribution from the TPS (Thoelking et al. 2016b). This feature was tested by making such a comparison between the aforementioned 18 plans. The subsequent distributions were then subject to 3D gamma index analysis as outlined by Stock et al. (2005) and Visser et al. (2013) with evaluation criteria of $\gamma < 0.4 =$ pass, $0.4 < \gamma < 0.5 =$ evaluate, $\gamma > 0.5 =$ fail. When comparing the CC dose, the dose generated from the DTD and the TPS dose, 16 of the 18 plans passed with 2 falling in to the evaluate criteria. As the dose has been reconstructed inside the patient data set, in theory, a clinician could evaluate the delivered dose. However the study chose to quantify the clinical assessment by comparing DVH parameters on target and OAR structures. The average dose to the target and 95% coverage dose calculated from the DTD signal were compared to the CC-calculated dose; averaged over all plans the differences were found to be found to be -0.8 +/- 0.5% and -1.3 +/- 0.6% respectively. Dose to OARs was overestimated by 2%, but most of the specific indices were < 2%. These results demonstrate that the dose reconstruction inside the patient based on the delivered signal captured by the DTD is good enough to be used clinically (Thoelking et al. 2016b), the same conclusion was reached by Monseux et al. (2016)

In a method similar to that used by Pasler et al. (2017) the ability of the DTD to detect errors was assessed by intentionally introducing errors in to RTP files, delivering them to the detector and comparing the reconstructed dose by means of 3D gamma index analysis and DVH-evaluation. 3%/3mm gamma values increased from 0.29 +/- 0.03 for the unmodified plans to 1.03+/-0.32 for the plans with a 1mm shift of the leaf banks toward the central axis. This implies that these errors, according to the passing criteria already discussed, would be detected by the system. Like Pasler et al. (2017) unidirectional shifts were more difficult to detect and were not evident in the gamma assessment. However, the authors claimed that the errors were evident when the dose distribution was compared to the unmodified dose as there was a significant difference around the edge of the treatment volumes. Output errors of -2 and +3% resulted in gamma values of 0.79 and 0.61 respectively demonstrating that they would both be detected.

Interestingly, although DVH analysis highlighted the output errors by showing an equivalent increase / decrease in dose, the leaf bank errors were not evident (Thoelking et al. 2016b). This is unexpected as errors of this magnitude have been shown to have clinical impact (Zhen et al. 2013, Rangel & Dunscombe 2009, Oliver et al. 2010). It would have been useful to see the modified plans calculated in the TPS and compared to the unmodified plans also calculated in the TPS. This way difference in dose calculation could have been ruled out and the difference, or lack of difference, between the plans could have been attributed to the characteristics (eg modulation) of the plans.

The Dolphin-Compass combination provides a complete patient-specific system for IMRT QA. The equivalence between the results from the MatrixX and the DTD show that pre-treatment QA with just the Dolphin array is possible. It is the intra-treatment monitoring where the solution excels. While 3D gamma index analysis is possible, it is often difficult to relate gamma values to clinical consequence and in some cases, errors with clinical effects can pass gamma criteria (Appendix A). The ability to reconstruct the delivered dose in the patient CT provides the opportunity

for well informed clinical decisions to be made in the event of a delivery error.

2.4.4 The DAVID

Construction and design

The Device for Advanced Verification of IMRT Deliveries (DAVID) is a transmission detector consisting of two Perspex blocks that sandwich a vented air gap; inside the air gap are thin wires that line up with MLC pairs (Figure 2.2). Incoming radiation scatters inside the Perspex blocks, the electrons move into the air gap and are influenced by the potential between the aluminium (evaporated on to the surface of the Perspex and held at a +ve voltage) and the collection wires held at ground (Figure 2.3). Consequently, a charge builds up on the collection wires that is proportional to the radiation incident on the detector. By making the device out of clear Perspex and ensuring that the aluminium layer is very thin, the device is transparent; allowing the ODI to be used with the DAVID in place (Poppe et al. 2006).

Figure 2.2: The DAVID

The DAVID device is specific to the linac head model as each detection wire needs to line up with a MLC pair: more MLCs – more wires. As a result, the cross-field resolution of the DAVID depends on the MLC head in use (Figure 2.4). Poppe et al. (2010) point out that in moving from a DAVID with 37 wires to one with 80 as a consequence of a new head; the resolution went from 10mm to 5mm.

Error detection capability and clinical implementation of the DAVID

Poppe et al. (2010) performed pre-clinical testing on the DAVID for its error detecting capabilities by purposefully introducing errors to a 25 fraction treatment with a fraction taking place daily. These artificial errors included:

- The mis-positioning of three leaves of the MLC – leaves 14, 17 and 20 were displaced by 2mm in the 15[th] fraction.
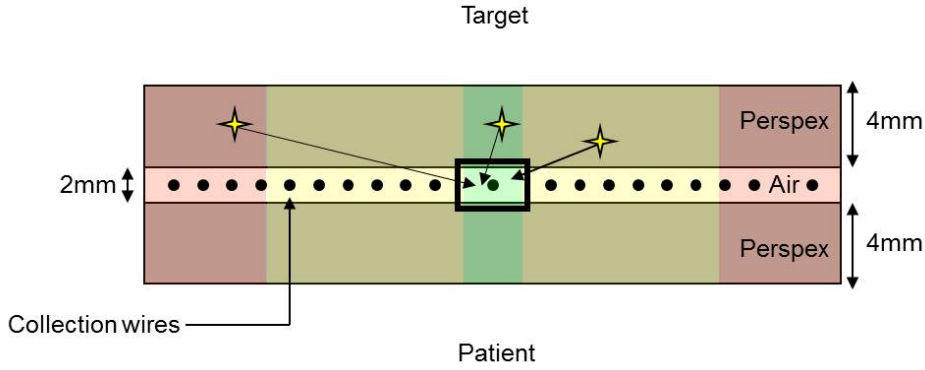
Figure 2.3: Primary and Compton-scatter events contributing to the signal measured at the central wire.
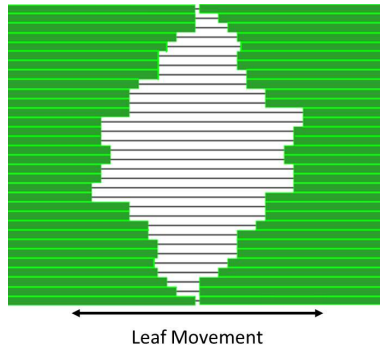
Figure 2.4: Beams-eye view of an IMRT segment showing the MLCs (green) and the collection wires - each lining up with a leaf pair, in the direction of leaf travel.

- The omission of a single segment – during the 19th fraction a segment was omitted

- The sudden change of the photon fluence simulated by removing the cross hair reticule

- The change of photon energy – the 25th fraction was irradiated with a 10MV beam instead of a 15MV beam.

The record and display system employed by Poppe et al. (2006) was designed so that the integrated signal for the whole fraction was displayed on the screen; the individual columns – indicative of the signal received by each wire – were subdivided by horizontal lines indicative of a segment. For each of the artificially-introduced errors, the DAVID system detected the fault. Poppe et al. (2006) also showed that the "relative deviations remained below 1% of the mean signal except for very low signals" indicating that the normal fluctuations of the DAVID are below the signal changes caused by errors. In another test of the devices' stability, the fluctuations of all the wires in an 8 fraction H&N treatment were never found to deviate outside 3% alert threshold.

The DAVID was used over an 8 month period during which 35 IMRT plans were subject to daily verification the results reported in Poppe et al. (2010). In this time the device performed well exhibiting good levels of stability and detecting three clinically relevant errors. The first one demonstrated the DAVID's sensitivity,

detecting a 1% change in the linac output. The dose per MU was changed due to the findings of a monthly QC session; the treatment continued but a new reference measurement was taken with the DAVID to account for the change. The DAVID consistently detected a fault at the field border, when this was investigated it was found that a collimator block was incorrectly calibrated. The final case involved a plan modified midway through a treatment. When the new field was imported the MLC information was lost and the linac irradiated the patient with a square field that was instantly picked up by the DAVID.

When Poppe et al. (2010) was published the DAVID system had verified over 50,000 segments.

Origin of DAVID signal

The signal measured at any wire is made up of three components (Poppe et al. 2006) (Figure 2.5):

- The fluence of radiation above that particular wire.

- The scatter radiation formed in the Perspex by radiation incident on other parts of the surface.

- The leaf and inter-leaf leakage.

Poppe et al. (2006) quantifies these signals empirically. The main signal is characterised by opening the leaf pair associated with the test wire. The scatter signal is analysed by keeping the test-wire leaf pair closed and varying the square field size with all the other MLCs and jaws. The leakage radiation was found by collecting a signal with the MLCs closed – this was deducted from the other results. The analysis consisted of plotting the signal as a function of leaf separation and then as a function of square field length - both exhibited a near-linear fit. The non-linearity manifested itself, in both cases, as a slight curvature of the response - this was attributed to the effect of the flattening filter.

The current paradigm for using the DAVID to perform in-vivo dosimetry involves planning the treatment and verifying the treatment plan with the DAVID in place; if the plan passes the verification then the DAVID signal recorded during the verification is used as a baseline for the treatment. Subsequent treatments are then performed with the DAVID in place and the resulting signal compared with the baseline acquired during verification.

Looe et al. (2010) went a step further by demonstrating that the signal can be described mathematically as the convolution of the photon fluence with the lateral response function (Equation 2.6).

$$S(x) = \sum_{\zeta=1}^{D} P(\zeta) f_\zeta(x) \tag{2.6}$$

Where $P(\zeta)$ is the photon fluence, ζ is the number pertaining to the leaf pair, $f_\zeta(x)$ is the lateral response function (LRF), x is the wire associated with the response and D is the total amount of wires.

The LRF describes the fractional contribution to the signal at any wire (position denoted by x) as a result of the fluence through any leaf pair (position denoted
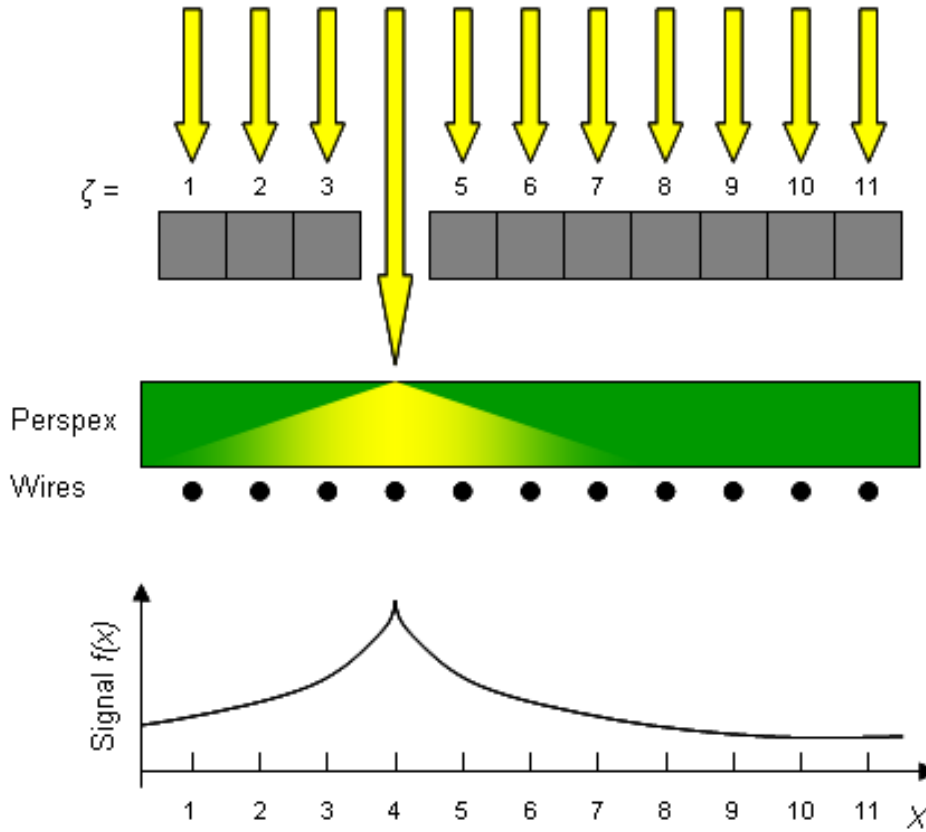
Figure 2.5: Lateral response function. Signal as a consequence of one open leaf pair is recorded at a number of wires. The signal reduces as the distance from the open leaf pair increases.

by ζ) and is normalised as shown in Equation 2.7

$$\sum_{x=1}^{D} f_\zeta(x) = 1 \tag{2.7}$$

The LRF can be characterised by closing all the MLC leaves apart from one, irradiating the DAVID and recording the response from all the wires. Although only one wire will be irradiated by primary radiation the wires next to the wire associated with the open leaf pair will measure considerable scatter radiation (Figure 2.5) (Looe et al. 2010). Poppe et al. (2006) observed that the sum of the signals from the adjacent wires was approximately equal to the response from the wire under the open leaf.

The DAVID Conclusion

Out of the devices discussed in this section, the DAVID is the only one that is transparent, meaning that it can be in place whilst the ODI is used. Unlike the COMPASS solution (Section 2.4.3) the DAVID lacks the ability to translate the signal into patient dose. However, it has been shown to be sensitive to small changes in leaf and jaw position (Poppe et al. 2006) and has even been used in practice to detect such errors (Poppe et al. 2010).

The detector array discussed in Section 2.4.2 has been designed to be spatially

sensitive to errors in the leaves in the direction of leaf travel, in contrast, the DAVID has been designed to be spatially sensitive to errors in the direction perpendicular to leaf travel. Unilateral leaf-bank errors will not be detected by the DAVID, however unilateral errors of 2mm were not detected by the slanted transmission detector and it has been shown (Section 6.1) these have limited clinical impact. The DAVID, like the slanted array and COMPASS system, is sensitive to leaf bank shifts towards and away from the isocentre (Poppe et al. 2010, 2006). Although all the detectors are sensitive to individual leaf errors the DAVID can detect very small miscalibrations (∼1mm) and isolate the associated leaf pair.

The DAVID, though lacking the complex analysis of the COMPASS detector is sensitive to clinically relevant leaf and jaw errors. The device is easy to use and produces a real-time output allowing the treatment to be interrupted in the case of serious errors. Poppe et al. (2010) have shown the device to be appropriate and useful for routine clinical monitoring of radiotherapy treatments.

Discussion on Transmission detectors

Is it in-vivo? Transmission detectors, by virtue of their position, have the potential to accurately describe the characteristics of the radiation beam. Their major shortcoming, at least with regard to in-vivo dosimetry, is that the response will contain no information about the patient (Edwards & Mountford 2009). Post-patient transmission dosimeters and ones that lie on the patient's skin will have a response that is dependent on the position of the patient while, with upstream transmission detectors, it would be possible to deliver the treatment and verify it even if there was not a patient on the couch.

Electron contamination and the effect on the beam Any transmission detector will, to some extent, attenuate the beam. There are two methods of compensating for the attenuation: increase the monitor units delivered or recollect commissioning data for the TPS with the device in place. The first method is a similar method to that used to compensate for the Perspex tray that was used to hold lead blocks in early conformal treatments. Increasing the monitor units has the advantage of being simple to implement but lacks the sophistication to deal with potential beam hardening effects, resulting in the behaviour of the radiation inside the patient not being correctly modelled in the TPS. Measuring the TPS commissioning data with the device in places offers a complete solution to the beam hardening problem, however any treatment planned with the software using the data collected from the detector must be delivered with the detector in place. This could have practical issues in a department where not every machine is using a detector all the time, meaning two sets of treatment planning data would need to exist for each linac. This over complication increases the likelihood of an error, treatments being delivered with/without the detector in place when the detector should not/ should be in place. The DAVID system addresses this by having an accessory coding system so the treatment beam will not deliver if the DAVID is/is not in place and it should not/ should be (Poppe et al. 2010).

Electron contamination is present in any external beam radiotherapy treatment. Electrons will be produced by the interaction of the radiation with the air, ionisation chamber, head shielding etc. Contaminant electrons are undesirable as, if they hit the patient, they increase the surface dose which can cause or contribute

to conditions like erythema, fibrosis, necrosis and epilation (Asuni et al. 2011).

In the case of the beam being blocked by a material with thickness greater than the range of secondary electrons, the majority of the secondary electrons created in the head apparatus are stopped and the material itself becomes the primary source of electron contamination (Asuni et al. 2011, Venkataraman et al. 2009). So, transmission detectors, by their very nature, will change the characteristics of the electron contamination when compared to an open field.

Asuni et al. (2011) examined the effect the COMPASS detector (Section 2.4.3) has on the electron contamination. They state that specific information such as energy spectra, angular distribution and fluence distribution of the contaminant electrons are required to understand the effects fully. Acknowledging that this information is difficult to determine experimentally Asuni et al. (2011) quantify it with Monte Carlo modelling using the EGSnrc distribution (Kawrakow et al. 2000). It was shown that modifications needed to be made to the electron contamination models of commercial treatment planning software to accurately model the beam with the COMPASS detector in place. In the case of the Eclipse anisotropic analytic algorithm (AAA) (Varian Medical Systems, Palo Alto, CA) this modification was a 12% and 15% increase in the Gaussian widths used in the contaminant electron source model. In the case of the Collapsed Cone algorithm used by Pinnacle, the off-axis coefficients needed to be reduced to accurately model the device. It was concluded that modifying the treatment planning dose-calculation algorithm was "a necessary step for accurate beam modelling when using the device". Though its analysis is limited to surface effects, no attention is given to clinical target doses or OARs other than the skin. Without this information, it would appear that drawing such strong conclusions is a little premature.

The findings in Asuni et al. (2011) match those of Venkataraman et al. (2009) and Stasi et al. (2010) who found experimentally that the presence of the COMPASS detector affected surface dose and the dose in the build up region; the effect beyond dmax was found to be negligible. Venkataraman et al. (2009) showed that the attenuation was about 3.3% (they conclude a transmission factor of 96.7%) and Stasi et al. (2010) show that the beam is attenuated by between 3.5 and 4%. Venkataraman et al. (2009) and Stasi et al. (2010) differ from Asuni et al. (2011) in that they suggest that the problem can be dealt with by introducing a transmission factor for deliveries with the COMPASS attached as opposed to modifying TPS variables.

The DAVID is 10mm thick (Figure 2.3); the section that the beam passes through is almost entirely Perspex - the metal wires used for signal collection are thin enough to have a negligible effect on the dose once scatter and beam divergence have been taken in to account. Despite the thin Perspex some attenuation will occur; transmission factors of 0.953 +/- 0.001 and 0.968 +/- 0.001 have been measured at energies of 6 MV and 10 MV respectively Poppe et al. (2006). A surface dose increase when the DAVID was attached was also described by Poppe et al. (2010) who found that it was proportional to field size. Interestingly, Poppe et al. (2010) also showed a surface dose decrease for very small field sizes and noted no significant difference in transmission factor observed for different SSDs.

Poppe et al. (2010) states that the MU delivery was altered "by a correction factor applied at the end of the treatment planning procedure" to compensate for

the DAVID detector but admits that "another option would be the commissioning of the planning system in the presence of the DAVID system in the beam".

For the purposes of comparison, Poppe et al. (2010) includes a table of the normalised dose enhancements to tissues at two depths as a result of transmission detectors (Table 2.1)

	Energy	SSD 70		SSD 150	
		D = 1.5cm	D = 3.5cm	D = 1.5cm	D = 3.5cm
DAVID	10	1.13	1.08	1.05	1.03
DAVID	6	1.06	1.04	1.02	1.01
COMPASS	6	1.37	1.16	1.1	1.04

Table 2.1: Dose enhancement as a result of two different transmission detectors at two depths (D) copied from (Poppe et al. 2010)

2.5 Conclusion and discussion of intended work

In-vivo dosimetry is now a recommendation for all centres providing radiotherapy treatment in the UK (BIR 2007). There is growing demand for an affordable device that will perform accurate IVD on many-field/arc treatments. Point detectors, although a viable option for traditional treatment regimes, are not practical for advanced treatment. EPIDs present a viable solution and there are currently several projects under way to realise the EPIDs potential. Transmission detectors are quite a new concept in the in-vivo field and at present this developing technology lacks large-scale uptake and investigation. The DAVID detector has been shown to be a stable and reliable device, sensitive to small errors in leaf position and is less costly than other array-like detectors. However, the DAVID still needs a baseline, acquired on the linac, to compare subsequent on-treatment measurements to. This is a time-consuming procedure as it not only requires that DAVID to be in place, but the linac delivery, while the DAVID is recording the baseline, needs to be verified by another device - otherwise, the baseline signal could be incorrect (maybe as a result of a transfer error) and subsequent treatments could be verified against this incorrect measurement. Furthermore, with the increasing prevalence of Online Adaptive Radiotherapy (OART) (Franks & McNair 2012, Roussakis 2016), where there is limited time between plan creation and plan delivery, the need for an accurate baseline reading without the use of a linac is desirable.

In this work I intend to investigate the characteristics of the DAVID device with the intention of being able to predict the response of advanced deliveries in order to facilitate intra-fraction monitoring of radiotherapy that is both efficient and safe.

Chapter 3

Preliminary work for predictive algorithms

3.1 Introduction

At the time of writing, published work on the DAVID has been limited to description of its design (Poppe et al. 2006), manipulation of its recorded signal (Looe et al. 2010) and discussions about its use and implementation (Poppe et al. 2010, Chandraraj et al. 2010, Karagoz et al. 2014). In the absence of published work on the topic, this chapter is concerned with ascertaining if the device behaves in a predictable fashion. Initially acting as a proof of principle by showing that the DAVID's response can be predicted in limiting cases, this chapter is a precursor to developing a predictive algorithm for the more general case in Chapter 4.

3.1.1 Simple Deliveries

Preliminary tests were performed that related the response of the DAVID to the machine output and the field size. While maintaining a square field size of 10cm 10, 50, 100 and 200 MUs were delivered by the linac, the response from the two central wires was averaged and plotted as a function of MU (Figure 3.1). Then 100 MU was delivered through square field sizes of 2.5, 5, 10 and 20cm^2, again the average response of the two central wires were plotted as a function of field size (Figure 3.2)

The linear relationship of the DAVID response between machine output (Figure 3.1) and radiation field size (Figure 3.2) implies that it should be possible to predict the response of the DAVID if sufficient information is provided about the radiation field.
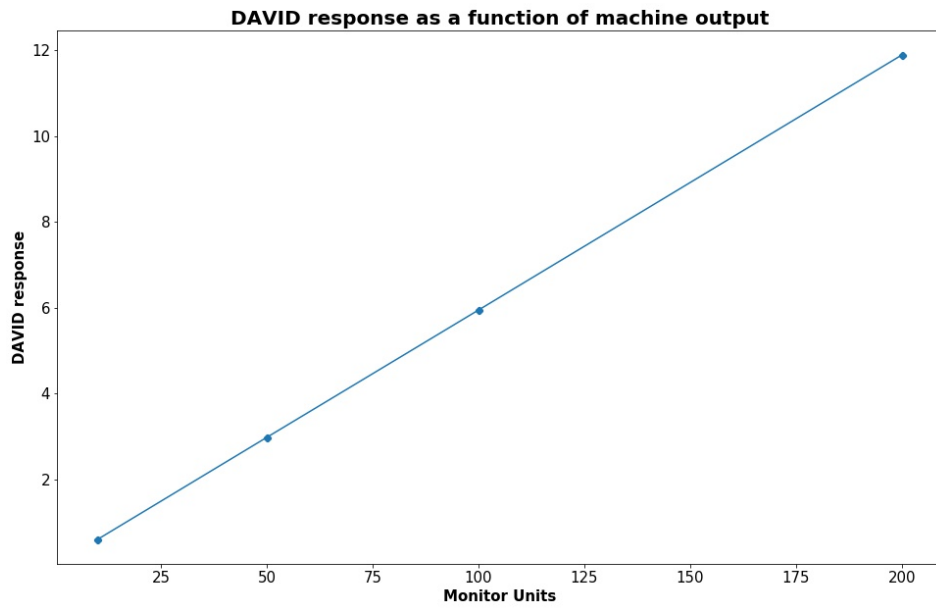
Figure 3.1: DAVID response (average of the two central wires) as a function of delivered MU for a 10cm square field
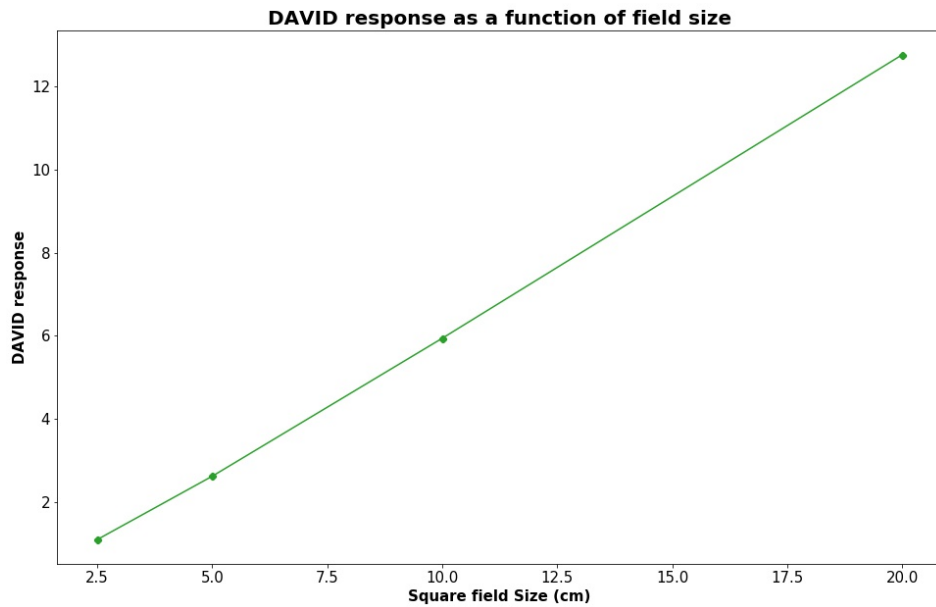
Figure 3.2: DAVID response (average of the two central wires) as a function of field size

The DAVID response at a single wire is a consequence of primary and scatter radiation. The radiation incident on the detector from the aperture created by the leaf pair associated with the wire (leaf separation), is defined as the primary radiation. Radiation incident on the detector elsewhere, but scattering inside the Perspex plate to the collection volume of the wire is defined as the scatter radiation (Figure 3.3). Figure 3.4 shows the response of the DAVID for the case where one leaf pair is open and the others are closed.
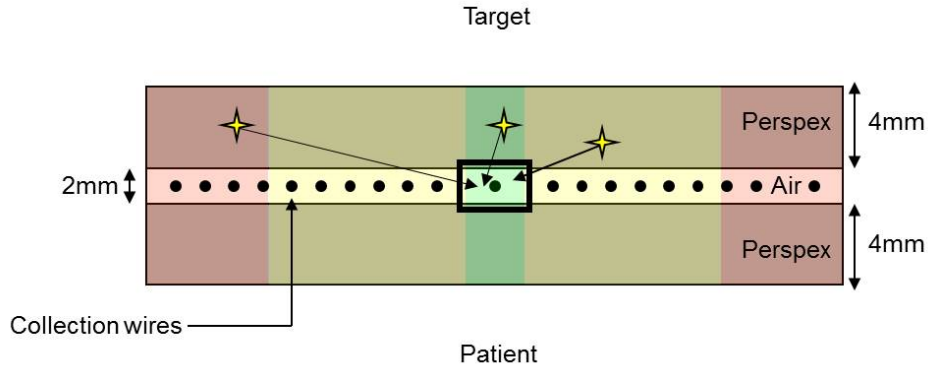
Figure 3.3: Primary and Compton-scatter events contributing to the signal measured in the central wire.
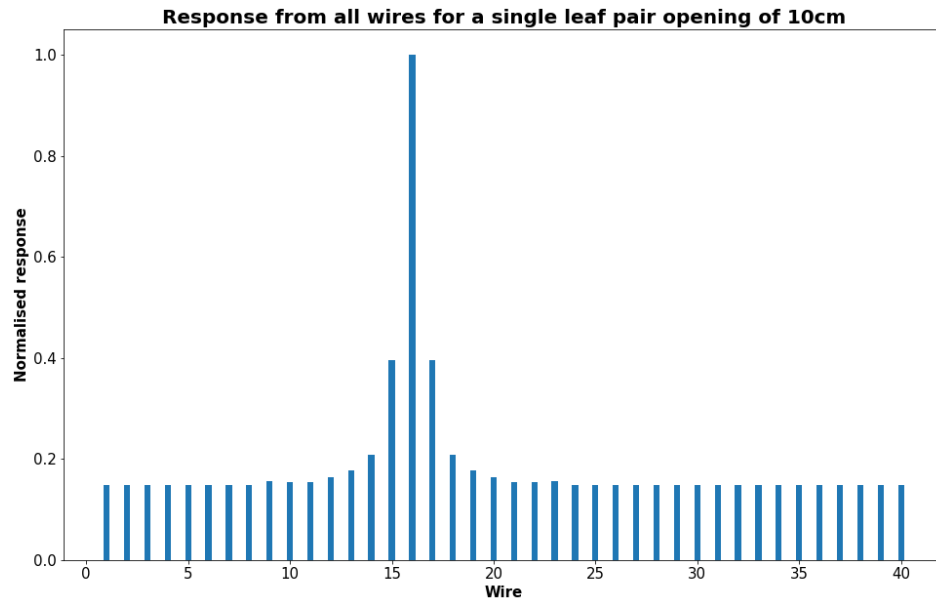
Figure 3.4: Response of all DAVID wires for a single leaf pair opening of 10cm (a leaf pair with each leaf separated from the central axis by 5cm)

The wire associated with the open leaf pair is obviously distinguished by having the largest response (wire 16). The wires adjacent to it (15 and 17) show a response that is clearly above background, and the wires adjacent to these (14 and 18) also show a response that is higher than the background. Apart from wire 16, the signal recorded in this plot is the result of scatter, leakage and penumbral radiation. Although these will contribute to the signal seen in wire 16, the main source is primary radiation.

3.2 Two basic predictive algorithms

A DICOM plan was created that moved equivalent leaves in opposite banks. All leaves started off closed [1]; there was a 10MU delivery then a leaf pair was moved so that the separation was 5mm; after a 10MU delivery each leaf was moved 2.5mm so that the separation was 1cm and another 10MUs were delivered. This process was repeated until each leaf had moved 5cm and the separation was 10cm, all leaves were closed and the process started again. This was performed for each of the 10 central leaf pairs (leaf pairs 16 to 25), the results were sorted into an array referred to hence forth as the R-matrix (Figure 3.8). Figure 3.5 shows the response of the wire associated with leaf pair 20 during the deliveries through leaf pairs 18 to 22. The scatter signal associated with the opening of adjacent leaf pairs can be seen both before and after the primary response.

The response of wire 20 in Figure 3.5 to the opening of leaf pair 20 demonstrates what shall now be referred to as the primary response - the response of a wire as a consequence of the leaf pair above it opening. The response of wire 20 to the opening of leaves 19 or 21 is the adjacent response. The response of wire 20 to the opening of leaves 18 or 22 is the second adjacent response. As described earlier in this chapter, leaf pairs 16 to 25 were sequentially opened and 10MUs delivered through each leaf separation, the data shown in Figure 3.5 for wire 20 was collected for wires 16 to 25. The primary, adjacent and second adjacent responses for each of these wires were collected and averaged to give a general relationship between these responses and the leaf separation (Figure 3.6). The collected responses and the relationships shown in Figure 3.6 were used to create two algorithms that would predict the response of the DAVID to simple fields that only used the 10 central leaves for 10MU deliveries.
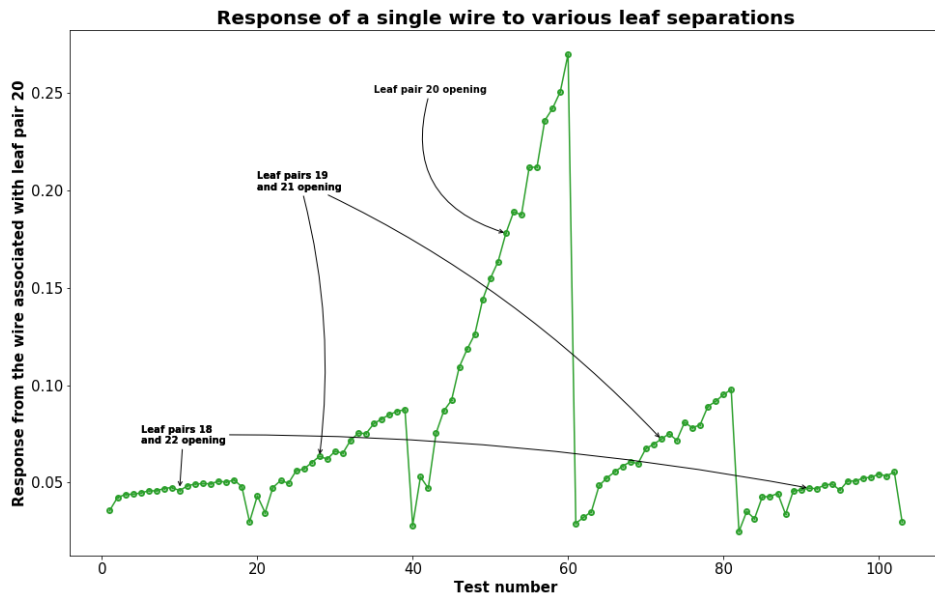
Figure 3.5: The response of the wire associated with leaf pair 20 during deliveries through leaf pairs 18 to 22

---

[1]Closed leaf positions corresponds to a separation of 2.5mm to avoid crashing the leaves into each other due to uncertainties in the linac control systems
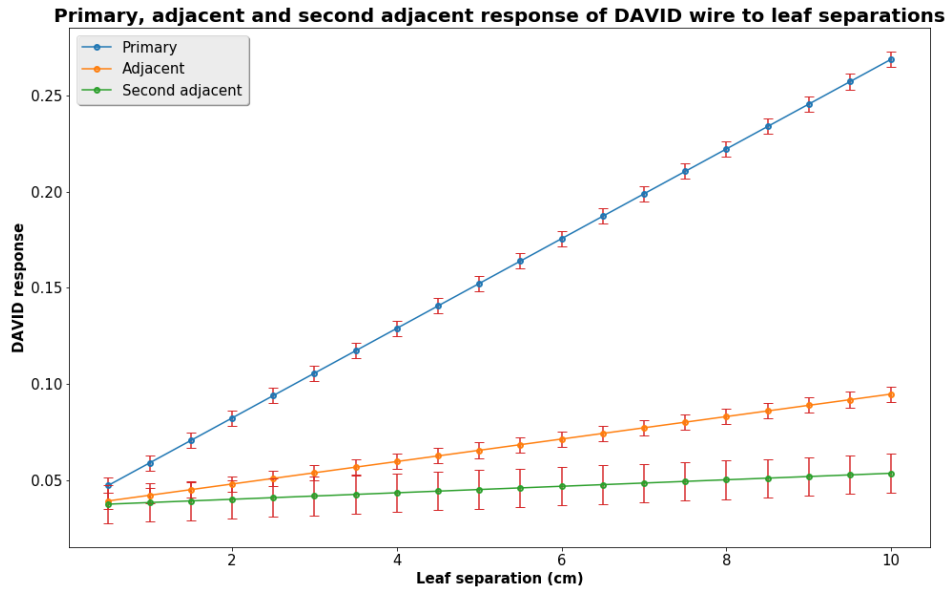
Figure 3.6: Primary, adjacent and second-adjacent responses to a leaf pair opening with deliveries of 10MU. Errors represent +/- 2 standard deviations of the results associated with the response. The second adjacent response was not used in the linear fit algorithm as its inclusion did not improve the accuracy of the predicted signal. This is likely to be because the second adjacent signal is small and the uncertainty associated with it is higher when compared to the primary and adjacent relationships.

3.2.1 Method

Two different approaches were used to write two different algorithms, both with the intention of predicting the DAVID's response. The programs were tested with basic field shapes shown in Figure 3.7. The results are shown in Figures 3.11 to 3.14.
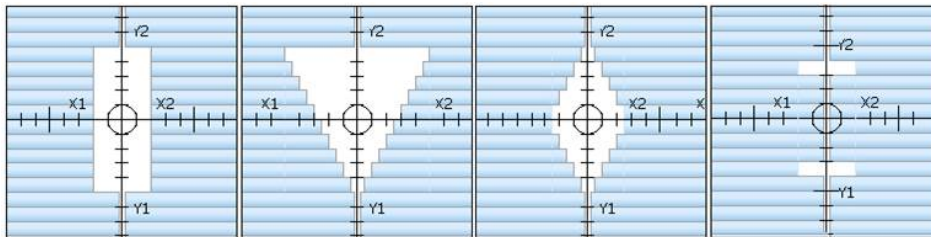
Figure 3.7: Leaf apertures used to test the algorithms

Algorithm A: look-up method

The raw data used to generate Figures 3.5 and 3.6 were sorted into an array (Figure 3.8). A program was written to use this array - the R-matrix - to predict the DAVID response. The algorithm worked by first defining the S-matrix. This is a 1×10 matrix where each element corresponds to the predicted response of wires 16 to 25; initially all these values are 0 but the matrix was populated as the algorithm generated a response due to each open leaf. Then, the algorithm scanned through the MLC positions to identify the leaves that were open and calculate their

separations. For each open leaf pair there was a row in the R-matrix that gave the response for each wire, these values were extracted, the background subtracted and added to elements 1 - 10 in the S-matrix. This process was repeated for each open leaf pair and a background signal added to each of the final S-matrix values. The R-matrix was populated using steps of 5mm. For instances where the leaf separation in the test field was not a multiple of 5mm a linear interpolation was performed between the two nearest values. A flow-chart outlining the whole process is shown in Figure 3.9.
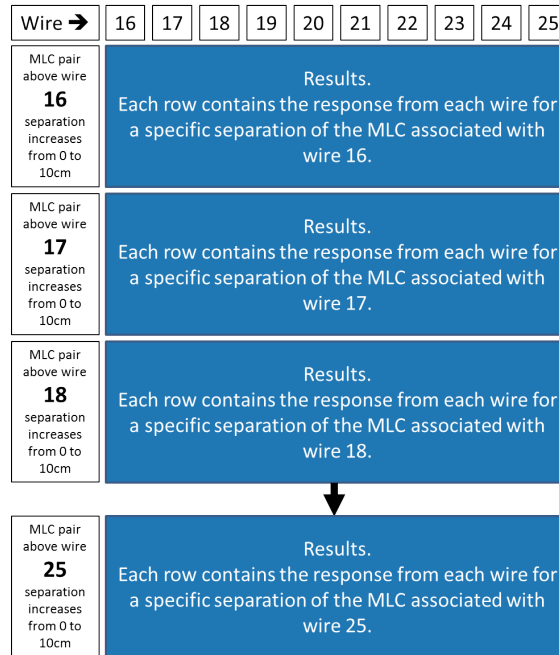
Figure 3.8: The R-matrix. How the data used to create Figures 3.5 and 3.6 was stored to be interpreted by the look-up algorithm.

Algorithm B: Linear-fit method

This approach uses the fact that the primary and adjacent responses of the DAVID are linear when plotted as a function of the leaf separation (Figure 3.5). The gradients used to predict the response were found by averaging all the gradients of the main and lateral responses for all the MLC openings (Figure 3.6). Like the look-up algorithm the linear-fit approach started by first identifying the position of the open leaves and calculating their separation. Assuming that a position of an open leaf pair was above wire 19, the predicted signal of wire 19 could be calculated using the linear relationship between primary response and the leaf separation shown in (Figure 3.6). The predicted signal of wires 18 and 20 could be calculated using the relationship between leaf separation and adjacent response again shown in (Figure 3.6). The predicted signals of wires 18, 19 and 20 as a result of the leaf pair above wire 19 would be added to the elements 3, 4 and 5 in the S-matrix. This process is repeated for each open leaf pair (Figure 3.10). The S-matrix is a 1×10 matrix where each element corresponds to the predicted response of wires 16 to 25. Initially all these values are 0; the matrix is populated as the algorithm predicts the response of the DAVID to the open leaf pairs.

MLC positions are imported and leaf separations calculated. S-matrix defined as empty 1×10 array

Select first MLC pair. $n = 1$.

Next MLC pair. $n = n + 1$

Round all leaf separation **UP** to the nearest 5mm

Round all leaf separation **DOWN** to the nearest 5mm

Find the row in the R-matrix that corresponds to the leaf separation determined in the previous step

Find the row in the R-matrix that corresponds to the leaf separation determined in the previous step

Subtract background from each value

Subtract background from each value

Interpolate between equivalent positions

Add the row to the S-matrix.

Does $n = 10$?

Yes

No

Add on background value and perform empirical linear correction to each element in S-matrix

S-matrix now contains the predicted values

Figure 3.9: The R-matrix contains the response from all 10 wires for any separation of a single leaf pair where the separation is a multiple of 5mm up to a maximum of 10cm. Given a field shape, the DAVID response for each of the leaf separations can be looked up and added together to determine the total. Subtracting the background each time ensures that it is included once - not 10 times. Interpolation allows the algorithm to more accurately predict the response when the leaf separations in the field are not a multiple of 5

48
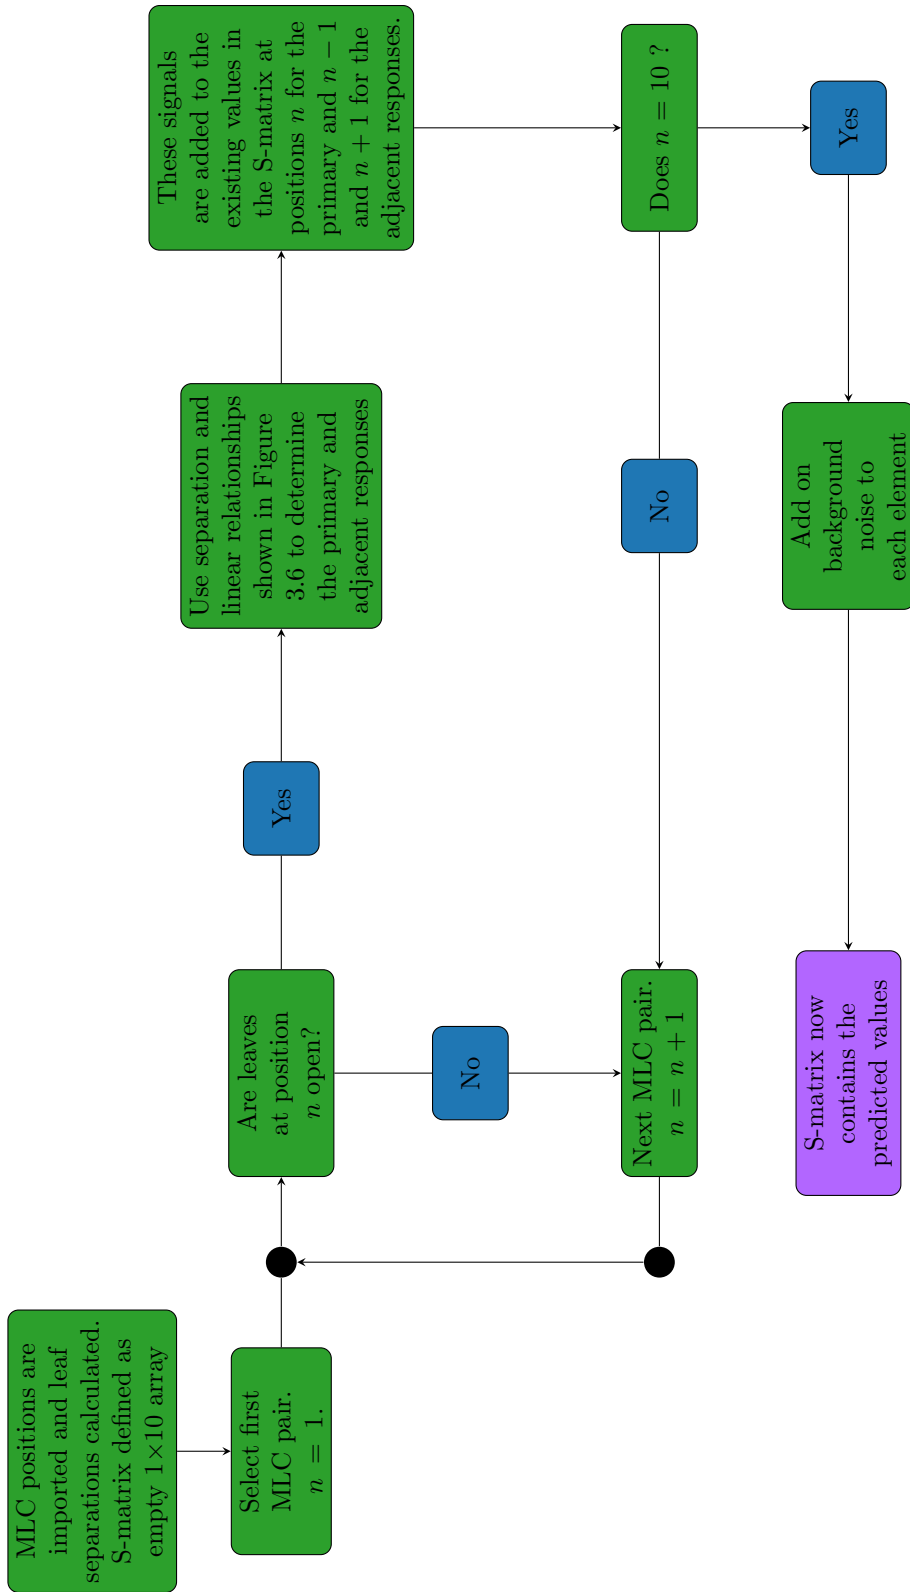
Figure 3.10: A flow diagram outlining the linear-fit algorithm.

Results

Four different test patterns (Figure 3.7) were used to compare the two algorithms. In all four cases, the DAVID was attached to the head and exposed to the four fields dictated by the MLC shape. In each instance, a 10MU delivery was used. The results for algorithms A and B are shown in Figures 3.11 to 3.14.
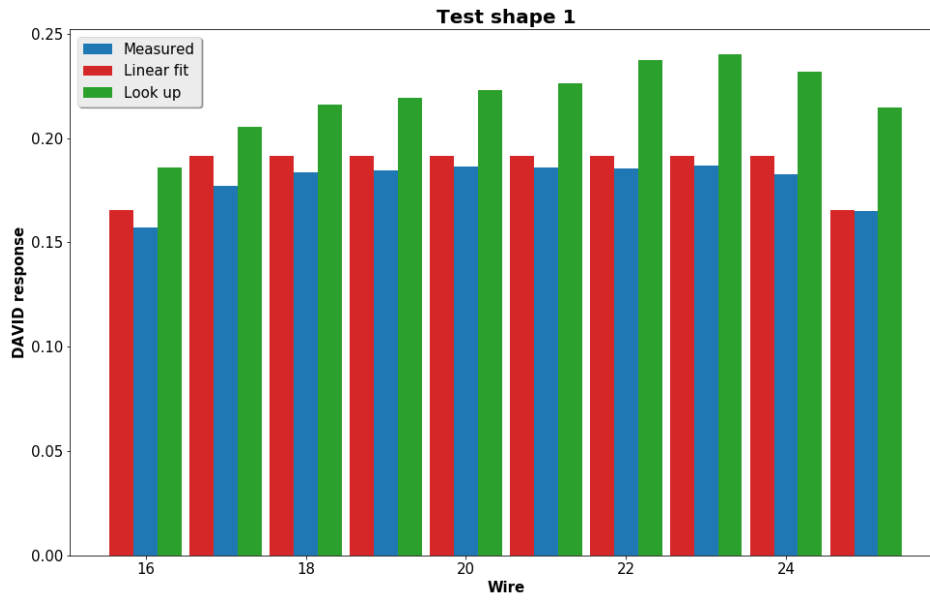
Figure 3.11: Measured DAVID response and signal predicted by the linear fit and look up algorithms for test shape 1
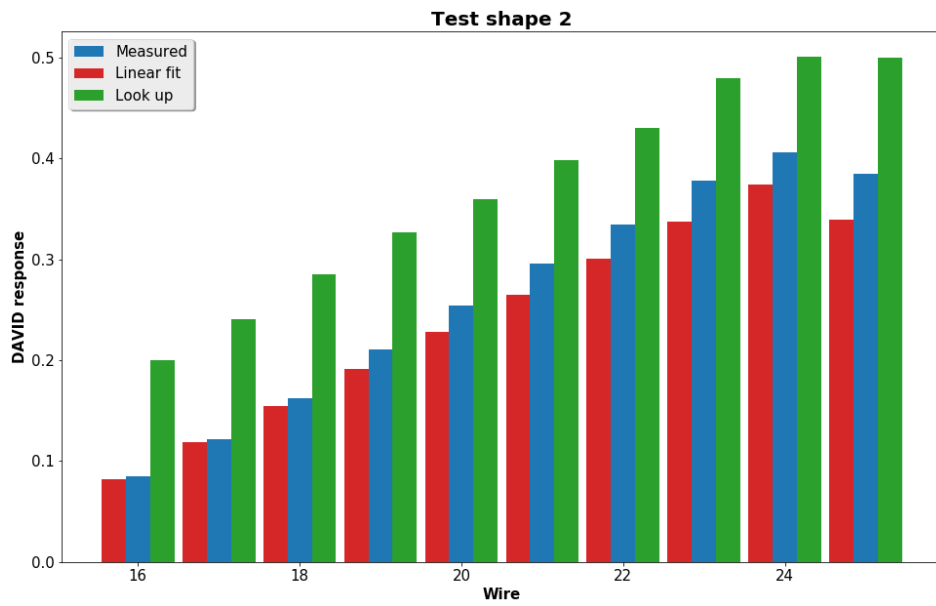
Figure 3.12: Measured DAVID response and signal predicted by the linear fit and look up algorithms for test shape 2
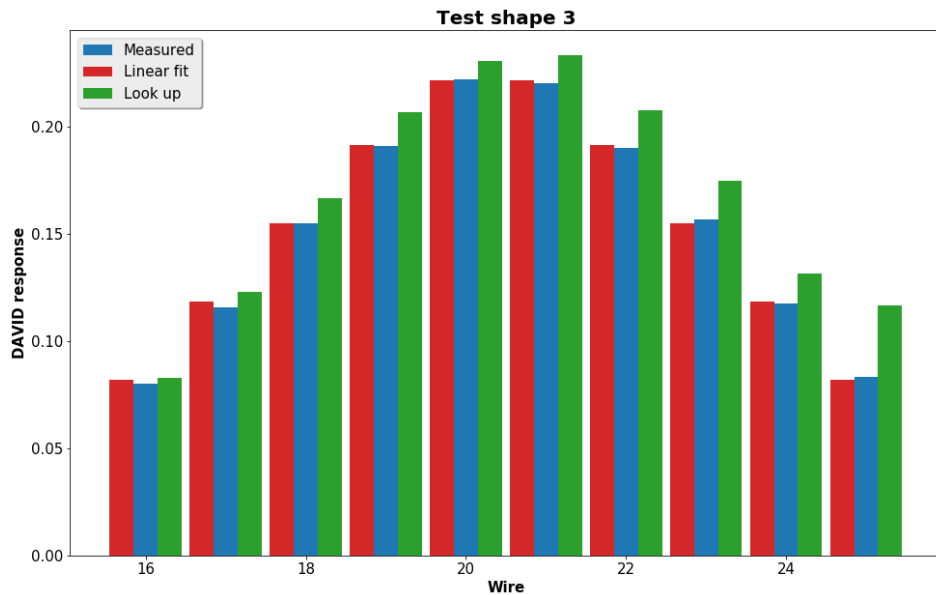
Figure 3.13: Measured DAVID response and signal predicted by the linear fit and look up algorithms for test shape 3
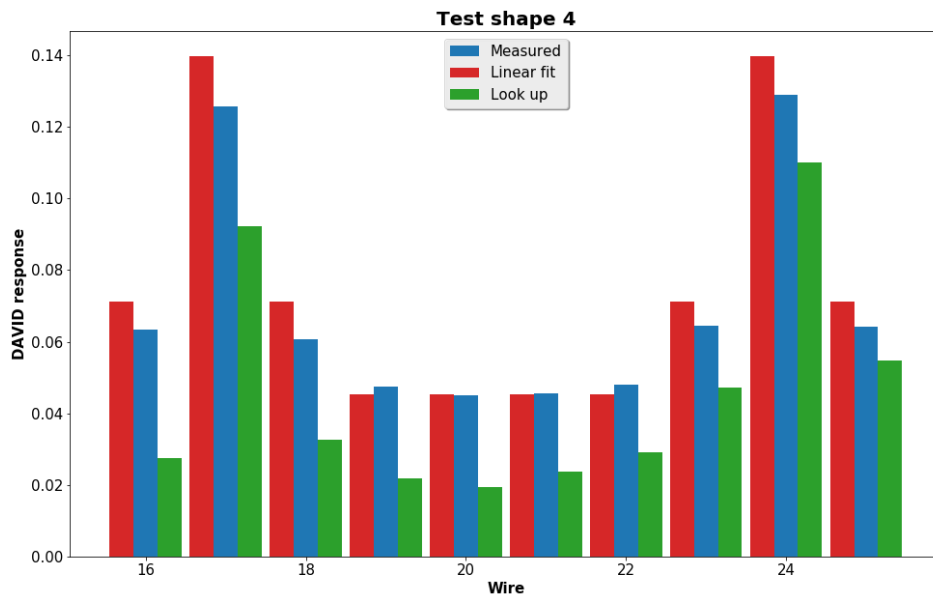
Figure 3.14: Measured DAVID response and signal predicted by the linear fit and look up algorithms for test shape 4

3.3   Discussion

Qualitatively it is clear, even through visual inspection, that algorithm B (linear fit) predicts the DAVID response more accurately than algorithm A (look up). The results shown in Table 3.1, where the sum of the absolute differences for each wire between the predicted and measured signal are shown as a percentage of the total measured signal, show this quantitatively. The difference is thought to be a result of the way the background signal is handled. In the case of algorithm A, the background is subtracted from each wire's response associated with that leaf opening; all these signals are then placed at their allocated position in the 10×10 matrix. The net signal is found by summing the columns and adding, to the resulting line matrix, the background. Poor estimation of the background signal will have a cumulative impact on this calculation. The effect of the background value is less for algorithm B as this only uses the background value once for each wire (as opposed to 10 times with algorithm A).

Algorithm	Test shape 1	Test shape 2	Test shape 3	Test shape 4
Look up	23%	41%	9%	34%
Linear fit	4%	9%	1%	9%

Table 3.1: Sum of the absolute differences between measured and predicted signals displayed as a percentage of the total measured signal.

This exercise has acted as a proof of principle, showing that it is possible to anticipate the DAVID response. The lack of accuracy reflects the simplistic methodology. By increasing the complexity and rigour when defining the variables, the accuracy of the algorithm should improve.

The algorithms, as well as being quite simplistic, are only designed to handle very limited cases. These include:

- Symmetrical displacement of the leaf pairs around the y-axis

- No MLC displacement outside the central 10 MLCs

- 10MU deliveries

For the project to move from a basic proof of principle to something with clinical applications, it will be necessary to improve both the accuracy and the capability of subsequent algorithms. This will be the subject of the next section.

Chapter 4

An analytical model for predicting the DAVID signal

4.1 Introduction

In chapter 3, it was shown that the DAVID response at a wire could be predicted in limited cases. In this chapter, the complexity of the algorithm is increased to include the effects of penumbra, jaw attenuation and varying MUs so that it can be applied more generally. The more advanced algorithm is tested by comparing the predicted response to the response measured by the DAVID for 10 clinical, H&N, step-and-shoot IMRT plans. A methodology for the DAVID that includes the predictive algorithm is proposed. The methodology's ability to detect known radiotherapy errors is discussed. Shortcomings in the algorithm, leading to a loss in accuracy, are shown and suggestions for developments are made. The work and findings discussed in this chapter were published during the writing of this thesis Johnson et al. (2014).

4.2 Increasing the complexity of the algorithm

4.2.1 Introduction

In chapter 3, two algorithms were devised that predicted the response of the DAVID for 10MU deliveries through basic, static, symmetric fields. In this section, the previous work is expanded on to allow the prediction of any field. Work was done to model the effects of:

- Varying MUs

- Scatter (more thoroughly)

- Jaw attenuation

- Penumbra

These are discussed in turn.

4.2.2 Varying monitor units and leaf separation

Introduction and Method

In chapter 3, only deliveries of 10 MU were modelled. This represents a typical IMRT segment; in practice however, there will be a range of MUs delivered per segment (Qi & Xia 2013, Stieler et al. 2011). This section is concerned with devising a scheme to accommodate varying MU deliveries.

A beam sequence was devised where 5, 10, 25, 50 and 100 MUs were delivered through leaf separations of 20, 40, 70, 100, 150, 200, 300, 350 and 400mm. The response measured by the collection wire associated with leaf pair 20 is shown in Figure 4.1. The response of the DAVID for a specific leaf separation is linear. The gradient associated with the response was calculated for each leaf separation, the results were collected and plotted in Figure 4.2 .

Results
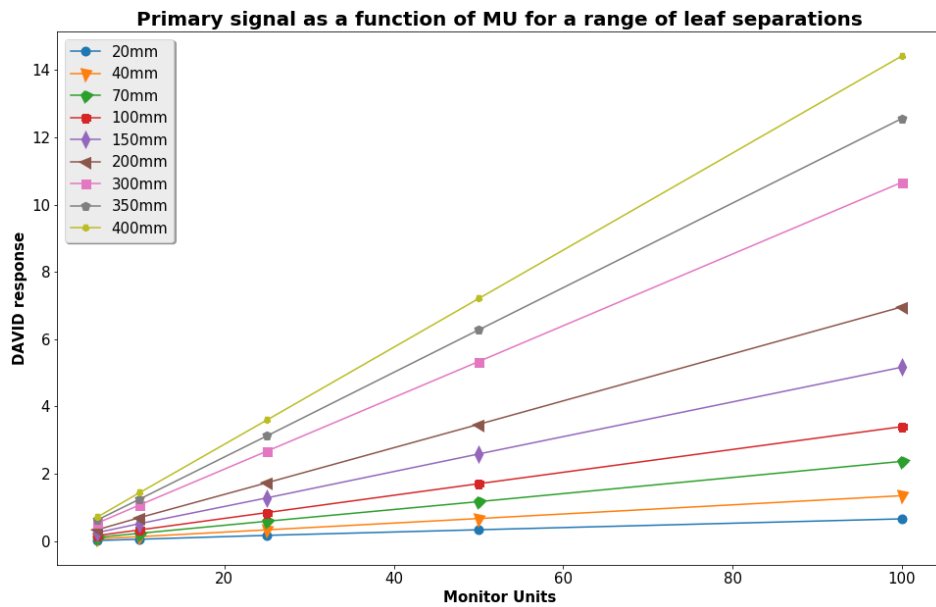
Figure 4.1: Signal as a function of MU for a variety of leaf separations (shown in the key).
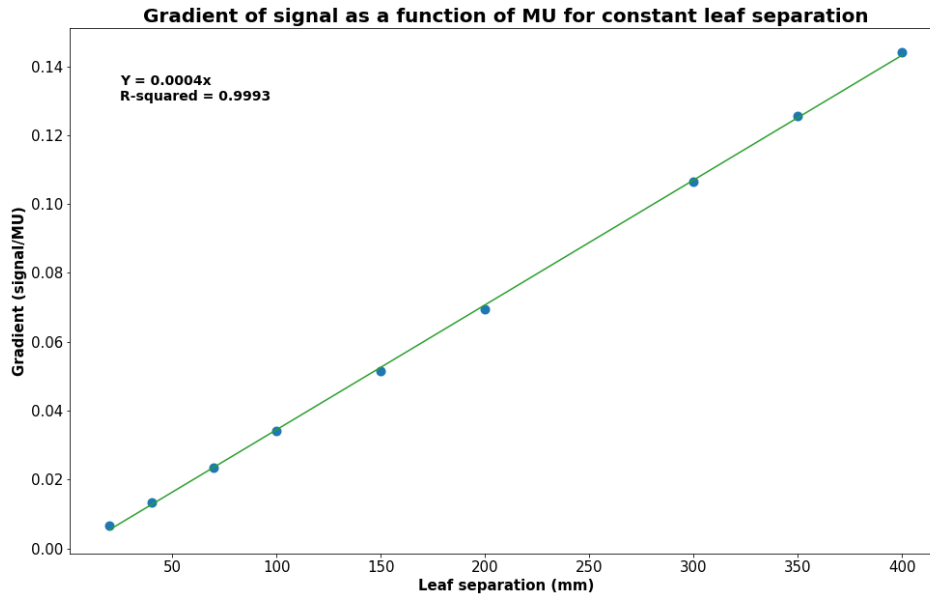
Figure 4.2: A graph showing the gradients of the MU response to different leaf openings (Figure 4.1)

Discussion

Figures 4.1 and 4.2 show that the response of the wire associated with an opening leaf pair (primary response) is linearly proportional to both the MU and leaf separation. So, for a given MU and leaf separation, Figure 4.2 shows that the gradient (G) of the MU response can be determined (Equation 4.1) and given the MU response Figure 4.1 shows that the primary response can be calculated (Equation 4.2).

$$G = C \times LS \tag{4.1}$$

$$Primary\ Response = G \times MU \tag{4.2}$$

or:

$$Primary\ Response = C \times MU \times LS \tag{4.3}$$

Where $C = 0.0004$ (Figure 4.2) , MU is Monitor units and LS is the leaf separation.

4.2.3 Scatter

Using the data collected in Section 4.2.2, it was possible to anticipate the signal output from the wire associated with an open leaf pair for any separation and for any amount of MUs delivered. However, leaf arrangements used in clinical practice will have more than one leaf open, meaning that the signal at any one wire will be a result of the primary fluence above the wire and the scatter caused by radiation incident on other parts of the detector. The scatter effect is seen in Figure 4.3. In the algorithms devised in Chapter 3, this effect was modelled very basically, neglecting the scatter from points further away than the wire immediately adjacent

to the primary wire.

Response of a 10MU delivery through a 10cm strip
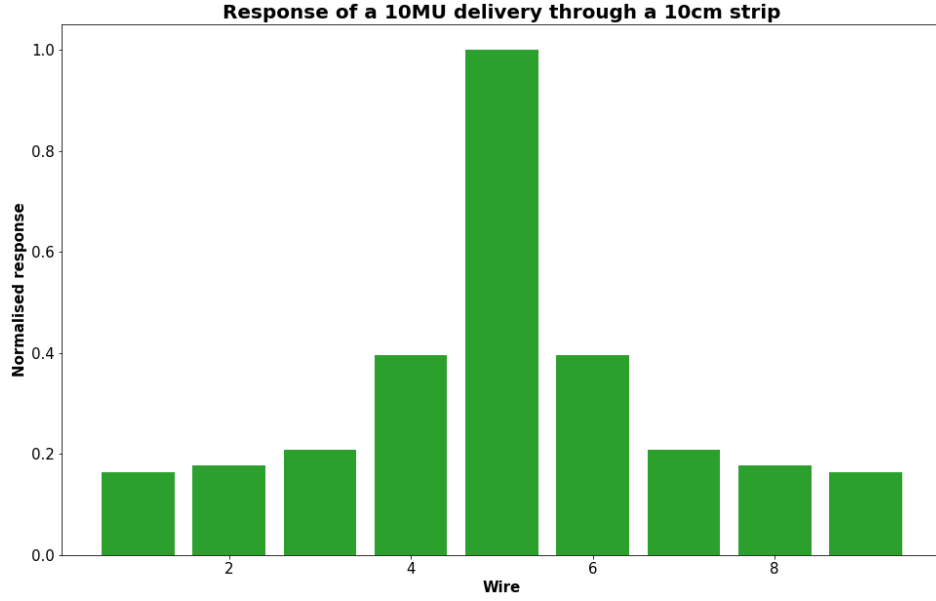
Figure 4.3: lateral response of 100MU being delivered through 100mm leaf separation with the jaws open. Taken from the data collected in the experiment performed in Section 4.2.2

The data collected in Section 4.2.2 were used to model the effect of scatter. The response of the four wires either side of the primary wire were normalised to the main wire for each exposure. This was repeated for all the MU deliveries at this leaf separation (10,20,50 MU etc.). Then, the average of these results was taken for equivalent wires. This procedure was repeated for each leaf separation and the results collected. Finally, an average was taken for the results at equivalent distances. The normalised and averaged results were plotted as a function of distance from the primary wire; the response was seen to be exponential (Figure 4.4). The lateral response shown in Figure 4.4 is characterised by equation 4.4.

$$L(d,w) = e^{(-1.447 \times D(d,w))} \tag{4.4}$$

Where:
L is the response of the wire w as a fraction of the signal at wire d.
$D(d,w)$ is the distance between wire w and wire d

Knowing from Equation 4.3 that the primary response of a wire is the product of some constant, the leaf separation and the delivered monitor units, it can now be said that:

$$R_w = C \times MU \times Sep_d \times L(d,w) \tag{4.5}$$

Where:
R_w is the scatter response at wire w from the radiation incident on the device at wire d
Sep_d is the separation of the MLCs above wire d - where the scatter response is measured from.
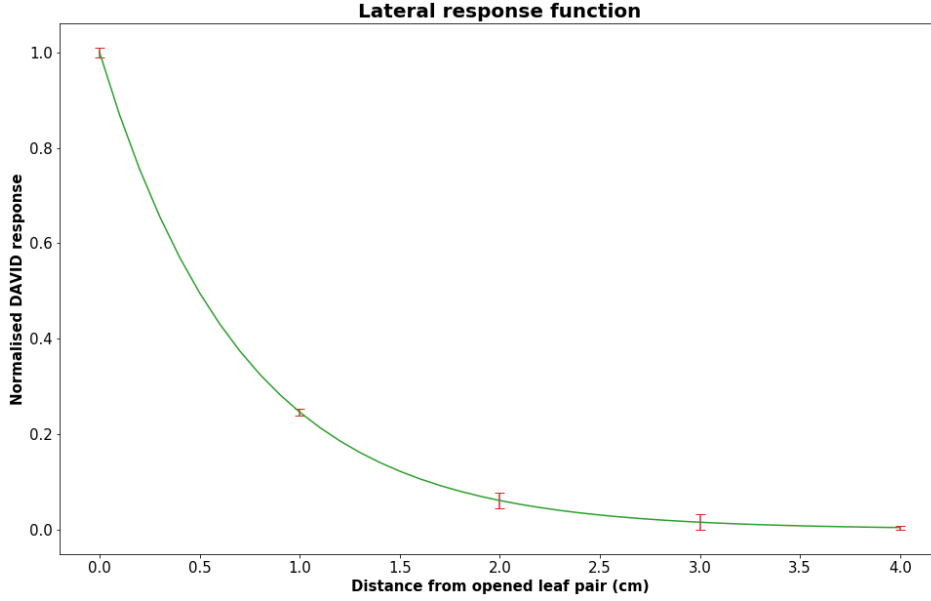
Figure 4.4: lateral response of 100MU being delivered through 400mm slit separation taken from the data collected in the experiment performed in Section 4.2.2, the errors represent 2 standard deviations of the measured data.

L is the lateral response function (equation 4.4)

C is a constant

MU is the delivered monitor units

So the total response at a wire can now be described as:

$$Total = \sum_d C \times MU \times Sep_d \times L(d, w) \qquad (4.6)$$

Where d is the number of wires, in this case, 40.

Note that when the equation is calculating the response from the leaf pair immediately above wire w, the distance $(D(d,w))$ becomes 0 so $L(d,w)$ becomes 1 and the response is determined by:

$$Total = C \times MU \times Sep_d \qquad (4.7)$$

This was shown earlier (equation 4.3).

Looking at this result, it can be seen that the total response to an IMRT segment can be estimated through the convolution of the lateral response with the leaf separation matrix with the result multiplied by the MU and the gradient factor (C):

$$SegmentResponse = [LSmat * L(d, w)] \times C \times MU \qquad (4.8)$$

Where $LSmat$ is a 1×n matrix where n is the number of leaf pairs. Each element $LSmat$ contains the separation distance of the corresponding leaf pair. $L(dw)$ is the lateral response function, C is the MU response gradient and MU is the monitor units.

4.2.4 Leaf leakage

In Chapter 3, all the calculations were performed with the jaws fully retracted. This will not be the case for clinical deliveries where the jaws are used to reduce unnecessary dose to the patient as a consequence of leaf and inter-leaf leakage. The consequence of this is that leaves inside the field will have a different leakage contribution to the signal than those outside the field.

Initial tests were done using a test field where leaves 10 and 31 were opened (Figure 4.5) and 10, 20, 30 and 40MU were delivered through the set up. The DAVID response was predicted using equation 4.8. The predicted response was compared with the measured signal; the results can be seen in Figure 4.5. Figure
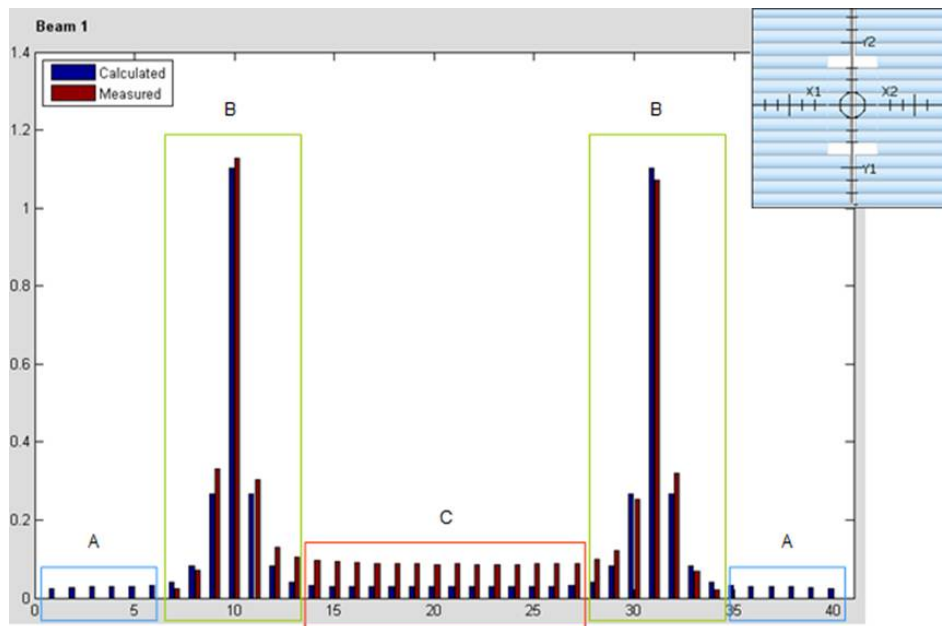
Figure 4.5: Convolution and measured results from the delivery of 20MU through two open leaves (field shape in top right).

4.5 shows a reasonable agreement between the predicted and measured results at the points where the leaves are open (sections b), but there is poor agreement in sections A and C. Despite wires 5 and 15 both being equally displaced from wire 10, wire 5 records no signal whereas wire 15 does. Initially it was thought that the signal at wire 15 could be a consequence of the scatter from both of the open leaves. Wire 5, although the same distance from wire 10, is much further from wire 31 than wire 15; this would suggest a shortcoming with the convolution. It was noticed, however, that the signal recorded in section C is much more than double the lowest signal recorded outside the field (wires 6 and 34). And, wire 20 is further from the open leaves than either of these low values are from the primary response. The key difference between the field in section C compared to that in section A is that, in section C, the leaves are not backed up by the jaws. To accommodate this, a MU-dependent leaf-leakage factor was added to each convolution result that was inside the open jaws. For plans where the jaws are fixed throughout the delivery, this factor will be applied to the same wires for each segment. Where the jaws do move, the wires having this factor applied will

differ from segment to segment.

It is also evident from Figure 4.5 that the DAVID does not detect low fluence scatter as no signal is recorded in the wires at the edges. The exponential nature of the lateral response function meant that the algorithm will always get some results outside the field. To stop this happening, thresholding is used to set the signal at the wires more than four wires from the field edge to zero.

4.2.5 Penumbra

Introduction

As a consequence of geometry and the finite attenuation of the beam, the edge of the radiation field is not completely distinct. The fluence drops from 80% of the main beam to 20% at the edge of the field over a finite distance. This region is termed the penumbra (Khan 2010).

The finite size of the source is responsible for the geometric penumbra – this is the region with a partially blocked view of the source. As shown in equation 4.9, the width of the geometric penumbra ($PGeo$) is a function of the width of the source ($dsource$), the distance from the source (F) and the distance between the source and the end of the collimator (S) see Figure 4.6

$$P_{Geo} = d_{source} \times \frac{F - S}{S} \tag{4.9}$$

The transmission penumbra, a consequence of photons from the distal end of a source passing through the collimator, will increase the width of the overall penumbra still further. Efforts are made to ensure that the focal spot in a linac are small ($\sim$mm), however the flattening filter also acts as a secondary photon source, broadening the apparent focal spot. The scattering of high-energy photons and the scatter of secondary electrons outside the beam edge (without being compensated for by electrons being scattered back in to the beam) also increases the penumbra (Mayles et al. 2007).

Effect on the DAVID signal

As a consequence of the penumbra, the DAVID signal is seen to drop outside the treatment field. Modelling this, however, poses a challenge as the signal is a consequence of:

- Scatter radiation - from the radiation incident on the device in the treatment segment

- Primary radiation from the segment's penumbra

- Scatter from the segment's penumbral radiation.

With the intention of making a simple algorithm that will be quick and easy to use, it is desirable to have limited input variables allowing easy set-up and implementation. With this in mind, an assumption was made that the out-of-field signal was proportional to the signal of the last in-field wire. This is justified, at least qualitatively, by looking at the components of both signals. The signal recorded at the last in-field wire will be proportional to the scatter as a result
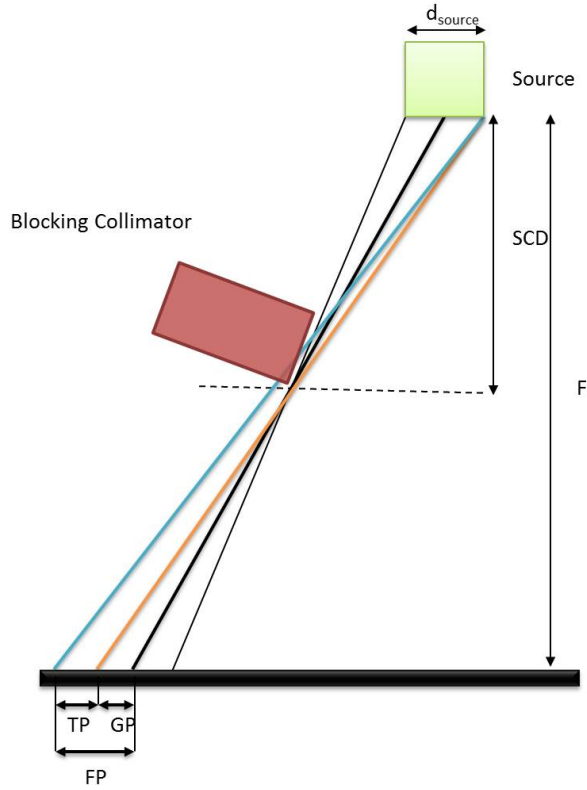
Figure 4.6: Diagram showing the origin of the penumbra; dimensions have been exaggerated to illustrate the effect. Geometric penumbra (GP), transmission penumbra (TP) and full penumbra (FP)

of radiation incident elsewhere on the device, and the primary radiation from the leaf pair associated with the wire. The penumbral radiation is made up of scatter radiation from elsewhere in the device and some component of the primary radiation from the last, open, in-field leaf pair. To test this assumption, the out-of-field responses for 16 segments of an IMRT treatment beam were measured. These were normalised to the last in-field signal. Results are shown in Figure 4.7; the error bars are two standard deviation of the normalised results for the respective points. This was implemented in the code by first calculating the in-field responses then the out-of-field responses by multiplying the last in-field response by the fractional values shown in Figure 4.7. For example, in the case of a 10×10 square field symmetric about the beam centre, the signal associated with the five leaves either side of the centre would be calculated using the lateral response function and multiplying by the delivered MU and MU factor. The penumbral component of the predicted signal for the wires associated with leaves 6, 7, 8 and 9 would be calculated by multiplying the signal calculated at the wire associated with leaf pair 5 by 0.82, 0.22, 0.05 and 0.01 respectively. The uncertainty associated with the penumbral signal is much higher than the in-field signal. This is due to the spread in the results that determined the size of the error bars in Figure 4.7. This uncertainty was summed in the algorithm to calculate the total uncertainty associated with the predicted signal for each wire, this is discussed in Section 4.4.
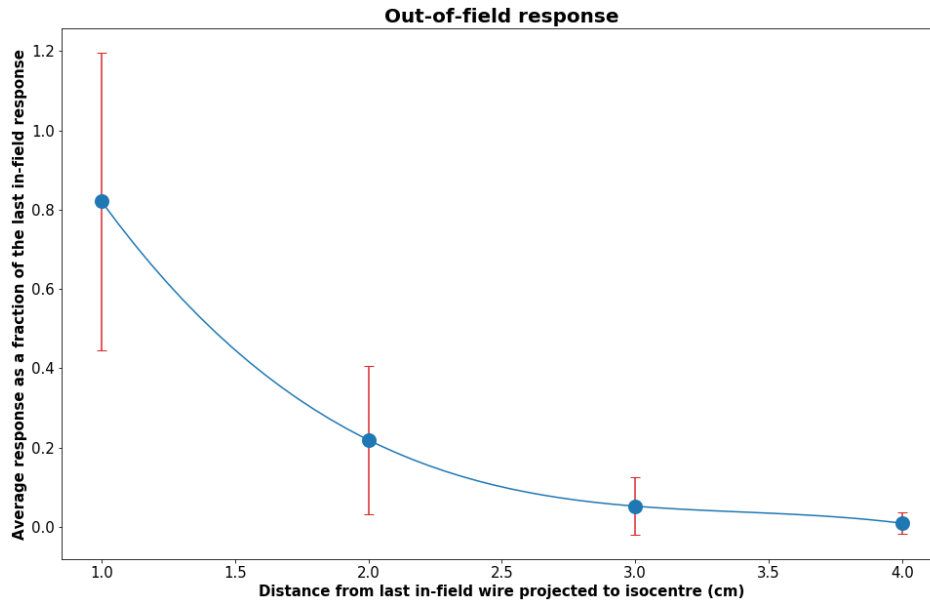
Figure 4.7: Average penumbral response as a fraction of the last in-field response for 16 segments of an IMRT field.

4.3 Iterative assessment

The preceding sections in this chapter have outlined several parameters that, given a treatment plan (that includes MU/seg and each segment's leaf positions), are able to anticipate the DAVID response. To summarise, the factors are: MU factor (Section 4.2.2), lateral response exponent (Section 4.2.3), leaf leakage (Section 4.2.4) and 4 penumbral factors (Section 4.2.5). The values for each of these factors were used in a MATLAB algorithm and applied to an imported treatment plan. The subsequent predicted signals were compared with treatment plans delivered on the linac. H&N plans are typically highly modulated making them a good choice for use in this development work as they provide of more rigorous test of the algorithm. 10 clinical H&N plans were selected and delivered three times; the average for each beam and entire delivery (5 beams) were taken to reduce any variability in linac output or DAVID response (the difference was never seen to be >1%). Initial results were disappointing; Figure 4.8 shows the response for an entire H&N IMRT delivery (5 beams ~70 segments).

To try and improve the predicted response, a program was written that produced a predicted response and compared it to the measured DAVID signal. For each of the algorithm factors listed above (lateral response exponent, 4 penumbral fitting factors and the MU constant, C), the iterative process used a value lower than the original as a starting point - this meant that if the optimal value was less than the original it would not be missed by the stepping process. The difference between the predicted and measured signal was calculated using the starting value, then a single unit was added to the starting value and the difference between the predicted and measured signal calculated again. If the result was better using the second value another unit was added on and the process repeated until the results starting getting worse. Then the unit was divided by 10 and was subtracted from the test value and the process repeated - this way each decimal place was dealt
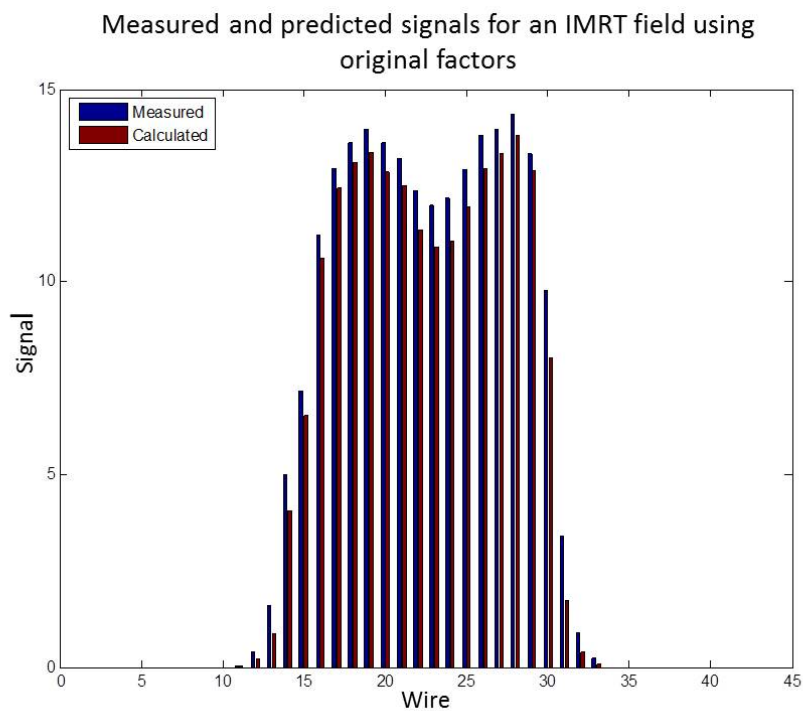
Figure 4.8: Predicted and measured signals for an entire H&N IMRT delivery (5 beams ~70 segments). Average local percentage difference = 23.4% Average global (when the absolute difference was compared with the smallest of either the maximum predicted signal or measured signal) percentage difference = 5.3%

with in turn, Figure 4.10 outlines the process. Let us assume that the original value for a parameter was 2 and the optimal value 2.82. 1 would be used as a starting value and this would be compared to the result gained using 2. The result with 2 would be better so then 3 would be used, this would be better than 2 so 4 would be used however, this would be worse than 3. This process demonstrates that the optimal value is somewhere between 2 and 3, to determine the second decimal the order of precision was increase by a factor of 10 the results were calculated using 2.9, then 2.8 then 2.81 and ultimately 2.82 (Figure 4.9). This process was done for each factor in turn and could go on for as many orders of precision it was necessary until the impact the changes were making on the results were negligible.
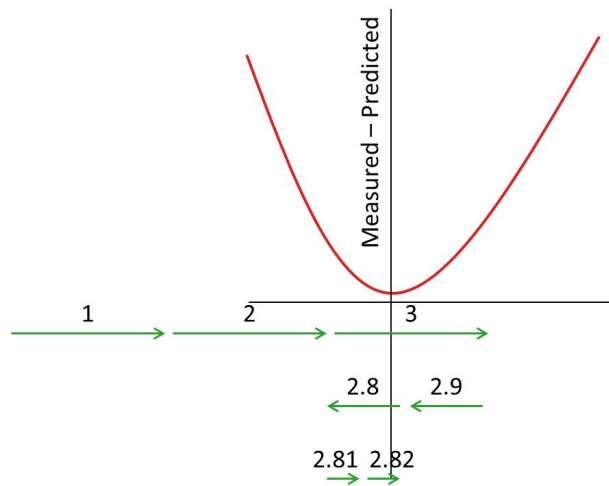
Figure 4.9: The stepping process associated with achieving optimal value for each model factor. The local minima - where the difference between the measured and predicted values is smallest is found by changing the value for each factor involved in the predictive algorithm by increasing levels of precision
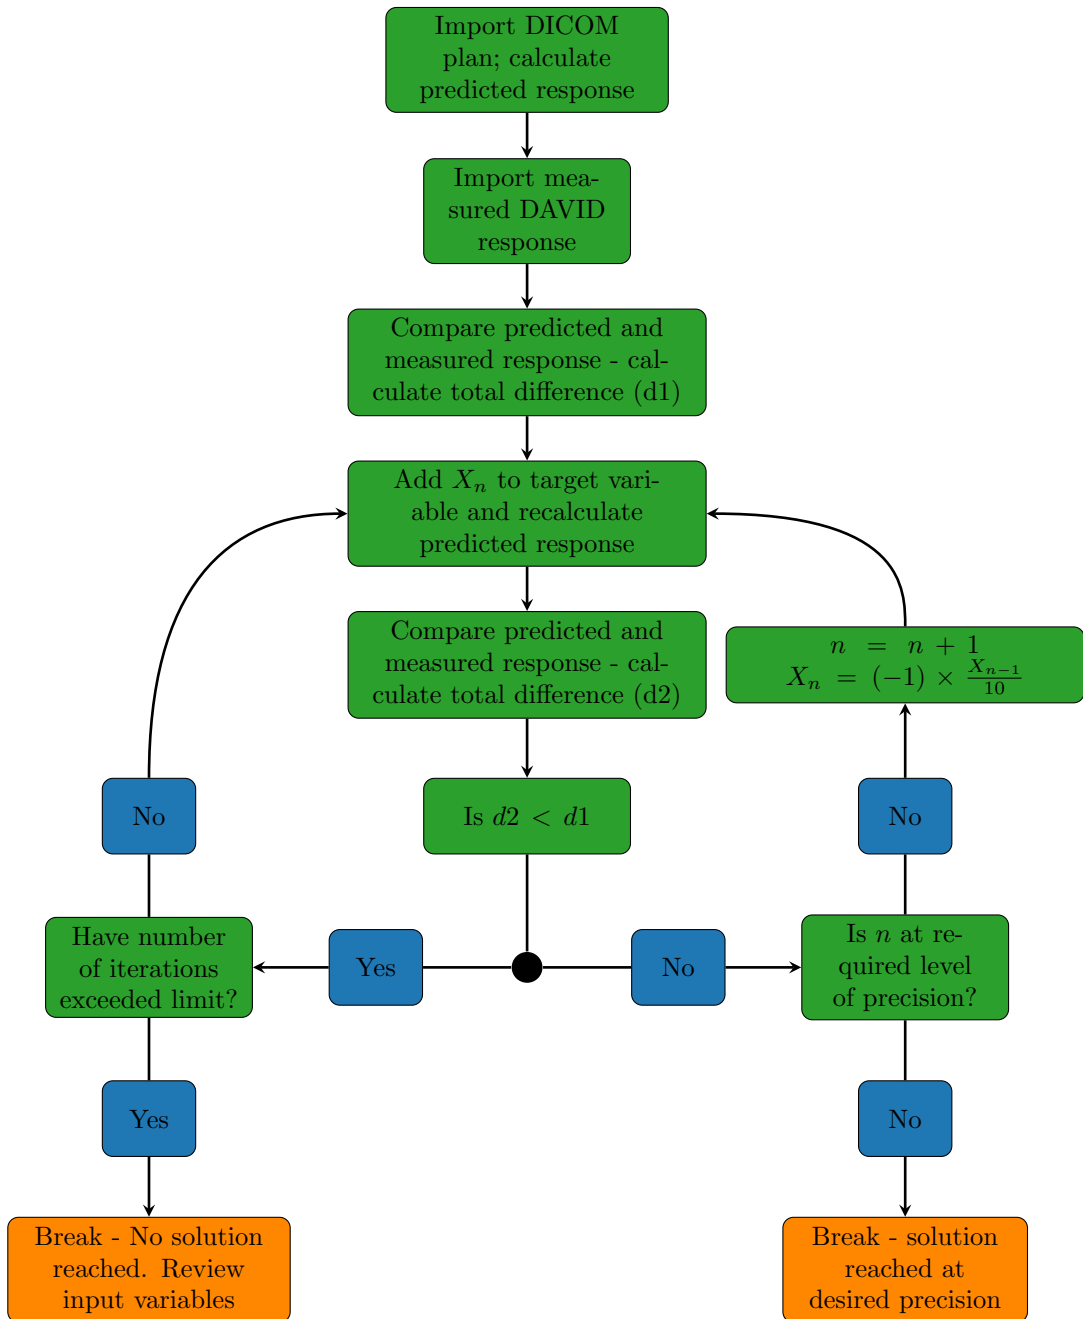
Figure 4.10: Flow diagram outlining the iterative process for tuning the variables associated with the predictive algorithm

Although the iterative correction produced good results (Figure 4.11), are the values it produced still sensible and valid in the context of the work that has been done? If the values produced through the iterative method were wildly different from expected values, it would suggest a shortcoming in the theory. The values that were seen to change significantly were the lateral-response exponent and the penumbral factors. The predictive algorithm was modified to include the iteratively-optimised values. The MU response factor (C from Equation 4.8) was not seen to vary with the optimisation, so kept the same. A summary of the factors used in the algorithm can be found in Table 4.1.

Measured and predicted signals for an IMRT field using iteratively-derived factors
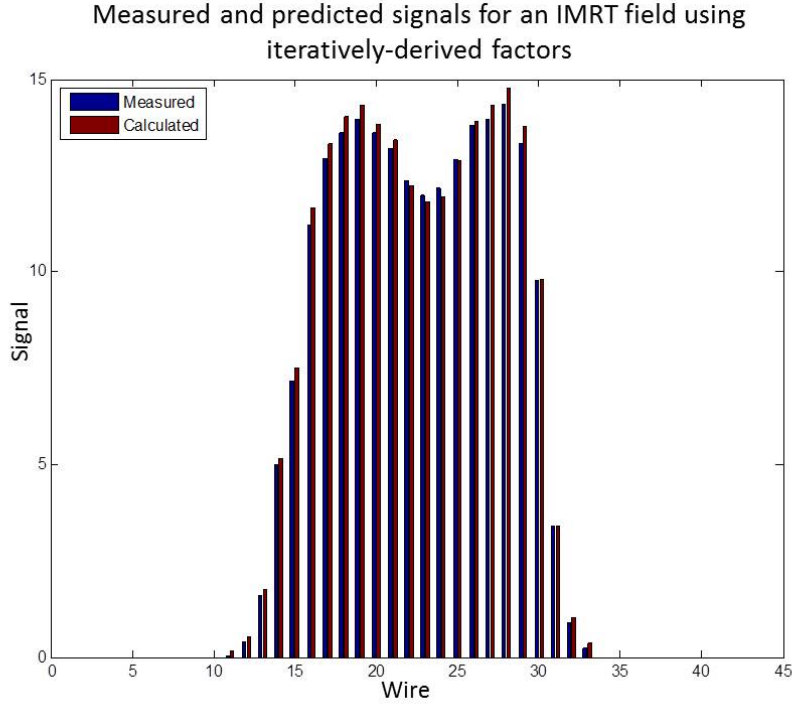
Figure 4.11: Predicted and measured signals for an entire H&N IMRT delivery (5 beams ∼70 segments). Average local percentage difference = 10.9% Average global (when the absolute difference was compared to the smallest of either the maximum predicted signal or measured signal) percentage difference = 3.1%

Factor	Original	Iterative result
Penumbra 1	0.821	0.932
Penumbra 2	0.219	0.347
Penumbra 3	0.052	0.106
Penumbra 4	0.009	0.039
C (MU constant)	0.0004	NA
Lateral response Exponent	1.447	1.404

Table 4.1: Original and iteratively-derived values for the factors in the algorithm

To ensure that the new values were in line with the theory discussed in this chapter, they were compared with the values derived from measured data. The lateral response derived from incrementally opening leaf pairs 1, 5 and 10 was plotted against the lateral response derived through the iterative process. They were

found to be consistent (Figure 4.12). These results also highlighted that the same lateral response function is valid across the whole detector. The penumbral values, derived through iteration, were also plotted against the ones found experimentally; they were found to be consistent (Figure 4.13).

.
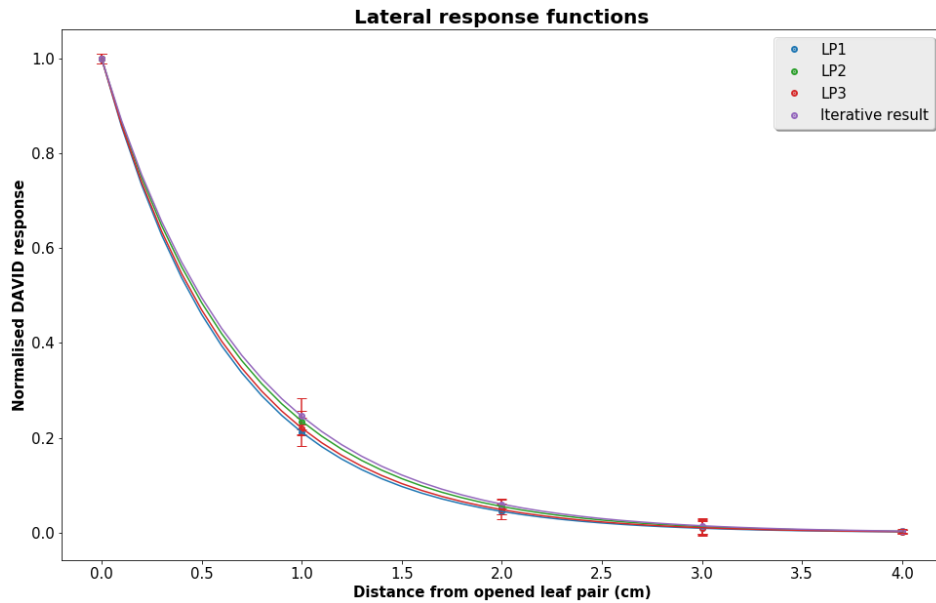
Figure 4.12: The lateral response derived through experiment by sequentially opening leaf pair 1 2 and 5 plotted against the lateral response derived through iteration.

The iterative correction worked well and, should a commercial piece of software ever be made, including this would make implementation quite easy. The user would need to get some approximate values for the predictive algorithm's variables and then the software, given a verified test plan, could tune itself.

.
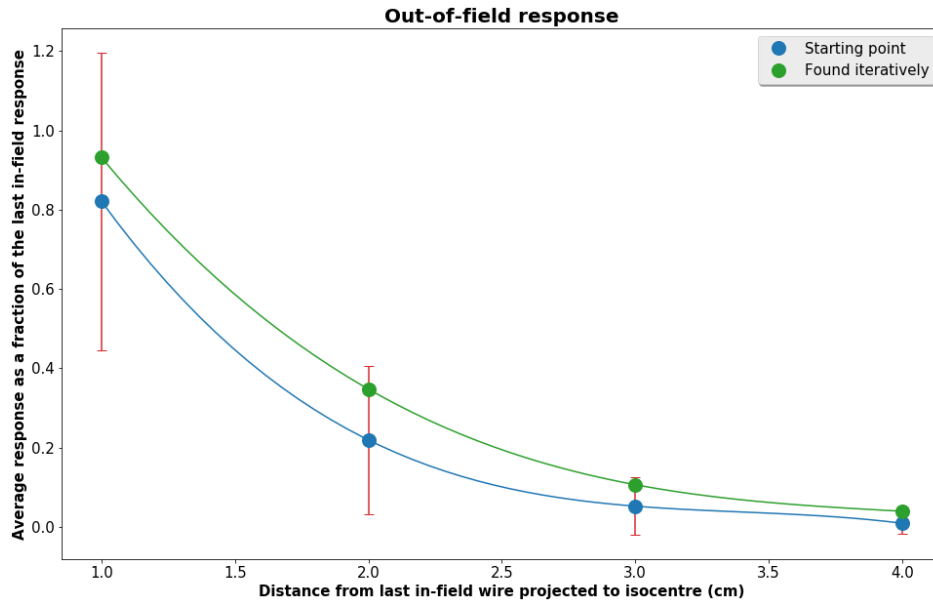
Figure 4.13: Measured penumbral values (blue points and line, from Figure 4.7) plotted against those derived through iteration.

4.4 Errors and assessment

4.4.1 Introduction

There needed to be a way of assessing the signals that was stringent enough to highlight an error, while at the same time, not triggering when minor discrepancies were detected. The percentage difference between two wires seems a sensible place to start. However, in the case of out-of-field signals, that have small values – and little clinical impact – the percentage difference can often be quite high (>50%) (Figure 4.14) and not indicative of a fault in the machine's output or delivery. One possibility would be to ignore the out-of-field signals, although this approach might miss some machine errors – eg leaves/jaws staying open outside the field throughout the delivery. With this in mind, it was necessary to design a form of assessment other than absolute percentage differences in the measured signals.

Normal distribution method

One idea was to look at the normal distribution of the difference in the results – any distribution not centred on zero or with a high standard deviation would indicate an error in the delivery. While this is an attractive idea in principle, the fact that there are <40 results means that creating the normal distribution is not always statistically justified (Figure 4.15).
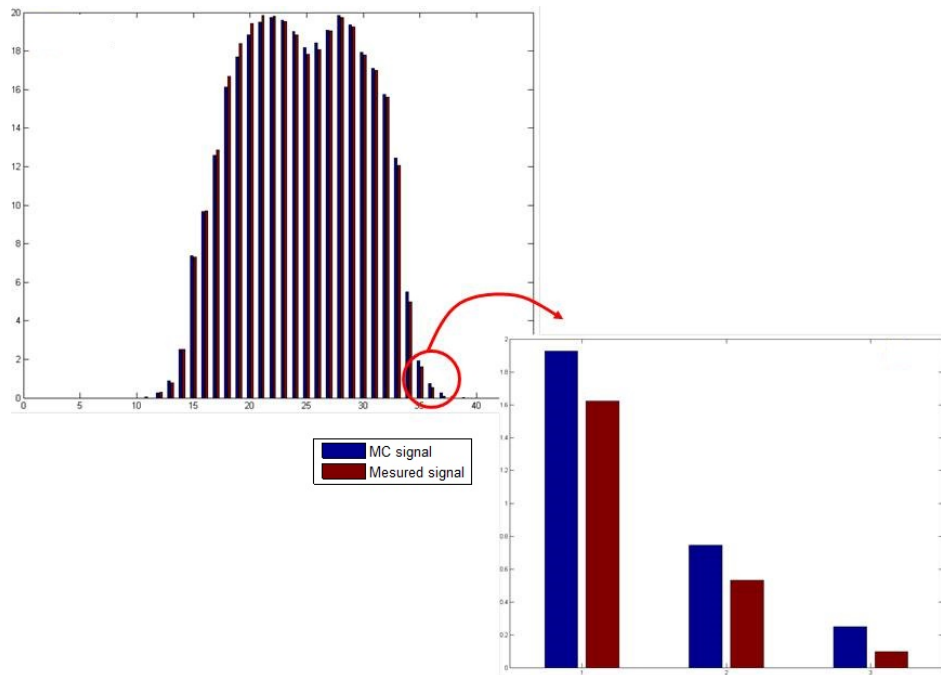
Figure 4.14: Beam 4 of a five-field H&N IMRT delivery (top left) and a close-up of the measured and predicted signals from wires 35–37 (bottom right). Despite having large relative percentage differences, the absolute differences will have negligible clinical impact. Alerting the user to this error would be pointless
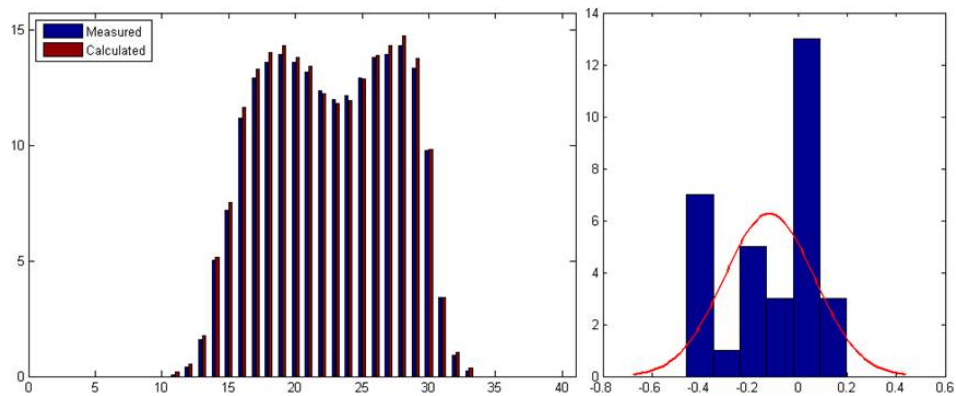
Figure 4.15: The measured and predicted response for an entire fraction are shown on the left, the corresponding histogram of the differences between the two is shown on the right. Although Matlab was able to fit a normal distribution, visual assessment is less than convincing. This is typical of both the beams and fractions for ten treatments.

Global assessment

Although the local percentage difference in out-of-field signals might be large, the absolute difference is quite small. Assessing this in the context of the entire delivery proved a sensible approach. X% of the smallest maximum was used to set a tolerance on the difference between the measured and predicted values for each wire. What is meant by the smallest maximum? Both predicted and measured signal sets have a maximum signal – it is the smaller of these two maxima that is used to define the assessment metric. The reasoning behind this being that, if there was a calibration error with the linac that resulted in a grossly increased output, the higher maximum could potentially be large, thus increasing the size of the assessment metric and, possibly passing the gross error. Using the lower maximum avoids this possibility and makes a tighter tolerance overall.

Position-dependent errors

Both the measured and predicted results need to presented with a level of confidence. Repeat measurements of the same plan on different days on different linacs (three sets of data, two different linacs) showed that the uncertainty on the measured DAVID response to be of the order 1% which is consistent with the results found by Poppe et al. (2006).

For the predicted signals, due to the variability in the penumbral signal (Figure 4.7), the uncertainty associated with these values was much greater. Uncertainties of 1% and 3%, respectively, were given to the in-field and penumbral regions for each segment. These values were summed in quadrature so that, in the case of the final beam, the regions of the field primarily made up of in-field radiation had a smaller uncertainty than those primarily made up of penumbral beam (Figure 4.16)
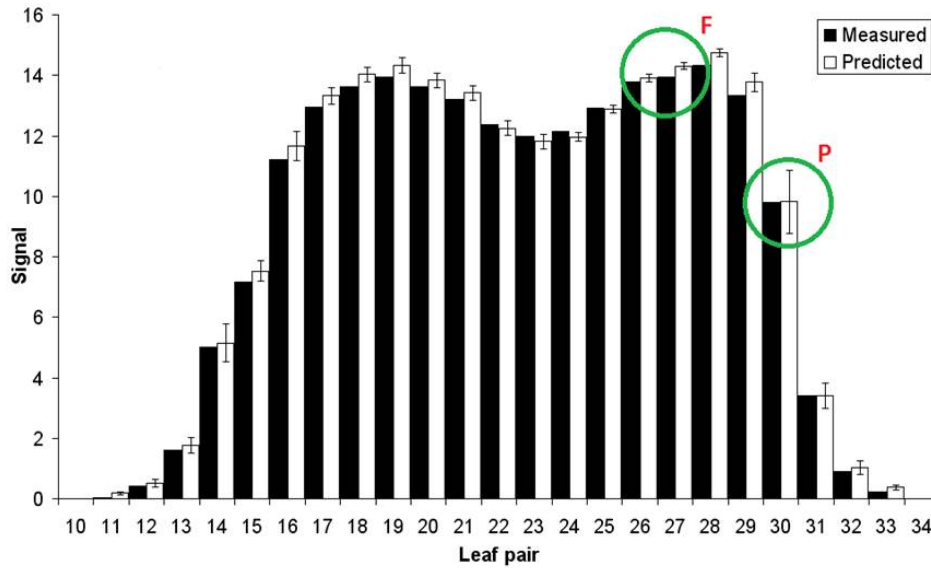
Figure 4.16: Position-dependent uncertainties on a full beam ($\sim$ 10 segments). The error bars in region "F" are lower as the response in the section was largely made up of in-field signals; the response in region "P" was made up of mainly penumbral signals. More penumbral regions will results in higher levels of uncertainty - because each beam is made up of a number of differently-shaped segments, these levels of uncertainty are unlikely to be symmetrical.

4.5 Implementation

The signals of ten, five-field H&N treatments were measured with the DAVID and compared to predicted values using the workflow outlined in Figure 4.17 . Each plan was delivered three times and the average taken. The DAVID had some difficulty resolving individual segments – some low-MU deliveries were often not recorded as distinct segments. These were included in the signal for the next or previous segment. Even if this only happens once, subsequent segments will no longer be associated with the right position in the delivery order, making comparison difficult. With this in mind, comparisons were made between measured and predicted beams (about 15 segments) and whole fraction (5 beams) (Figure 4.18). A tolerance (X% in Section 4.4.1) of 5% of the maximum was used for the beams and 2.5% for the whole fraction (Figure 4.17).

In order to test the sensitivity of the algorithm, a calibration error was simulated by both increasing and decreasing the MU of the imported plans by various percentages. In addition to this, to asses the algorithm's ability to detect more coarse delivery errors, predicted results were compared to the incorrect measured signals for beams and whole deliveries.

4.5.1 Results

The algorithm successfully predicted all the DAVID response for all ten plans, for the whole fraction and for each beam in accordance with the tolerances discussed in Section 4.5.

Where the predicted results were compared to the measured results from a different beam or whole delivery the major difference between the two was obvious
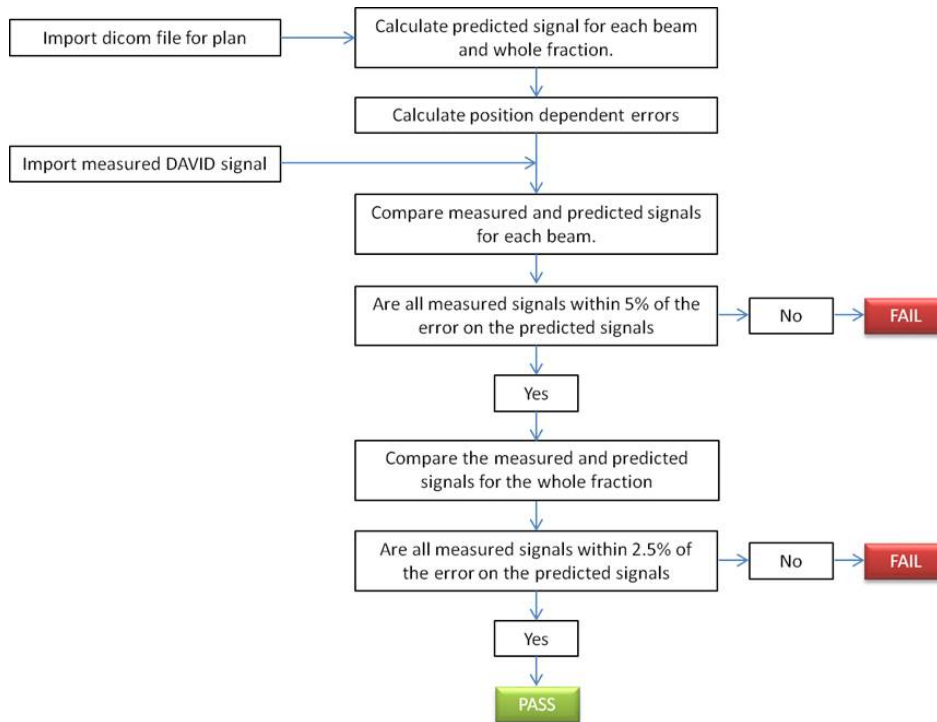
Figure 4.17: Pass/Fail workflow for the comparison algorithm.

and easily detected each time.

The simulated calibration error resulted in all of the plans failing when at both plus and minus 5% errors were introduced.

The algorithm provides a quick way of predicting the DAVID signal and the tolerances and errors chosen are appropriate for detecting gross ($\geq$5%) errors (Figure 4.19).

4.5.2 Application of the analytical model for predicting the DAVID signal and initial conclusions on its use.

Figure 4.20 shows a suggested paradigm for the algorithm's use. It is expected that the implementation of this methodology would catch most major errors that can occur in radiotherapy. The predictive algorithm developed here was intended to remove the pre-treatment base-line step. It was able to predict the measured signals for the whole treatment, and for individual beams. Differences between the predicted and measured signals at each wire were <5% and <2.5% of the maximum for each beam and each fraction respectively. It was also shown to be capable of detecting calibration output errors at the 5% level.

Verification with checking software and signal generation using this algorithm would detect gross errors on the first treatment fraction, without the time cost of pre-treatment verification on the LINAC (Figure 4.20). While having the same tolerance as that typically implemented for IVD diode readings (Section 2.3.1), this method is 2D and does not require accurate, patient-surface positioning, making it much more suitable for IMRT deliveries. The suggested method (Figure 4.20) would detect:

- Errors in treatment planning software – detected by independent checking
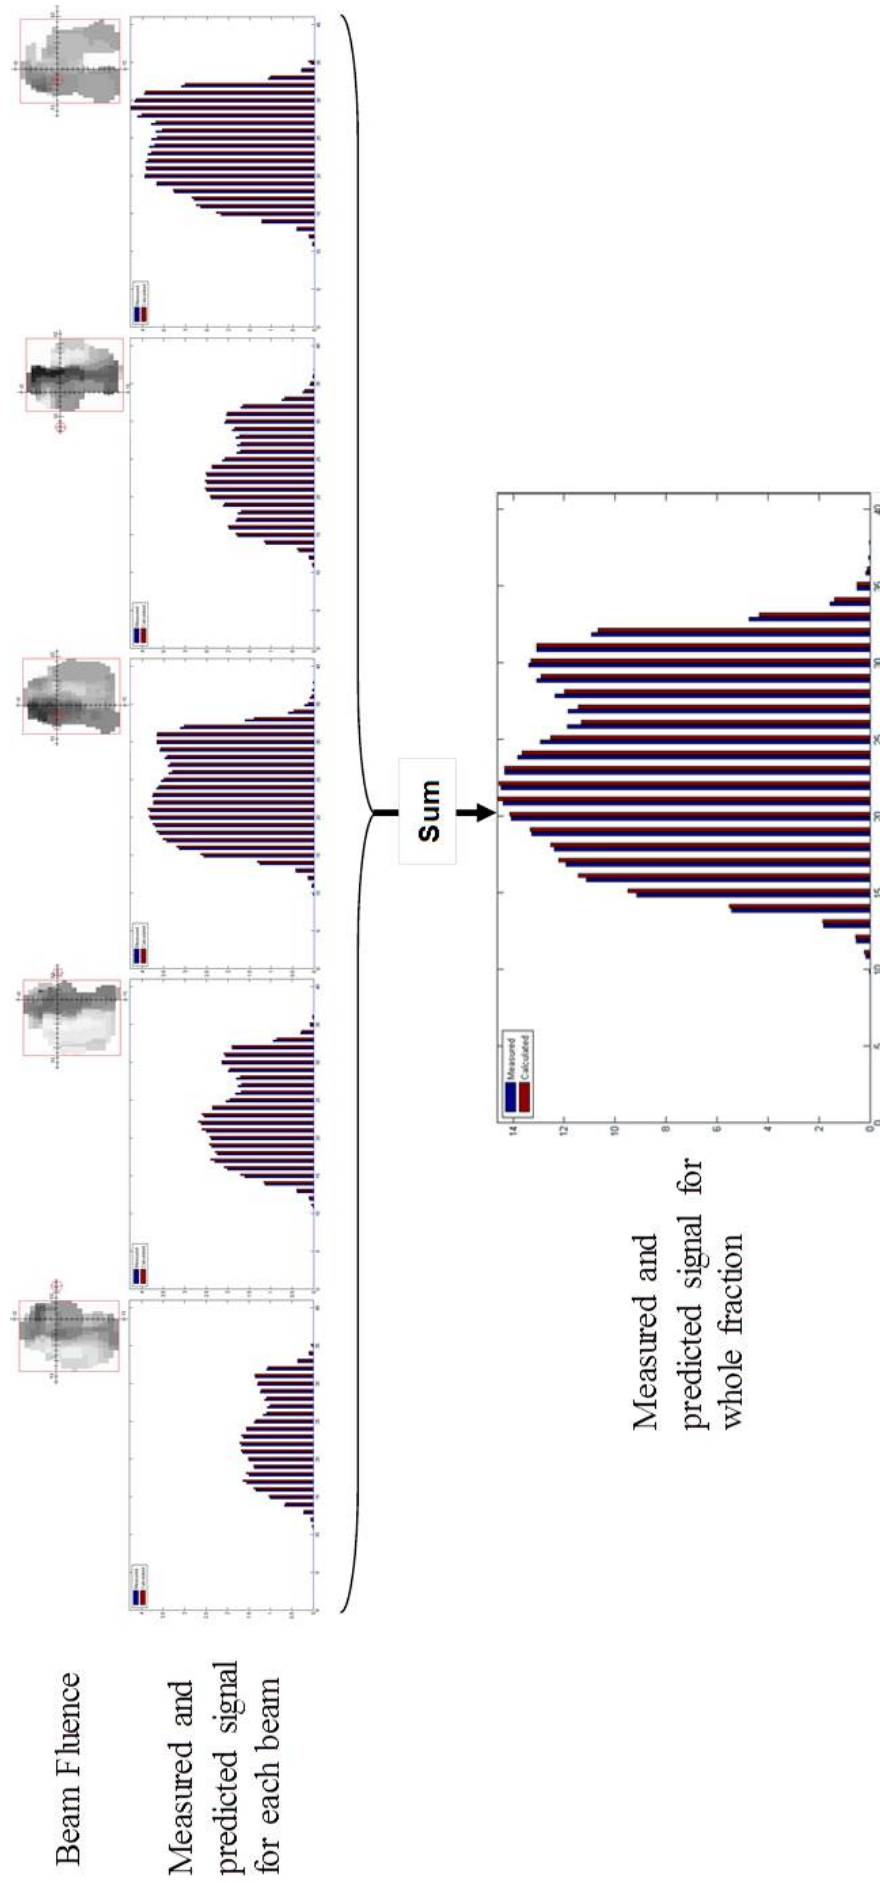
Figure 4.18: Measured and predicted signals for each beam of a five-field H&N IMRT delivery (individual beam fluences are also displayed) each beam is made up of several (about 15) segments and there are 5 beams per treatment fraction.
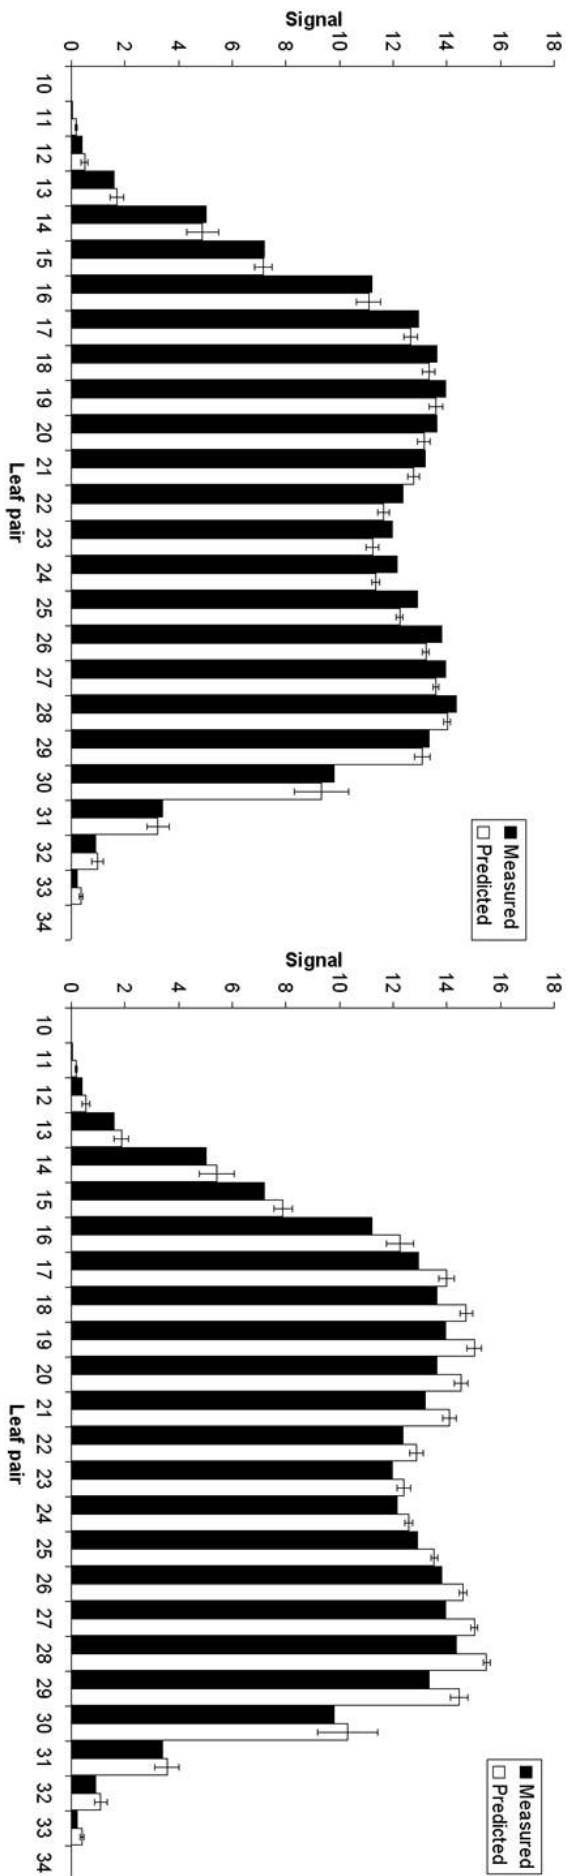
Figure 4.19: The left and the right images show the measured and predicted signals for a plan where the MU are modified by -5% and + 5%, respectively.

74
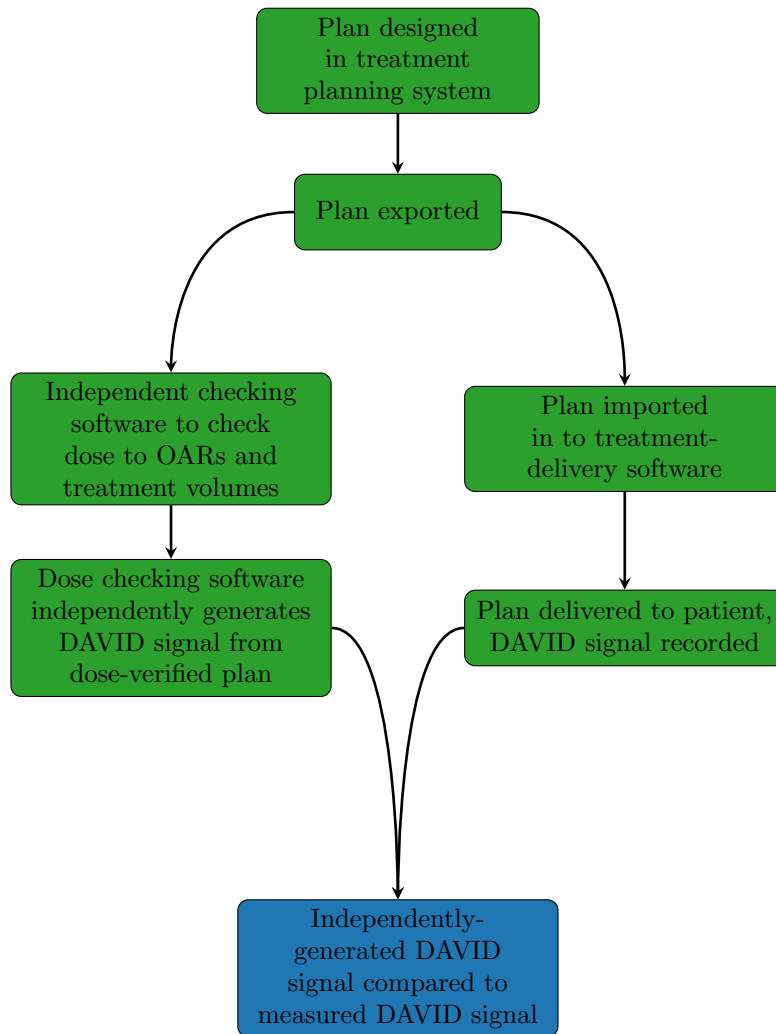
Figure 4.20: Suggested paradigm

software

- Errors in DICOM plan export – detected by independent checking software

- Errors in plan transfer, or upload to treatment delivery software (e.g. the New York incident discussed in Section 2.2.2) – detected by the comparison of independently-generated and measured DAVID signals. This can be done in real time allowing the operator to stop the treatment after the first IMRT beam or, in the case of VMAT, early on in the first fraction.

- Machine errors – detected by the comparison of independently-generated and measured DAVID signals

4.6 Leaf pair defined segment shape and position dependence - a possible limitation?

The analytical model has been shown to work with acceptable accuracy for typical clinical IMRT beams; Section 4.5 shows how it can be implemented in clinical practice. There is a potential limitation for complex segment shapes. The algorithm, in its current state pays no attention to the shape of the specific segments

and the relative position of the leaf gaps. The signal is calculated by the convolution of a 1d matrix of leaf separations with the lateral response function. The position of the leaves relative to each other is important as it dictates the amount of field that is penumbral and, as a consequence, has a lower fluence. A 10×10 cm square will only have four penumbral regions – one at each side; ten 1×10 cm strips with alternating 5/-5 offsets will have many penumbral regions resulting in a lower overall fluence. The effect of this on the DAVID signal was assessed by recording the signal from two wires for instances where the two wires were covered with a 1×5 cm strip; in each case however the two strips had a different relative offset (0 to 5 cm). The results with their respective segment shapes are shown in Figure 4.21; the most heavily offset segments result in the lowest signals.
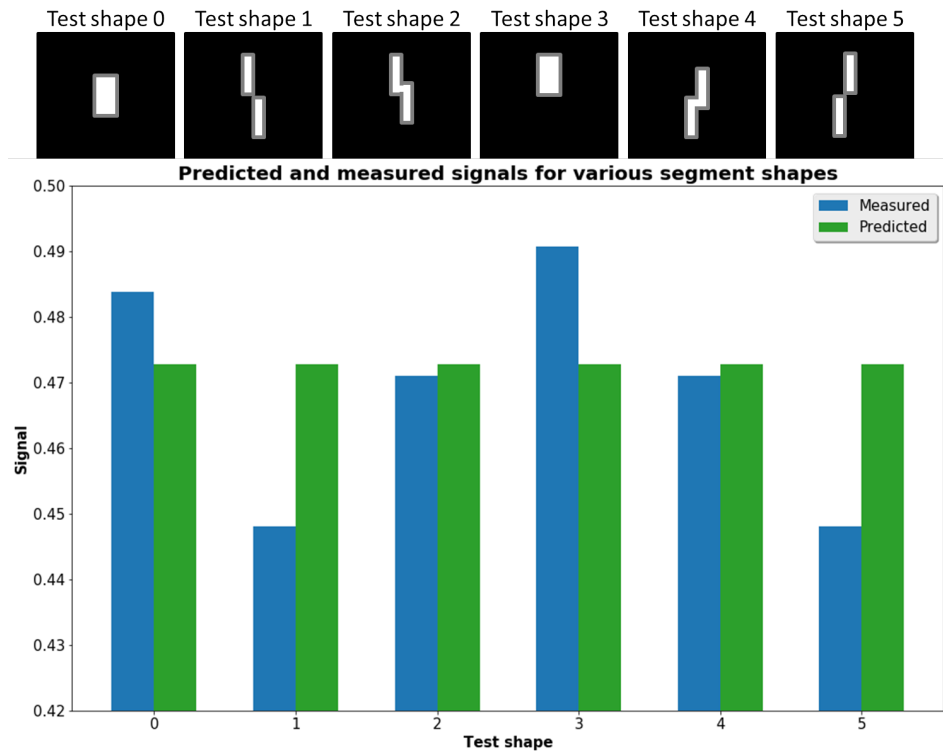
Figure 4.21: Segment shapes with their associated DAVID response for a 20MU delivery along side the predicted response. In each case the predicted response is the same, this is because the leaf separation matrix used in the algorithm will be the same in each case. The variation in the measured response is as a consequence of the varying amount of penumbra in each of the test fields.

4.7 Overall summary and discussion of the need for more sophisticated modelling approaches.

The analytical algorithm developed and tested in this chapter has only 6 variables (Table 4.1). Nevertheless is has been shown to be able to import a DICOM treatment plan and generate a predicted DAVID signal for complex treatments. The accuracy of the prediction, with the uncertainties discussed, has been demonstrated to be sufficient to detect gross errors, for example: delivering the wrong beam or other dose delivery errors greater than 5%. Section 4.5 provides a method

to implement an error detection workflow using the algorithm in clinical practice, without any need to have pre-treatment verification measurements. The analytical model and error-detection workflow has been shown to perform with acceptable accuracy for typical IMRT beam, although Section 4.6 discussed a potential limitation if the arrangement of the leaves relative to each other is particularly complex. Should the DAVID have been used as outlined in Figure 4.20 then the severity of treatment errors discussed in Section 2.2.1 could have been reduced, or maybe even averted altogether. The algorithm has been published as "A simple model for predicting the signal for a head-mounted transmission chamber system, allowing IMRT in-vivo dosimetry without pre-treatment linac time" (Johnson et al. 2014) and can be implemented or used for gross error detection, which is seen by many as the main purpose of IVD (Section 2.2.3). Maintaining dosimetric confidence of 5% is the minimum needed to ensure clinical objectives are met (Brahme 1984). However, a dosimetric certainty of 3% is desirable (Brahme 1984). The algorithm in this chapter lacks the complexity to predict the signal to the extent that the tolerance on the agreement between measured and predicted signals could detect 3% errors in machine output. This is largely down to the necessary simplifying assumptions of a limited-parameter analytical model (Figure 4.21) and the uncertainties associated with predicting the out-of-field response (Figure 4.13).

A 2D fluence map contains more information about the delivered field than an array of leaf separations. Subtleties, like the ones shown in Figure 4.21, will be included provided the software generating the fluence is suitably complex. The signal recorded at a wire in the DAVID device is proportional to the energy deposited in the collection volume. The energy deposited in a volume is proportional to the radiation passing through it and the stopping power of the medium. A radiation fluence map across the DAVID area will provide a relative energy distribution map as the medium is uniformly air. Summing the fluence along the collection volume of a specific wire and applying a scaling factor will give a prediction of the the measured signal of that wire (Figure 4.22).

The MU-checking software DIAMOND (PTW, Freiburg) was commissioned and used to generate fluence maps of IMRT deliveries, the details of this process are outlined in Appendix B. However, it was evident that specific and accurate penumbral modelling was not a feature of the algorithm. Also with the out-of-field area having the fluence set to zero, it was clear that scatter, leaf leakage and inter-leaf leakage were not considered in the calculation. As a consequence the fluence maps lacked the accuracy necessary to be used to generate a predicted signal through the method outlined in Figure 4.22. Attempts were made to modify the fluence maps and the interaction with the DAVID in an approach similar to that discussed in Greer et al. (2009), but these were unsuccessful. The idea of generating a 2D map that could be used to generate the DAVID signal was not abandoned, though it was clear that a more thorough approach was needed. Monte Carlo modelling is recognised as the most thorough and accurate way of simulating linear accelerators and their output (Verhaegen & Seuntjens 2003, Seco & Verhaegen 2016, AAPM 2007). Monte Carlo algorithms can be used to generate a dose volume output. Multiplying dose by mass yields energy deposited in a medium which relates directly to the DAVID response. The subject of the next chapter is the use of two Monte Carlo techniques to simulate IMRT deliveries

through the DAVID so that the energy deposited in the collecting volumes of the wires can be determined and used to predict the signal. It is expected that this much more sophisticated modelling approach will produce much more accurate results than those discussed in this chapter and in Appendix B where fluence-based model lacks detail.
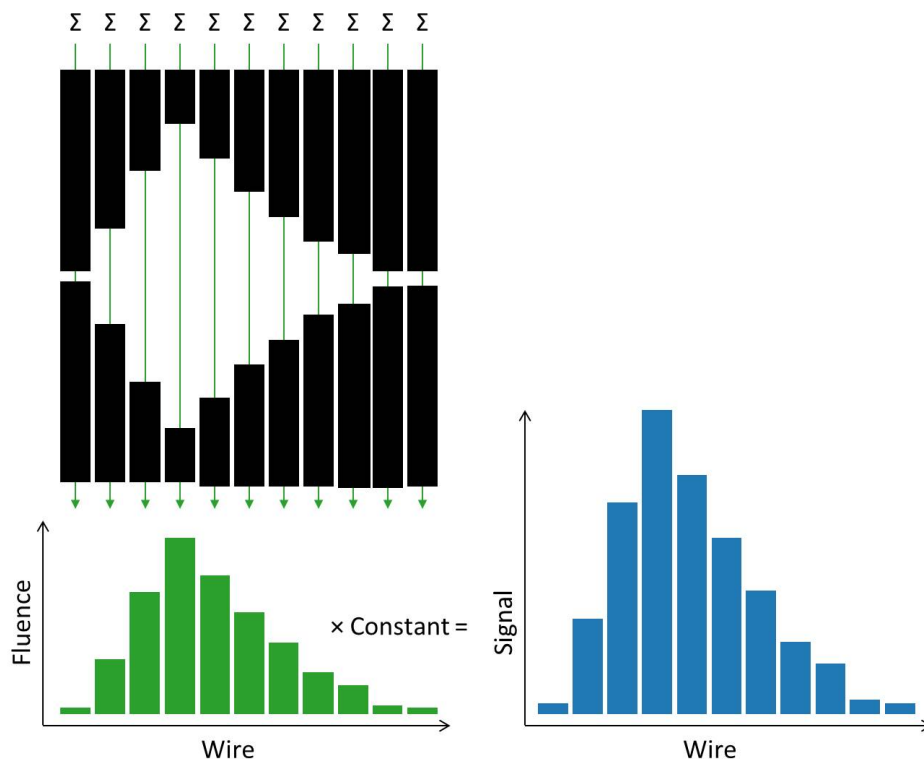
Figure 4.22: The fluence (white) as a result of the MLCs (black) for an IMRT segment is shown in the top left. Summing the fluence along the collection volume of the DAVID wires will give the total fluence the volume is exposed to. Applying a correction factor to this will give a prediction the measured signal.

Chapter 5

Monte Carlo Modelling of the DAVID

5.1 Introduction

This chapter is concerned with generating two Monte Carlo models of the DAVID. One will be generated in Monaco, a commercially available planning system, the other using EGSnrc. The EGSnrc approach discussed in this chapter is a full MC model of the linac and DAVID, this thorough approach produces the most accurate results but at the cost of time in both developing the model and running individual simulations. The dose calculation in Monaco makes a number of approximations, this speeds up the calculation considerably, but at the expense of some accuracy. Assuming that a beam model is already in place, as would be the case in a clinical department, implementing the Monaco approach to MC modelling of the DAVID is much easier in terms of both setting up the model and calculating the signal. It is expected that using Monte Carlo models will enable accurate predictions of the DAVID signal, and therefore avoid the need to perform pre-treatment measurement on a linac before validating an in-vivo patient measurement. Comparing the two models is expected to enable evaluation of the performance of the simpler approach against the fuller EGSnrc approach to consider the Monaco approach's accuracy and possible limitations.

5.2 Monte Carlo Modelling

Generally the Monte Carlo (MC) method can be described as a numerical approach to solving problems based on random number sampling (Seco & Verhaegen 2016). Take, for example, the value of π. If we take a circle of radius r and a square of side length $2r$ then the ratio of the surface areas will equal $\frac{\pi}{4}$. Let $2r$ be set to equal 1 and both shapes drawn so they share a common centre at 0.5,0.5 (Figure 5.1). A random number generator that produces numbers between 0 and 1 can be used to determine an x and y coordinate, this will define a point that will fall either inside the area defined by the circle (C) or inside the area defined by the square (S). The likelihood of a point of falling inside C or S is proportional to the area of the circle and square respectively. As the number of points increases, the ratio of the points inside C to those in S tends toward the ratio of the areas.

If enough points are used it is possible to use the ratio to determine an accurate value of π (Figure 5.1).

This example also demonstrates the statistical nature of MC calculations. The value of π is not determined analytically, but estimated through a ratio that changes with N. As N increases, the uncertainty associated with the value reduces. This uncertainty, associated with any MC calculation, is unavoidable and later on in this chapter manifests itself as noise. It can be shown (AAPM 2007, Feller 2008) that the statistical uncertainty is proportional to $1/\sqrt{N}$ meaning that statistical noise can be reduced by either increasing N (in the case of EGSnrc N corresponds to number of histories) or, in the case of voxel-based calculations, increasing the voxel size so that $N\,per\,voxel$ increases.

$$4 \times \frac{nC}{nS} = 2.6$$

N = 9
$\pi_{estimate}$ = -20%

$$4 \times \frac{nC}{nS} = 3.5$$

N = 32
$\pi_{estimate}$ = +11%

$$4 \times \frac{nC}{nS} = 3.125$$

N = 64
$\pi_{estimate}$ = -0.5%
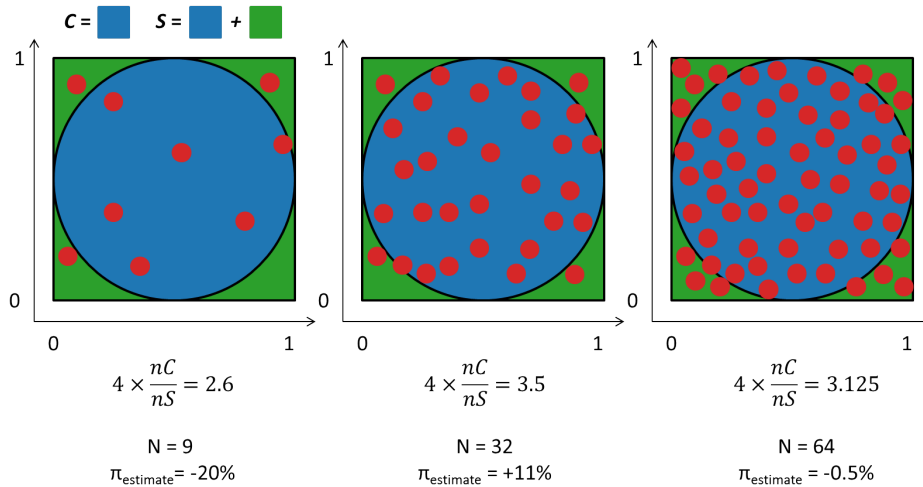
Figure 5.1: As the number of randomly-generated points increases, the estimated value of π becomes more accurate.

Monte Carlo simulations in medical physics attempt to model nature directly through the essential dynamics of the system. MC calculations offer a numerical solution to a problem that can be described as the temporal evolution of particles based upon their cross sections. Mimicking nature, the interactions are processed randomly and repeatedly until the numerical results converge on useful means. In this way MC models offer a solution to a macroscopic problem through simulation of its microscopic interactions (Seco & Verhaegen 2016).

EGSnrc introduction

There are a number of general-purpose MC codes used throughout medical physics. PENELOPE (Salvat et al. 2006), MCNP (Briesmeister et al. 2000), GENAT4 (Agostinelli et al. 2003) and EGSnrc (Kawrakow et al. 2000) are the four most popular with EGSnrc being the most widely used (Rogers 2006). All have various strengths and weaknesses; however, EGSnrc was used for this work for a number of reasons:

- Most widely used general-purpose code in radiotherapy physics.

- EGSnrc distribution comes with BEAMnrc - a package designed specifically to model therapeutic linacs.

- Of the 4 aforementioned codes, the use of the PRESTA algorithm (Bielajew & Rogers 1986) for particle boundary crossing in EGSnrc makes it the most accurate for ion-chamber calculations (Seuntjens et al. 2002).

- Local experience and knowledge of the software.

For the EGSnrc modelling in this chapter, BEAMnrc was used to produce phase-space files in the plane perpendicular to the beam axis at some distance outside the linac head (Figure 5.2). From here, the simulation was taken over by DOSXYZnrc. DOSXYZnrc is, like BEAMnrc, an extension of the EGSnrc code and is included in the software distribution. It can use the phase-space file generated by BEAMnrc to simulate the interaction of photons and electrons in a volume allowing it to score energy deposition in user-specified voxels (Walters & Rogers 2003). Section 5.3 is concerned with creating an accurate model of the
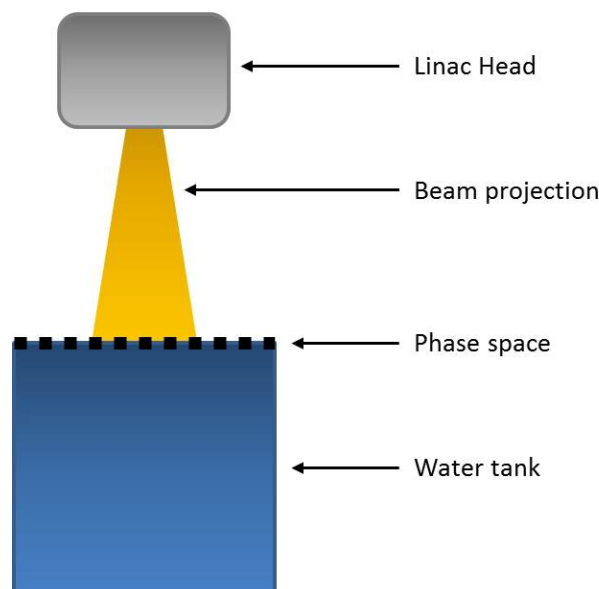
Figure 5.2: Position of phase space for water tank simulations

Elekta MLCi2 head in using BEAMnrc. Various iterations of different electron spectra (Section 5.3.1) and beam-shaping geometry (Section 5.3.2) were simulated and used to create phase-space files. DOSXYZnrc uses these phase-space files as a particle source for a water tank simulation (Fgiure 5.2). The accuracy of the BEAMnrc model of the linac was benchmarked by comparing simulations with water tank measurements made on the physical linac.

Having shown that the BEAMnrc linac model accurately simulates the actual MLCi2's output, Section 5.4.1 uses DOSXYZnrc to model the DAVID with and without the collection wires, assessing the wires' impact to determine if it is necessary to include them in the model. This is done in a similar way to Asuni et al. (2011), who assesses the electron contamination from a transmission detector. In Section 5.4.2 - DOSXYZnrc is used to make a model of the DAVID. Finally, in Section 5.5, the DOSXYZnrc model is validated by generating BEAMnrc phase-space files of clinical VMAT deliveries and using these to predict signal in the DAVID. These are then compared with equivalent signals collected by the DAVID on the linac.

All EGS-based simulations were run on the Linux-based supercomputer - Medical Advanced Research Computer 1 (MARC1). MARC1 is part of the High Performance Computing and Leeds Institute for Data Analytics (LIDA) facilities at the University of Leeds, UK. This is a large memory node cluster made up of 57 blades each housing one Intel Haswell node. Each node is dual socket with a 10-core Intel E5-2660v3(2.6GHz) processor per socket (20 cores per node); 256GB of DDR4 2133MHz memory per node (configured as 16 x 16Gb); a 500Gb hard drive and QDR Connect-X Infiniband. In total, the 57 blades contain 114 CPUs and 1140 cores. All user traffic is carried over the InfiniBand network (Leeds University 2018).

Monaco Introduction

"The Monte Carlo technique is potentially the most accurate method for the calculation of dose distributions in treatment planning if radiation sources and patients are completely modelled and a sufficiently large number of photon and electron histories are simulated"(Fippel 1999, AAPM 2007). At the time of writing, the only commercially-available MC-based TPS in widespread use was Monaco. Monaco incorporates the X-ray voxel Monte Carlo (XVMC) algorithm but with a number of approximations and simplifications to dramatically speed up calculation time. When benchmarked against EGSnrc calculations, no significant loss in accuracy has been observed (Kawrakow et al. 1996, Fippel 1999, Semenenko et al. 2008, Alber & Nüsslin 1999). At present, there is no published literature on modelling a transmission detector in Monaco. However, a TPS-based model of the DAVID would be much easier to use and would generally be quicker to calculate than one in EGSnrc. This would allow easier implementation pre-delivery baseline generation, as it is likely that more people in a radiotherapy department will be familiar with the TPS than the EGSnrc code. Section 5.4.4 discusses how the model is designed in Monaco, and it is compared with the EGSnrc-based model in Section 5.5.

5.3   Modelling the MLCi2 head using EGSnrc

The Electron-Gamma-Shower (EGS) software is designed to perform Monte Carlo simulations of equipment used in radiotherapy and calculate the dose that these simulations produce in user-defined media. It can simulate the coupled transport of electrons and photons through a wide range of geometries and materials for energies from a few keV to hundreds of GeV (Kawrakow et al. 2000). The EGSnrc code is widely used for radiotherapy simulations because it has been shown to be accurate, furthermore, its geometries, materials and energy range are tailored specifically to radiotherapy (Rogers 2006).

BEAMnrc is an extension of the EGSnrc code, specifically tailored to simulate radiation produced by therapeutic linear accelerators. The software includes component modules (CMs) that lend themselves to creating the geometry and parts integral to a hospital accelerator. In addition to this, a number of variance reductions tools are included in the software, such as: range rejection, photon forcing, Bremsstrahlung splitting, Russian Roulette, directional source biasing and Bremsstrahlung Cross Section Enhancement. All of these are calculation options

aimed at reducing simulation time with minimal, or unimportant, losses in accuracy. A more in-depth discussion of these options is included in Appendix C. A linac model, built in BEAMnrc, will consist of a series of component modules, each dealing with a specific geometry of the beam-delivery system. These are independent of one another, have a defined material and density and occupy a horizontal band in space. A BEAMnrc-generated graphic of the primary filter and flattening filter can be seen in Figure 5.4. In both cases the CM FLATFIL was used. FLATFIL allows the user to define cones of different width, thickness and material. In Figure 5.3 FLATFIL is used to make:

1. A single cone, with the bottom radius set to r_1, the top radius set to 0 and the height set to h_1

2. Another single cone, but instead of going to a point, a trapezium shape is defined by using the same settings as 1, but setting the top radius as r_2 instead of 0.

3. Two materials are used. The bottom cone, as before, has a radius of r_1, but now has length h_2 with a top radius of r_3; the other cone starts at height h_2 with radius r_3 extending h_3 to a radius of r_4. The two cones, as seen in this instance, can be defined as being of different materials.

4. Geometry 4 in Figure 5.3 shows that with suitable manipulation of the widths of cones, different materials (shown here as different colours) can be made to sit inside each other.
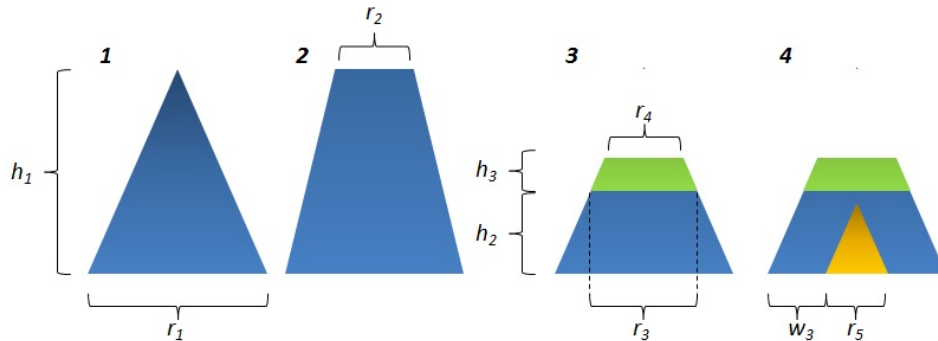
Figure 5.3: Possible geometries using the FLATFIL component module

By exploiting the versatility of the FLATFIL CM, it was possible to create a variety of shapes that have very different functions in the model. At the time of writing there are, in total, 25 CMs included in the EGSnrc distribution. A full description of each is outside the scope of this work, but between them, they can be used, with suitable manipulation (like that seen for FLATFIL in Figure 5.3) to define the physical components of a linac in the MC model. The horizontal banding of the component modules can be seen in Figure 5.5; this shows all the CMs in the static part of the linac model, labelled with their defined material (Rogers et al. 1995).

The principal output of the BEAMnrc code is a phase-space file. This contains the position, direction, energy, charge and history tag (Section 5.4.1) of all the particles in a user-defined plane. This file can be used as a source for simulations

/ calculations by various programs (Rogers et al. 1995) any number of times. This can reduce the need for repeat simulations of components of the beam-transport system prior to encountering a component of interest or a dynamic component, saving time simulating the entire linac for each calculation.

At the time of writing there were no published models of the MLCi2 treatment head using the EGSnrc code (See Section 5.7). The linac model built in the BEAMnrc software for this thesis contained 10 CMs. Table 5.1 is a brief description of each:

Module Name	Component Module used	Description
Target	FLATFILT	Copper rings surrounding a tungsten core
Primary collimator	FLATFILT	Conical air space flanked by tungsten shielding. See Figure 5.4
Flattening filter	FLATFILT	6, stacked, steel trapezoids
Monitor chamber	FLATFILT	A complicated 21-layer arrangement using Mylar, Aluminium and air to define Elekta's monitor and output chamber (Figure 5.32)
Backscatter plate	SLABS	SLABS are a simple CM that allow the user to define the material and relative position of numerous rectangular blocks. In this case 1 aluminium slab 2.9mm thick was used.
Mirror	MIRROR	The MIRROR CM is similar to SLABS but has the extra option to change the angle of the slab relative to the beam axis; the SLABS CM keeps this fixed at 90.
MLC	SYNCMLC	The SYNC modules allow the easy delivery of multiple-segment plans (IMRT, VMAT and DMLC) (Lobo & Popescu 2010). The 40 MLC tungsten leaf pairs were defined as having rounded tips with a radius of 15.02 cm and length in the beam-directions of 8.2cm.
Jaws	SYNCJAWS	Have the same SYNC capabilities as the leaves and are composed of tungsten.
Mylar	SLABS	A single, thin mylar sheet.

Air	SLABS	Air slab where the phase space is saved. The length of this slab was modified depending on the data being collected. For tank-based simulations, it extended to the water's surface (Figure 5.2); for DAVID work, it extended to the DAVID surface.

Table 5.1: The names and descriptions of the CMs used in the EGSnrc model of the linac

The SYNCMLC and SYNCJAW CMs (Lobo & Popescu 2010) used in the MC model both require a sequence file that defines the position of the jaws and leaves in the head. Software was written that took the calibration curves in Section 5.3, read the leaf / jaw position at isocentre, defined in a GUI, and converted it to sequence files for the SYNCMLC and SYNCJAW modules. In the case of a single, static delivery (Figures 5.11 to 5.15), the weight of the segment would be one. The software was developed further so that it could read a DICOM plan, extract the leaf positions from each segment, convert them to the position in the head and write them to a sequence file that could be read by the SYNC modules. This feature was used for the VMAT plans in Section 5.5. Each segment in the sequence contained the positions of the leaves / jaws as well as the fraction of the delivery that was delivered through that segment (segment MU / total MU). The SYNC modules use the MU weight to define a cumulative probability distribution function (CPDF) for all the possible leaf positions. For each initial particle history a random number is generated and used to sample the leaf positions from the sequence file with a probability determined by the CPDF (Liu et al. 2001). For VMAT simulations, when the random number falls between two segment indices, the leaf positions are determined through linear interpolation between the two adjoining segments; for step-and-shoot deliveries, only even-numbered segments are sampled, and there is no interpolation (Heath & Seuntjens 2003). SYNC modules not only have the advantage of allowing an entire VMAT delivery to be done in one go (previously it was necessary to model each individual segment and sum them at the end), it is also more accurate, as it models the delivery between the defined segments (Liu et al. 2001).

The geometry, materials and position of each component module can be set to match the parameters of the linac to be modelled. In this instance, Elekta provided a machine specification to work from. Graphics of the flattening filter and target are shown in Figure 5.4; Figure 5.5 shows the static part of the linac model.

Having compiled all the CMs to match the specification that Elekta provided the model was tested by comparing the output with results measured in a plotting tank. Where necessary the parameters in the model were modified so that it matched the true linac output (Chibani et al. 2011). The model output was generated by creating a phase space using BEAMnrc then releasing it through a dose-scoring medium using DOSXYZnrc. A phase space file contains all the particles generated by the accelerator as well as their position, energy and direction of travel. Once generated, the same phase space files can be used any number of times for any analysis they might be suitable for. Throughout the testing described
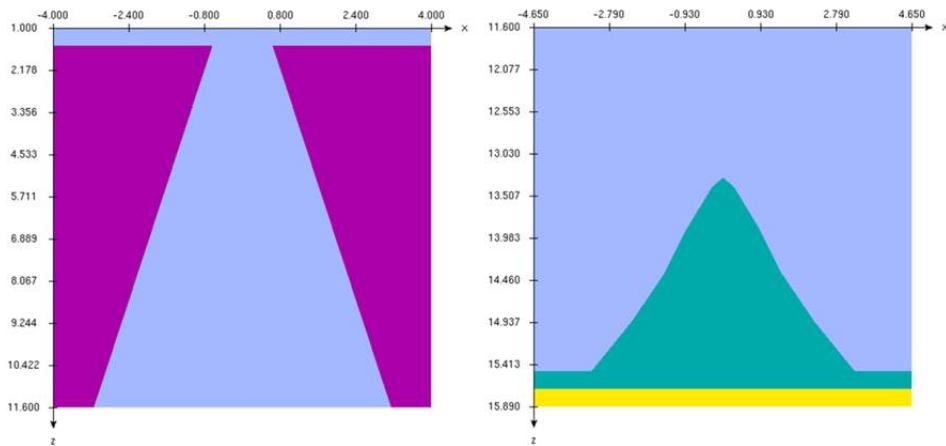
Figure 5.4: Primary collimator (left) and Flattening Filter (right) as modelled by BEAMnrc

in this chapter, the phase space file was positioned directly after the mylar film at the end of the linac head. DOSXYZnrc was then used to score the dose of the field in a water tank simulation. Figure 5.6 shows the DOSXYZnrc arrangements used to measure the simulated PDDs and profiles.

The simulated profiles were generated by obtaining phase-space files for 20×20, 15×15, 10×10, 5×5, and 3×3cm fields from the BEAMnrc model. 30×10^9 histories were used for each delivery, Directional Bremsstrahlung Splitting (DBS) was used as a variance reduction technique (see Appendix C) and the 700ICRU PEGS data set was used for particle interactions (Figure 5.5). PEGS is from the original application for EGSnrc which was in high-energy physics. PEGS4 (Preprocessor for EGS) was designed to generate photon cross sections for this purpose. For radiotherapy applications, the Storm & Israel (1970) data set is used by default (and throughout this work) for cross-section data, however PEGS4 is still used for material information, photon threshold energies and some of the electron data. The EGSnrc distribution includes the material information to provide all the cross section data for two data sets: 700icru.pegs4dat and 512icru.pegs4dat (Rogers et al. 2001). The numbers 512 and 700 refer to the energy down to which the cross sections are calculated, 0.512 and 0.700 MeV respectively. For both the DOSXYZnrc and BEAMnrc simulations in most of this work, unless stated otherwise, the 700 data set was used, because it gave acceptable accuracy and better calculation times.

The phase space files for the five field sizes (20×20 to 3×3 cm^2) were then used as sources for DOSXYZnrc simulations of a water tank at 90SSD. For the DOSXYZnrc simulations, ISMOOTH (see Appendix C) was turned on, the 700ICRU PEGS data was used for particle interactions and 1.5×10^9 histories were used for each data set. Profiles were extracted from the DOSXYZnrc .3ddose files using statdose (included in the EGSnrc distribution package McGowan & Faddegon (2013)). A 10×10cm phase space was also used as a source for DOSXYZnrc to generate the simulated PDDs in Figures 5.8 to 5.10. The same input criteria and extraction techniques were used as in the profile data. The DOSXYZnrc water phantom was simply modified to measure a PDD (Figure 5.6). For all the simulations, the ECUT and PCUT were set at 0.7 and 0.01 MeV, respectively (see
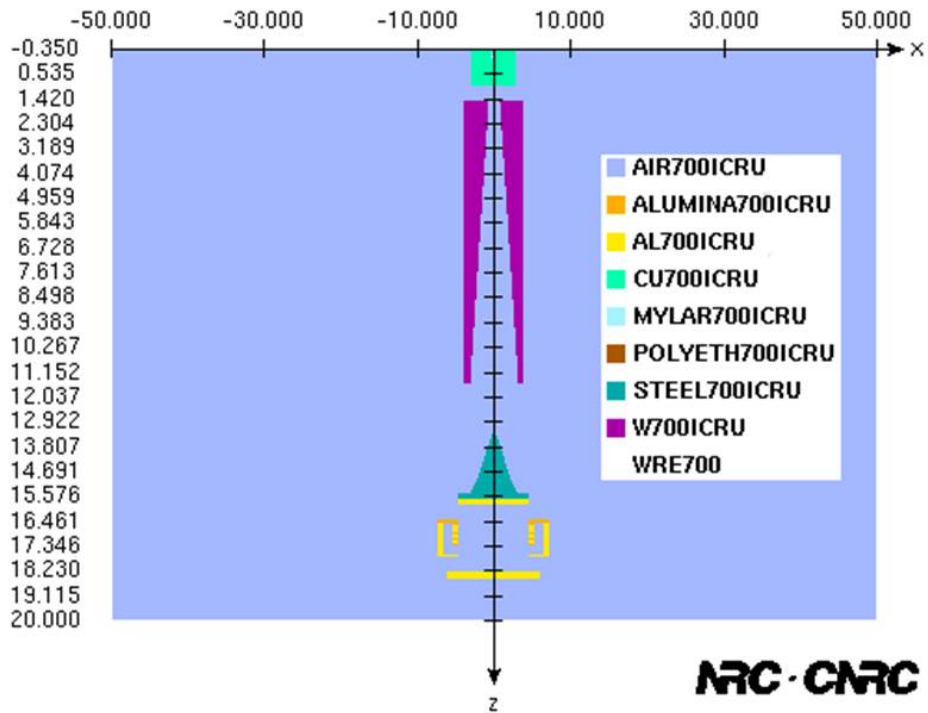
Figure 5.5: Static part of the linac model as modelled by BEAMnrc. Materials are defined by the 700icru.pegs4 data set. The cross-section data in this file is based on the density corrections contained in the ICRU report 37 (ICRU (1985)) and uses cross-section data down to 0.7 and 0.01 MeV for electrons and photons, respectively. Both have upper energy limits of 55MeV (Rogers et al. (2001), Kawrakow et al. (2000)).

Appendix C.2.1).

The water tank data was collected by employing similar techniques described in Appendix B.2.1; the MP3 PTW water tank and accompanying MePhysto software were used with semiflex 31010-type chambers as both field and reference detectors (PTW Freiburg) (Figure B.3). For the profiles, the semiflex field chamber was orientated in the vertical position, 2mm steps were used for the central part of the beam and 1mm steps were used in the penumbra. For the PDDs, the semiflex chamber was in the horizontal orientation and 1mm steps were used throughout.
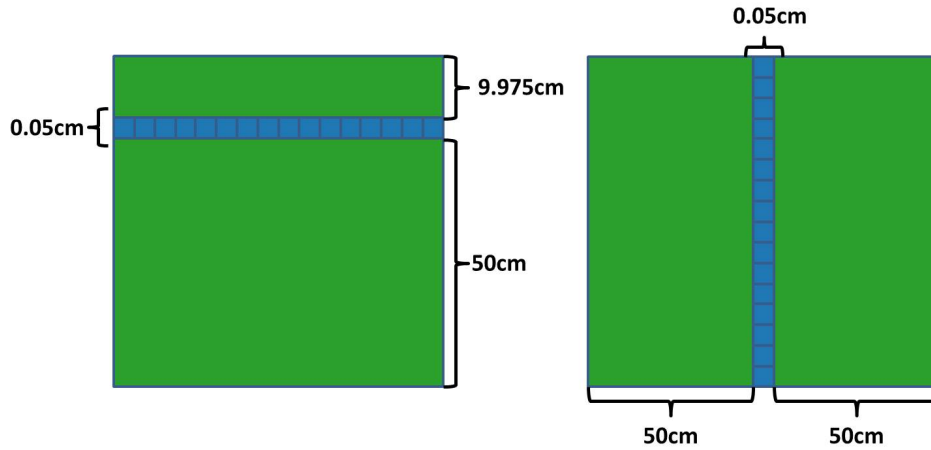
Figure 5.6: The voxel arrangement used to collect profiles (left) and PDDs (right). In the actual simulation there were 250, 0.05cm2 voxels where the dose was scored

5.3.1 Model Spectra and PDD

Modelling approach

To ensure that the dose deposition is simulated correctly, the beam energy in the model must closely match that of the physical linac. The beam energy is dictated by the energy spectra of the electrons hitting the target. In a physical linac, this is controlled by a number of components and the values of their associated control setting parameters – electron gun, magnetron bending magnets etc. The BEAMnrc simulation uses a file that contains various electron energies and the relative abundance of these energies. It has been shown that the central-axis depth-dose curve is very sensitive to on-target mean electron and energy distribution (Sheikh-Bagheri & Rogers 2002*b*, AAPM 2007). Consequently, matching simulated and measured PDDs implies equivalent on-target electron beams.

Results

Electron spectra are normally described as a Gaussian distribution centred on the beam energy (Khan 2010). Several Gaussians were tried (Figure 5.7), but none were found to create a beam that matched the measured data; instead the spectrum shown in Figure 5.10 was found to work best. Figures 5.8 to 5.10 show PDDs resulting from mono energetic electron spectrum, Gaussian spectrum and the final empirical spectrum plotted against the machine-measured PDDs. Figure 5.11 shows the measured and simulated PDDs for 5, 10 and 20cm^2 fields using the final electron spectrum shown in Figure 5.10. Gamma index evaluation (Appendix A) was used to compare simulated and measured PDDs, 100% of points passed with the 2D gamma index set to 1% / 1mm. Due to the good match between measured and simulated PDDs in Figure 5.11, the energy spectrum that generated them (Figure 5.10) was used for all subsequent simulations. It may be noted that the noise in the MC simulated PDDs is statistical noise, as discussed in Section 5.2.
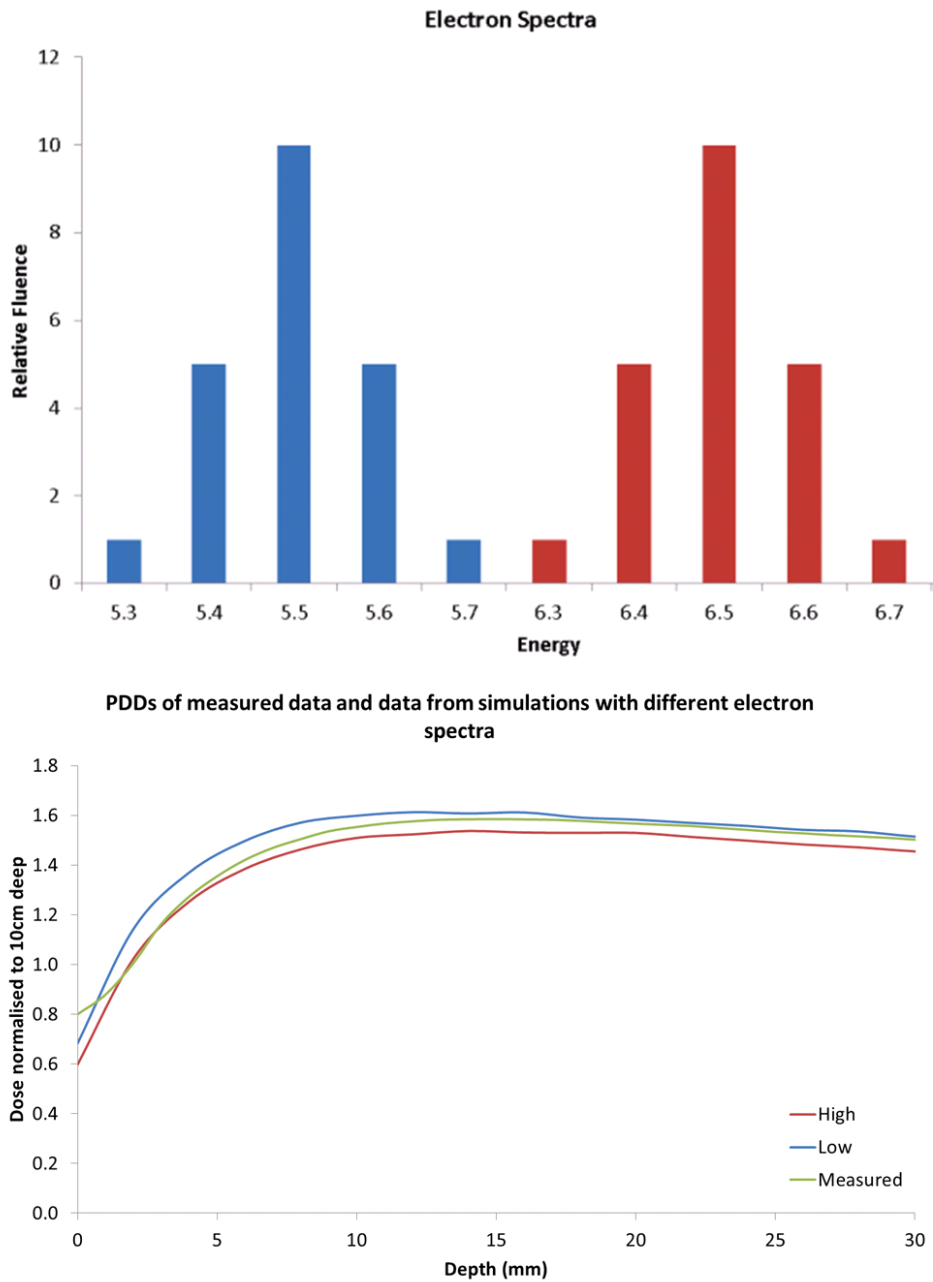
Figure 5.7: Simulated PDDs for low and high energy electron spectra with measured data; the colour of the PDD (bottom) corresponds to the colour of the spectrum (top) that was used to simulate it, with the measured data displayed in green. While it was possible to match the simulated data with measured data in the build-up region when a higher energy was used, and in the deeper regions when a lower energy was used; it was not possible to get a PDD match at all depths using a Gaussian electron energy distribution. All PDDs taken for a 10×10 field at 90SSD.
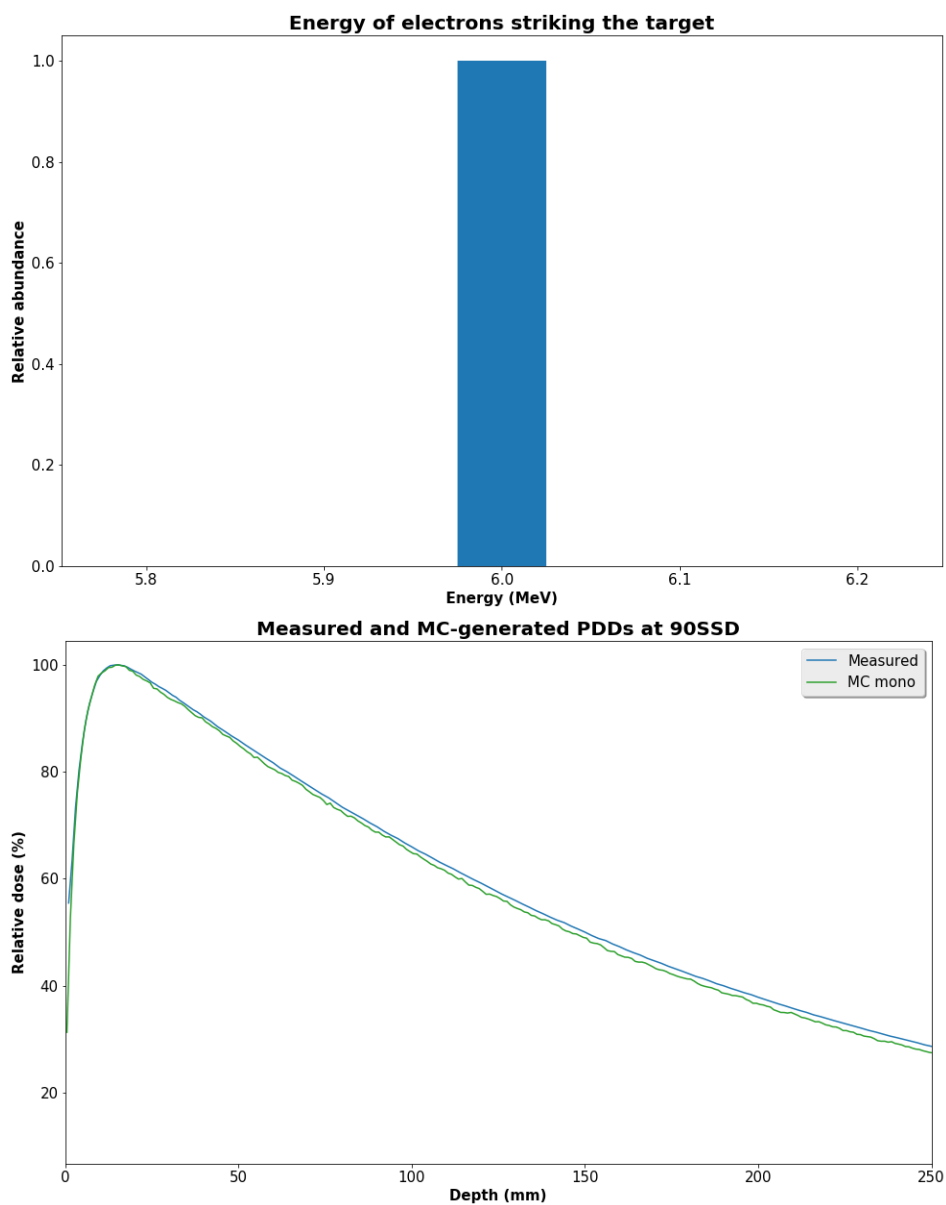
Figure 5.8: Monte-Carlo modelled PDD for a 10×10 field at 90SSD for a mono-energetic electron beam (top) incident on the target plotted with data measured using a plotting tank.
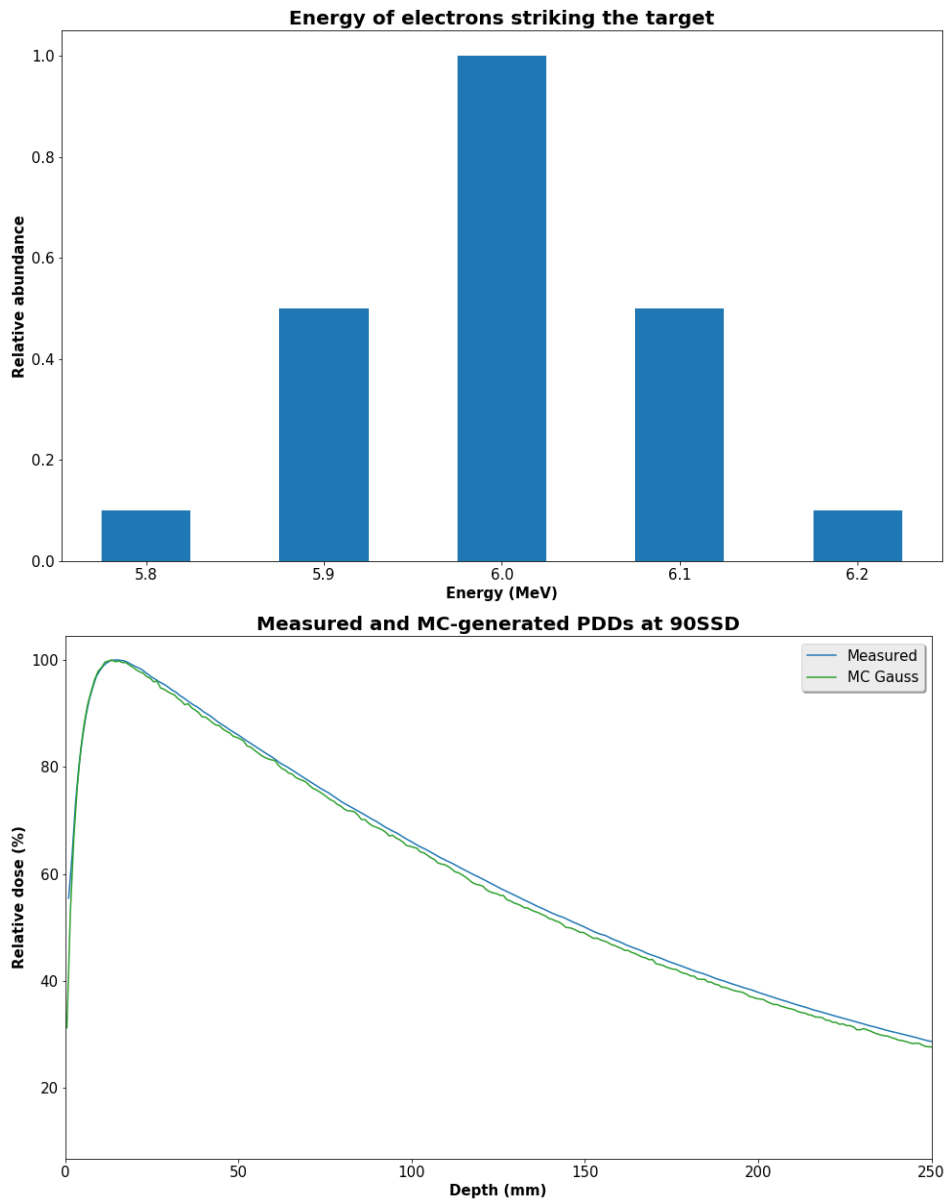
Figure 5.9: Monte-Carlo modelled PDD for a 10×10 field at 90SSD for a Gaussian electron beam (top) incident on the target plotted with data measured using a plotting tank.
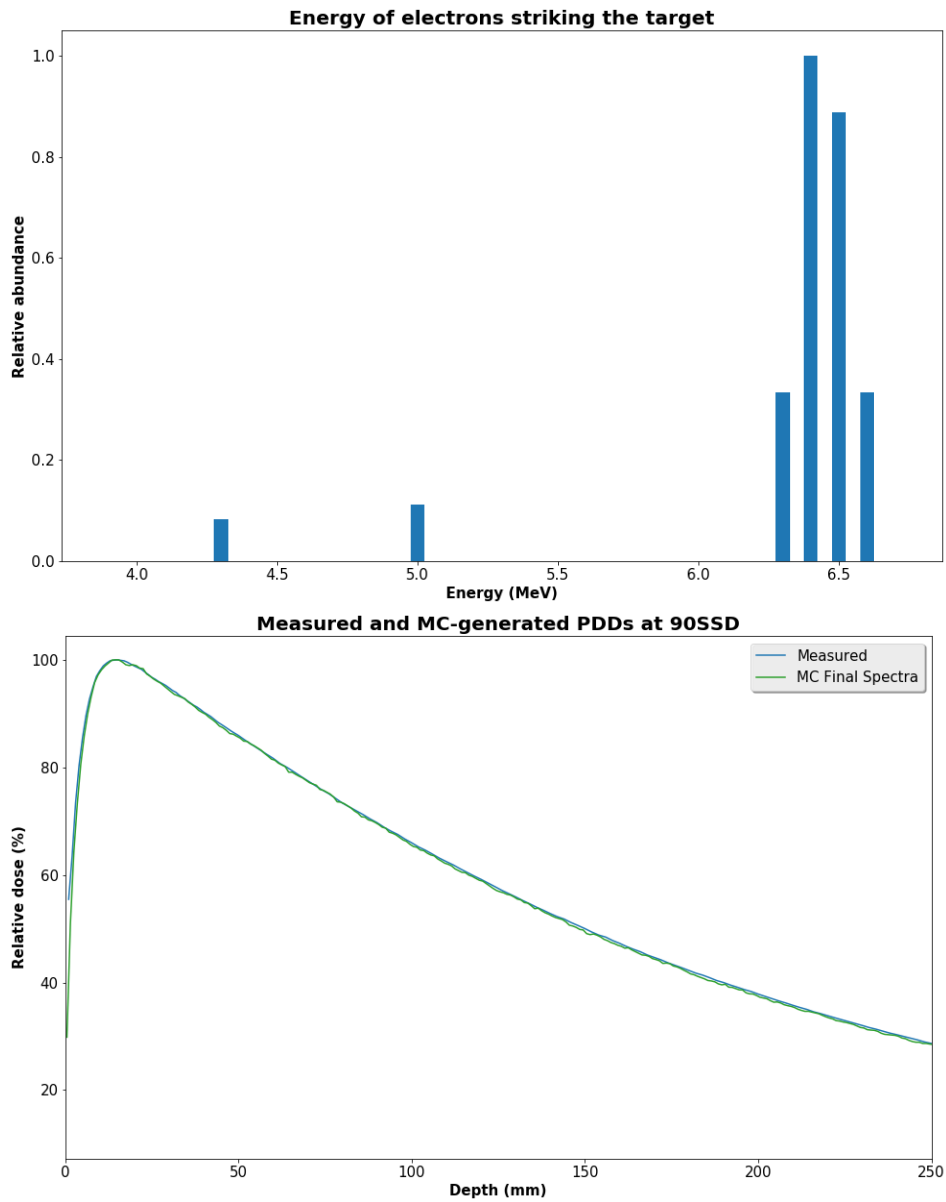
Figure 5.10: Monte-Carlo modelled PDD for a 10×10 field at 90SSD for the final empirically-selected electron beam (top) incident on the target plotted with data measured using a plotting tank.
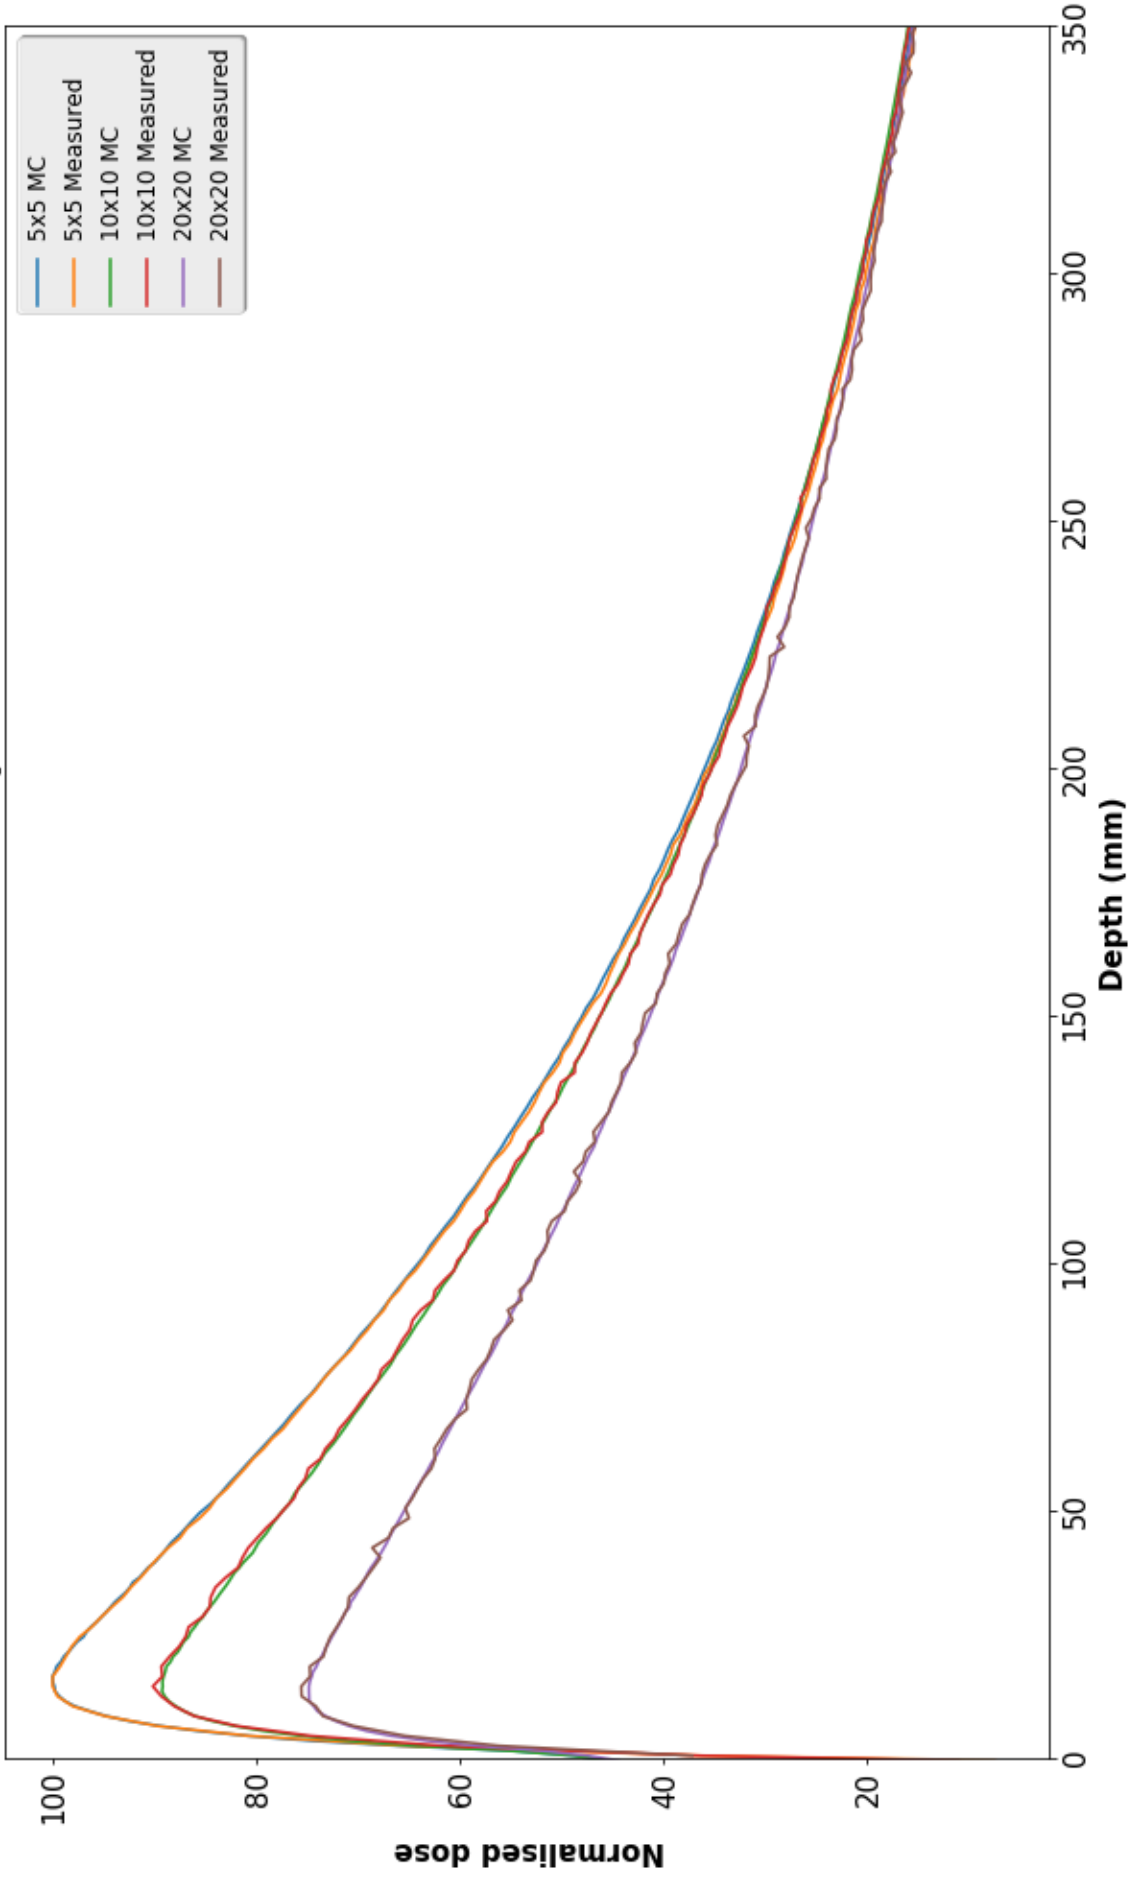
Figure 5.11: Comparison of measured and simulated PDDs taken at 100cm SSD for 5,10 and 20cm square field sizes. Agreement = 100% of points passing 2D 1% / 1mm gamma index criteria.

5.3.2 Modelling leaf and jaw positions and Profiles

Modelling approach

The position of the leaves and jaws in the linac model required to create a 3×3, 10×10 and 20×20cm^2 square fields were initially determined using geometry and the supplied schematic. Profiles generated from these field shapes were then compared with measured profiles and minor adjustments made to the model's leaf and jaw positions so that the profiles matched. The position of the leaf bank and jaws in the head were plotted as a function of the isocentric field size for these three field sizes. The relationship was characterised by a quadratic for both the leaves and jaws. These scaling equations were then used to determine the leaf and jaw positions for 5×5 and 15×15cm^2 square fields on the central axis as well as 5×5 and 3×3cm^2 square fields displaced by 10cm on the X and Y axis.

Results

The X and Y axis profiles from the model and measured results are shown in Figures 5.12 to 5.15. The X axis profiles were defined by the leaves with the Y axis profiles being defined by the jaws. All points were found to agree within the 1.5% / 1mm 2D gamma index acceptance criteria. The quadratics defined by this process were used to determine the leaf and jaw positions for all subsequent simulations.
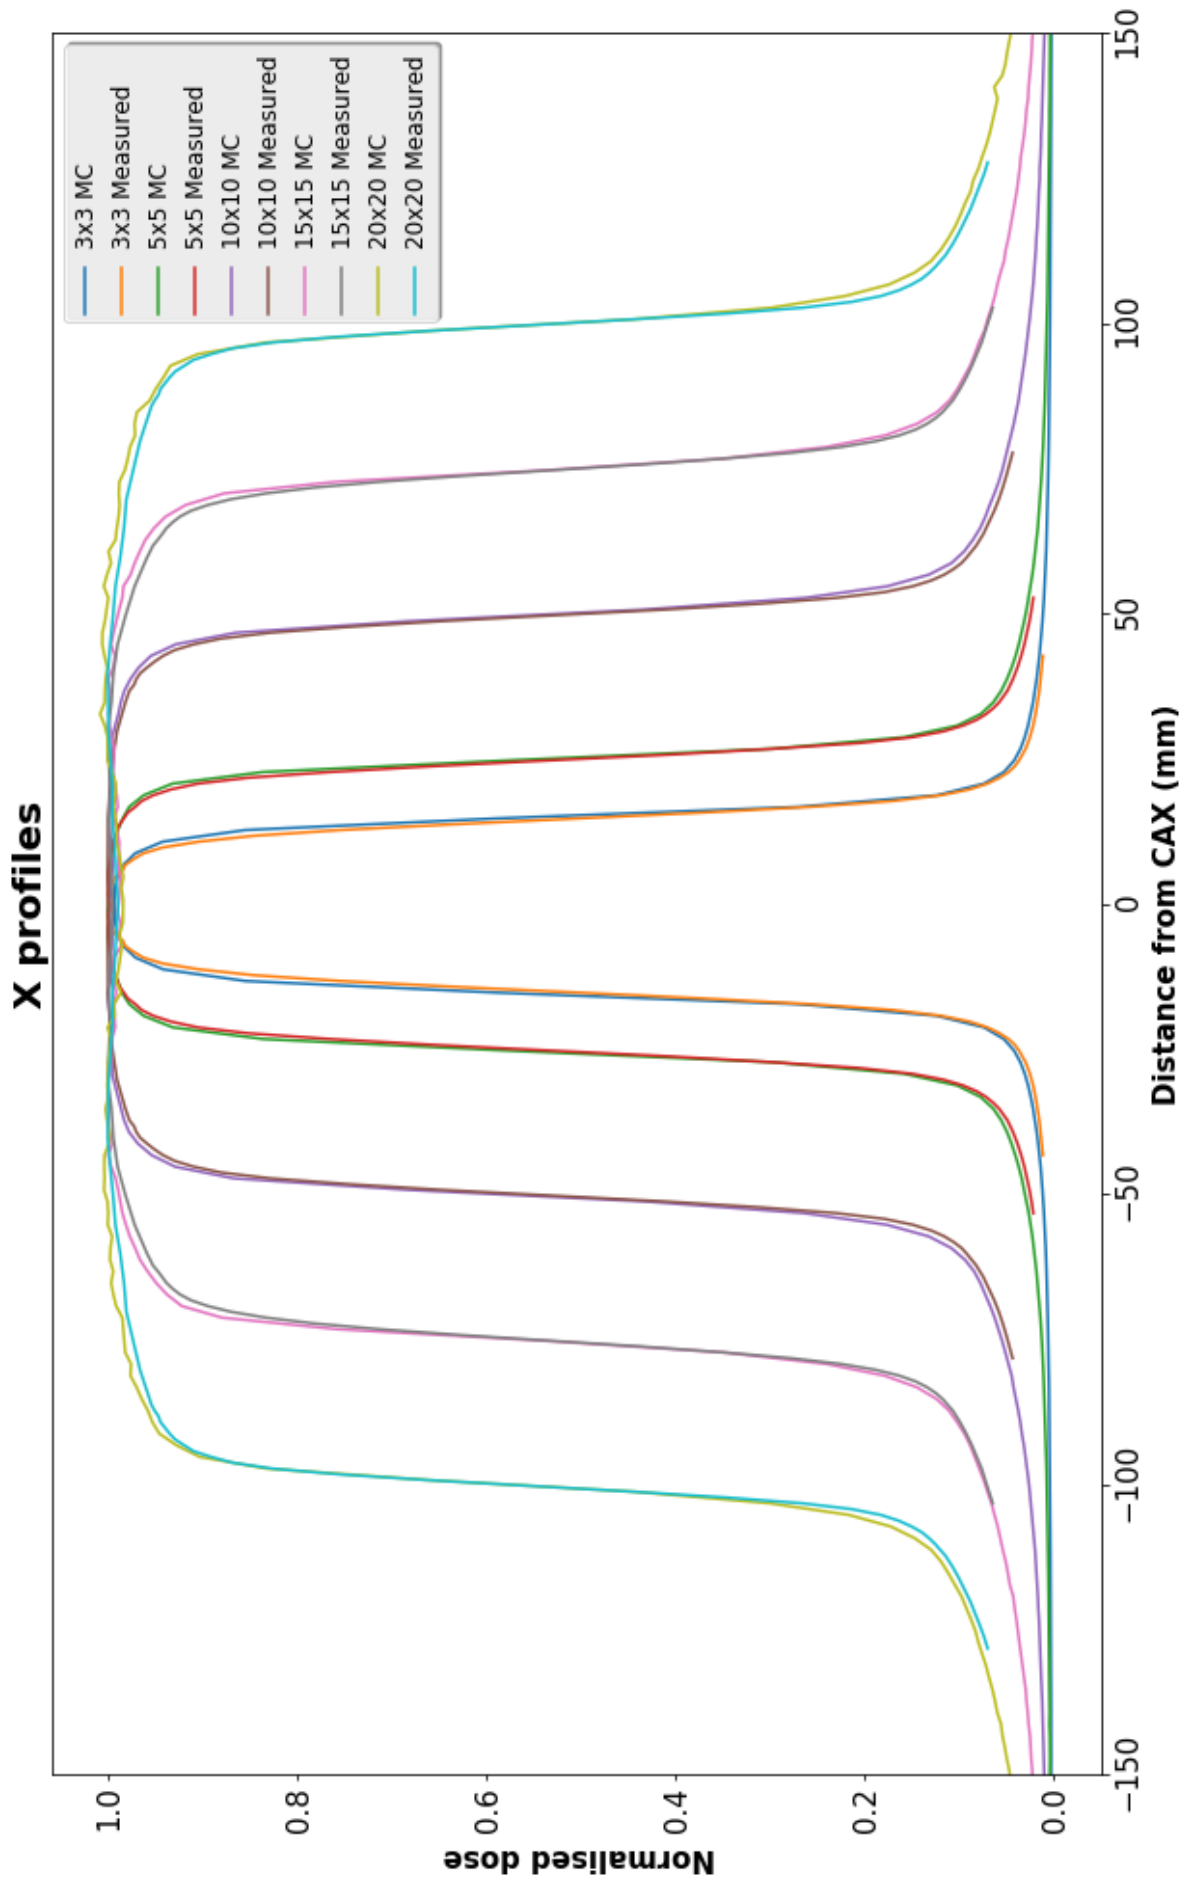
Figure 5.12: Comparison of measured and simulated X axis profiles, defined by the leaves. Agreement = 100% of points passing 2D 1.5% / 1.5mm gamma index criteria.
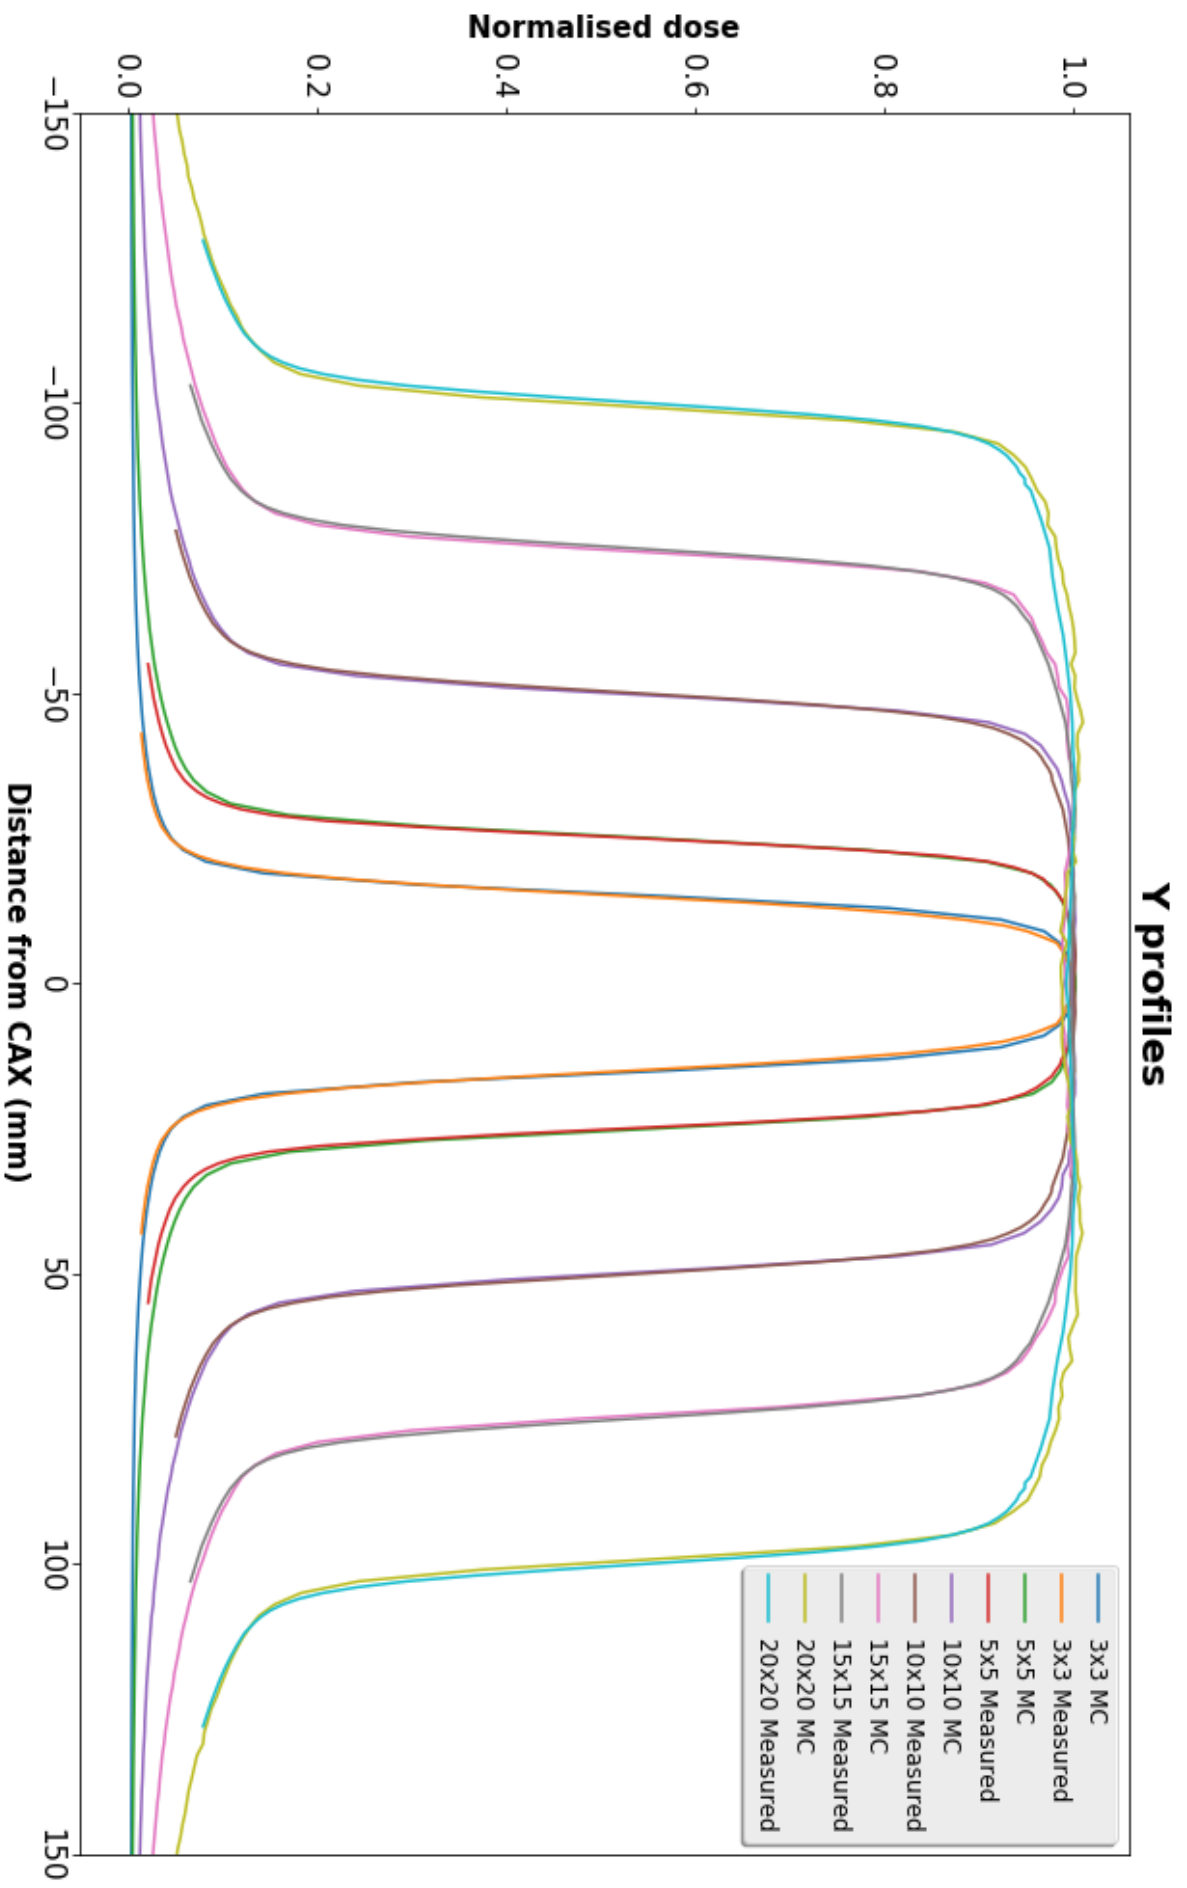
Figure 5.13: Comparison of measured and simulated Y axis profiles, defined by the jaws. Agreement = 100% of points passing 2D 1.5% / 1.5mm gamma index criteria.
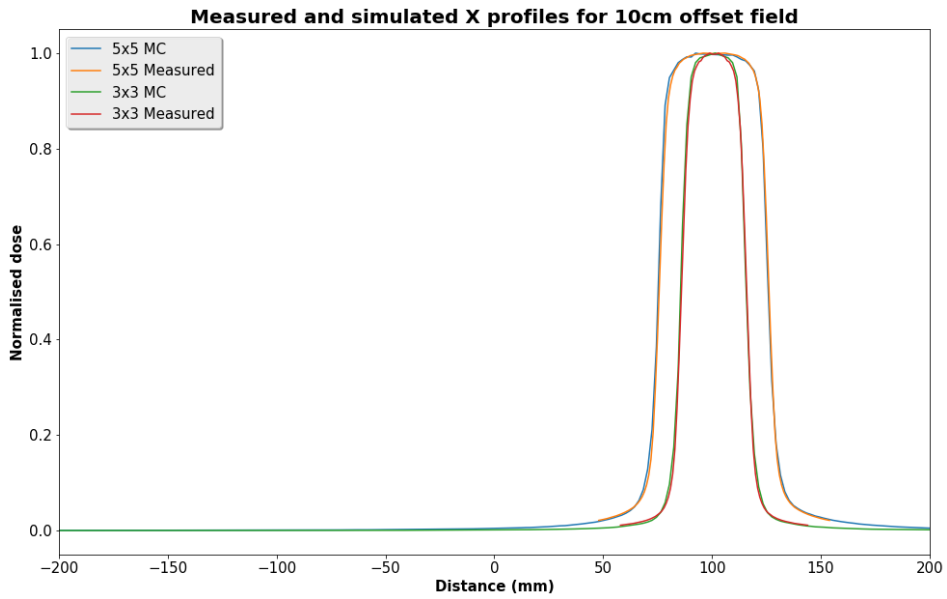
Figure 5.14: Comparison of measured and simulated X profiles for a 5 and 3cm field offset from the beam axis by 10cm in both X and Y directions. Agreement = 100% of points passing 2D 1.5% / 1.5mm gamma index criteria.
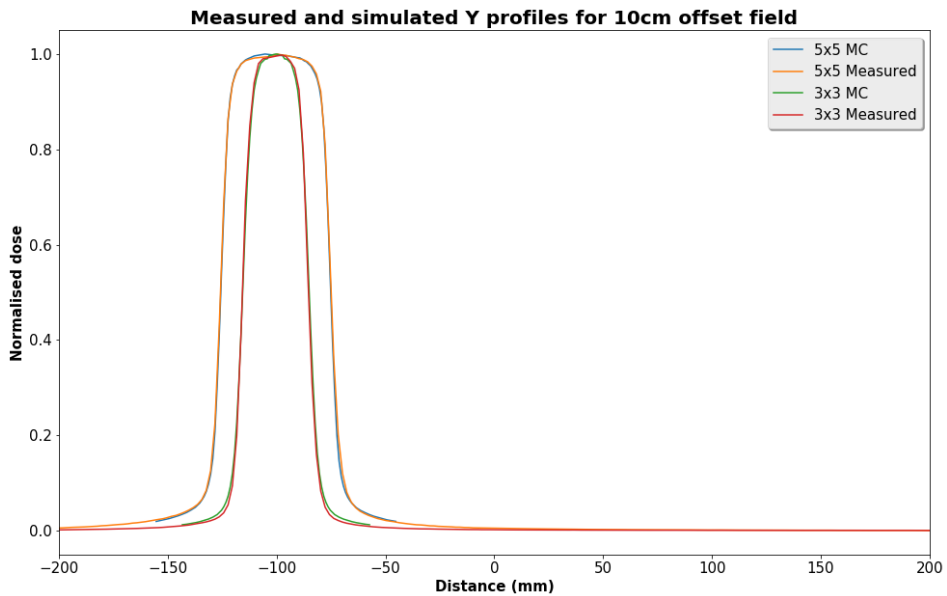
Figure 5.15: Comparison of measured and simulated Y profiles for a 5 and 3cm field offset from the beam axis by 10cm in both X and Y directions. Agreement = 100% of points passing 2D 1.5% / 1.5mm gamma index criteria.

5.3.3 Discussion of the linac Monte Carlo Model

Through a combination of vendor-supplied schematics and iterative processes, a Monte Carlo model of the MLCi2 linac head was made. The model-generated PDDs and profiles matched the data collected on the linac using the water tank within the gamma index criteria of 1%/1mm and 1.5% /1mm, respectively. As this match exceeds what might be deemed necessary to match a clinical treatment planning system to a linac (Venselaar et al. 2001, Fraass et al. 1998), the model can be assumed to be good. The linac model, and leaf /jaw position calibrations discussed in this section will be used for all subsequent EGSnrc simulations.

5.4   Monte Carlo modelling the DAVID

5.4.1   Assessing the impact of the collection wires

The DAVID collection wires have a diameter of 100 micro meters (Poppe et al. (2006)). Due to their small cross-section, including them in the model of the DAVID in either EGSnrc or Monaco would be difficult. This section is concerned with using the EGSnrc software to assess the impact of the collection wires on the signal and determine if their inclusion on subsequent models is necessary.

To determine the effect the collection wires had on the energy deposited in the collection volume, the DAVID was modelled in DOSXYZnrc with a single collection and wire volume (Figure 5.16). It is stated in Poppe et al. (2006) that the cross-sectional area of the circular collection volume is 0.03cm^2, this corresponds to a diameter of 0.2cm. Circular collection volumes are difficult to model, so the collection volume was modelled as a square of side length 0.2cm. Modelling the volume as a square instead of a circle will increase the charge collected, however a relative effect is being investigated, so provided the volumes are the same in both wire and non-wire cases, using a square instead of a circle should not matter. The
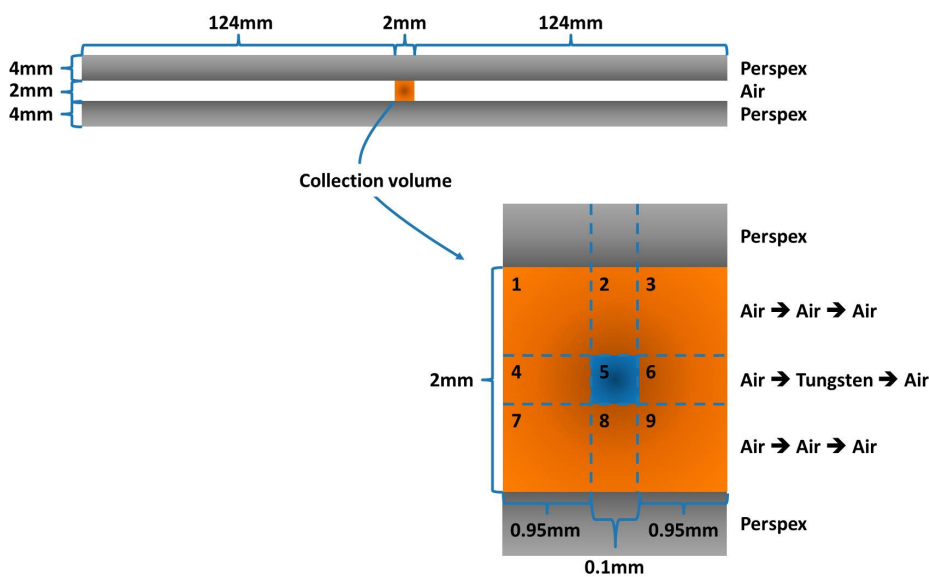
Figure 5.16: DOSXYZnrc model of the DAVID with a single collection wire and volume.

signal measured in the collection volume shown in Figure 5.16 was calculated by summing the energy deposited in regions 1-9. When modelling the volume without the wire, the material in regions 1-9 was defined as air. When modelling the wire, the material in region 5 was defined as tungsten and the dose deposited in the region was not included in the signal, similar to the approach taken in previous studies of ionisation chambers with central electrodes (Ma & Nahum 1993, Wulff et al. 2008, Buckley & Rogers 2006).

BEAMnrc was used with the model described in Section 5.3 to generate phase space files for three square fields of isocentric side-length 5, 10 and 20cm. For each phase space file, 1.5×10^9 histories were used, and DBS (Appendix C.2.4) was implemented as a variance reduction technique.

Two simulations were run in DOSXYZnrc, using the model of the DAVID shown in Figure 5.16, for each of the three square-field phase-space sources: one with the wire, one without, six in total.

For the profile and PDD simulations discussed previously in this chapter (Sections 5.3.1 and 5.3.2) $\sim 10^9$ histories were used. In this case (Figure 5.16) the volumes were considerably smaller. Consequently, it was necessary to run 1.5×10^{11} histories to ensure that there were enough particles to reduce uncertainty of the dose deposited in the collection volume to $< 0.1\%$.

As discussed previously (Section 5.3), the 700 PEGS data set has been used for most of the EGSnrc modelling in this work. However for the work in this section, modelling the small volumes of the wires and collection volumes around them, the 512 PEGS data set has been used. The number 512 refers to the energy down to which the cross sections are calculated, i.e. it allowed lower ECUT and PCUT values to be implemented than the 700 set, 0.512 and 0.01MeV respectively. In the case of DOSXYZnrc calculations this avoided the over deposition of energy in small voxels adjacent to large voxels. This is thought to arise from tracks of charged particles having been transported through the large voxel that happen to end in the small voxel (Walters 2016). Additionally, for both BEAMnrc and DOSXYZnrc calculation, the decision to use lower PCUT and ECUT values also meant that the particles were tracked for longer making the position of the dose deposition more accurate. This all came at the expense of increased calculation time.

The differences between the energy deposited with and without the wire present can be seen in Figure 5.17. For each field size, the difference in energy deposition for the case when the wire was present and when it was not was found to be 3%.

Origin of the increased signal

It is proposed that the increase in signal in the wire simulations was a result of more secondary particles being created in the high density wire than in the low density air. To show that this was the case, BEAMnrc's facility to track a particle's origin and history through a simulation using latch filtering was utilised.

Each particle in a phase-space file has a history tag; this is a 32-bit variable that can be used to record the latch values of the component modules it has interacted with (bits 1-23) or originated from (bit 0) (Rogers et al. 2001, 1995). By including the top half of DAVID in the BEAMnrc simulation and assigning a latch value to the wire / air volume (volume 5 in Figure 5.16) it was possible to apply a filter to
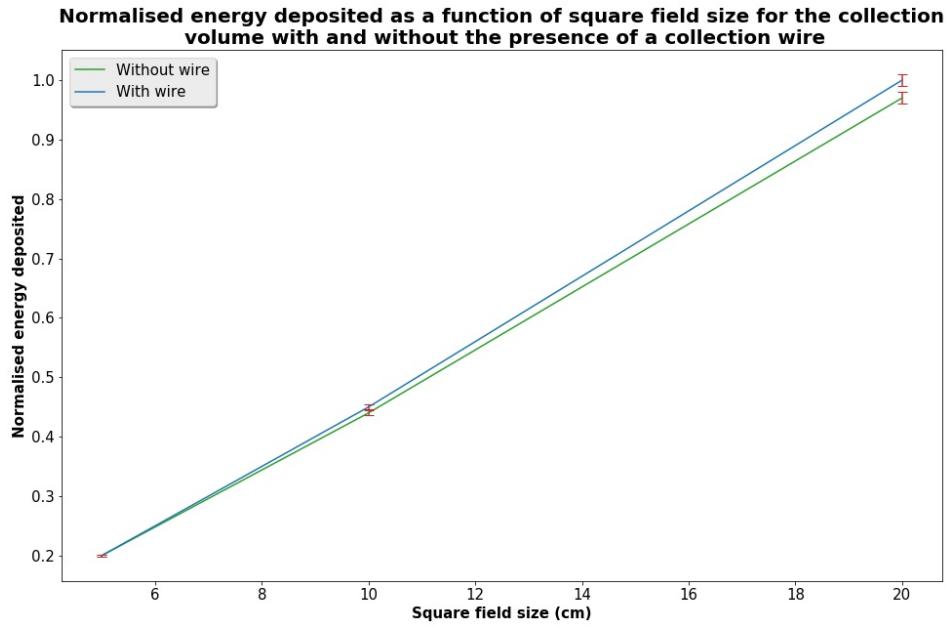
Figure 5.17: The energy deposited in the collection volume for different field sizes for instances with and without a wire. The energy deposited is normalised to the maximum energy deposited (wire 20, 20cm square field). The distances are in cm.

the lower half of the collection volume so that the particles generated in the wire were captured (Figure 5.18) and extracted using BEAMdp (a phase-space analysis program included in the EGSnrc distribution (Ma & Rogers 2018)). Figure 5.19 shows the number of particles passing through the bottom three collection volumes for both the wire and non-wire instances, with and without the latch filter applied. It can be seen that there is an increase in particles generated when the material is defined as tungsten (Figure 5.19, top). Applying the latch filter demonstrates that the increase in signal is a consequence of electrons being produced in the wire that are not present when the wire is defined as air (Figure 5.19, bottom).

The results shown in Figures 5.18 and 5.19 are from wire 20, the central wire. Additionally, data was collected for wire 22 for the $5{\times}5\text{cm}^2$ field; wires 22 and 26 for a $10{\times}10\text{cm}^2$ field; wires 22, 27 and 29 for a $20{\times}20\text{cm}^2$ field. All exhibited equivalent behaviour.
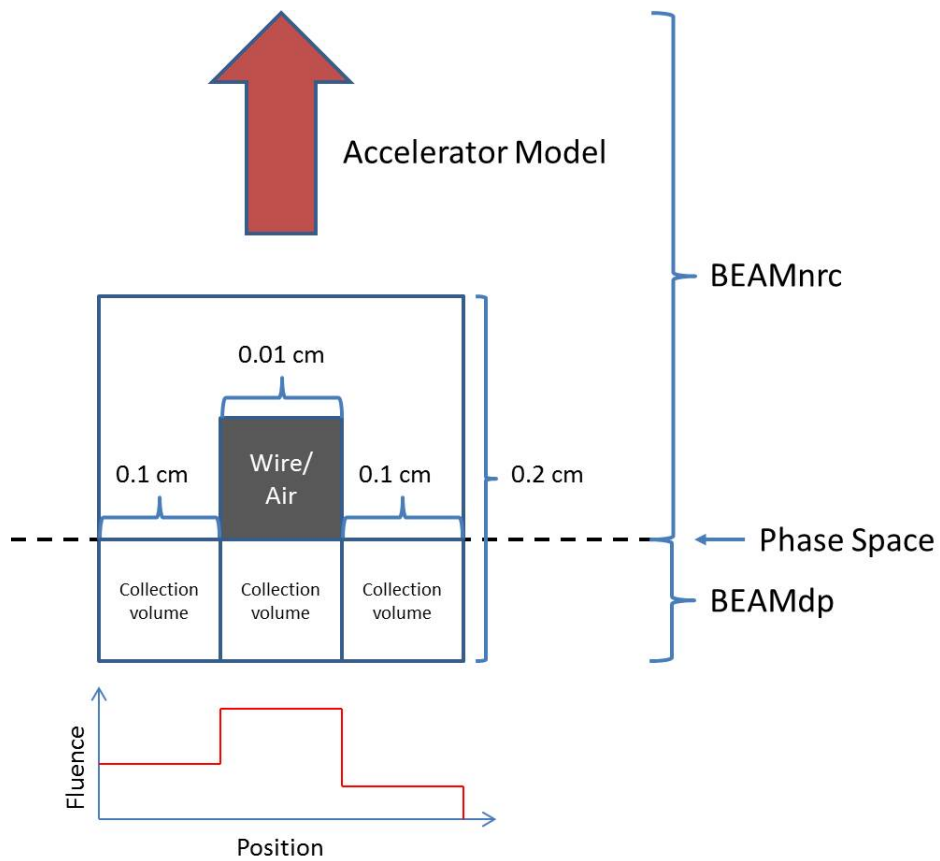
Figure 5.18: The wire and top collecting volumes were modelled in BEAMnrc along with the accelerator. The particle and energy fluence entering the bottom three collection volumes could be extracted from the BEAMnrc-generated phase space using BEAMdp.
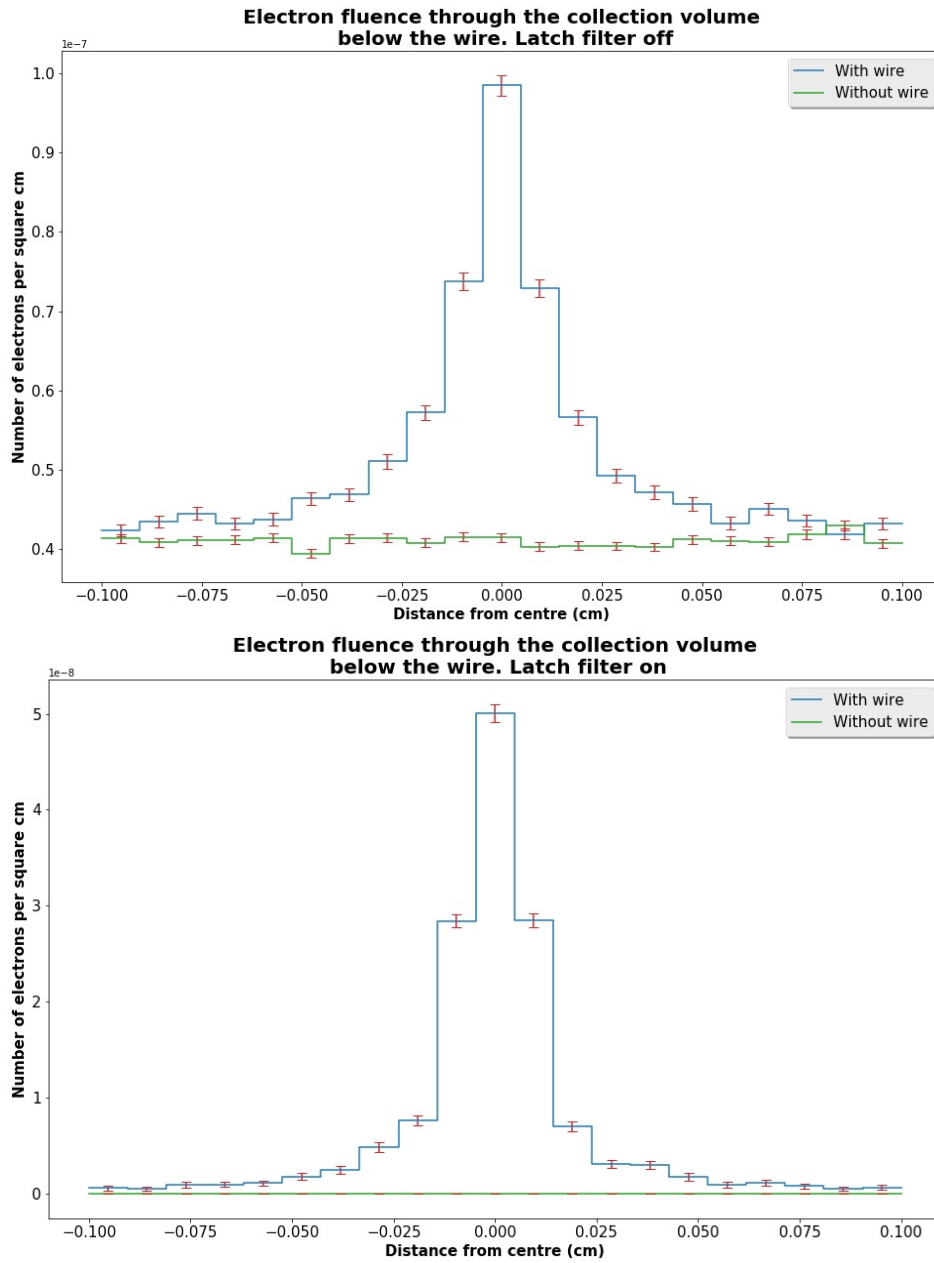
Figure 5.19: Electron fluence passing through the bottom collection volume (7-9 on Figure 5.16) without latch filtering applied (top) and latch filtered so that only particles created in the wire volume are scored (bottom).

5.4.2 Modelling the DAVID with EGSnrc

Materials and dimensions

The DAVID was modelled in DOSXYZnrc as two slabs of Perspex, each 4mm thick, separated by a 2mm thick air gap (Poppe et al. (2006)). In the plane perpendicular to the beam axis, the volumes were split into 250×250, 1mm volumes extending from -125mm to +125mm, the beam axis being at (0,0).

As with Pena et al. (2006) and Liu et al. (2000), energy deposited in the collection volume was used as a surrogate for the signal. This was calculated first by extracting the dose from the central air gap using software written by the author. The dose was summed, at the position of the collection volumes, in the direction of leaf travel along each 250×1mm scoring volume. These values were then multiplied by their volume and the density of air (1.2041kg/m^3) to give energy deposited. As the collection volume is 2mm across (Poppe et al. (2006)), two of these 250×1mm volumes were used to generate the Monte Carlo signal used to compare against the measured signal.

Calibrating the EGSnrc linac model output

Every quantity scored in an EGSnrc simulation is normalised by the number of primary histories, in this case, the number of electrons hitting the target. For BEAMnrc simulations, this number is known to the software; for DOSXYZnrc simulations, this number is estimated using information in the phase space file (Rogers et al. 2001, Walters & Rogers 2003). Knowing this, it follows that the value of any scored quantity should not really depend on the number of histories, however, increasing the number of histories will reduce the level of uncertainty associated with the value. For the purposes of this work, it was necessary to convert the simulated dose to a measured signal. It has been shown that the DAVID signal is linearly proportional to the MUs for equivalent fields (Section 4.2.2). This made it possible to define a MU-based scaling factor that could be applied to the simulated dose so that it could be compared to measured signals. This was done by delivering 10, 50, 100, 200 and 500 MUs to the DAVID through a static 5×5cm^2 open field. The 5×5cm^2 field was then simulated in BEAMnrc and the DAVID signal determined in DOSXYZrnrc. For the simulation, 1.5×10^9 histories were used and the values of ECUT and PCUT were 0.7 and 0.1MeV respectively. The ratio between the DAVID signal and the five different MU deliveries was then calculated and the linear relationship, shown in Figure 5.20, was used to convert subsequent simulated signals to predicted measured signals.

Scaling factor as a function of MU for 5x5cm fields

y = 8.91e+14x + 1.51e+14
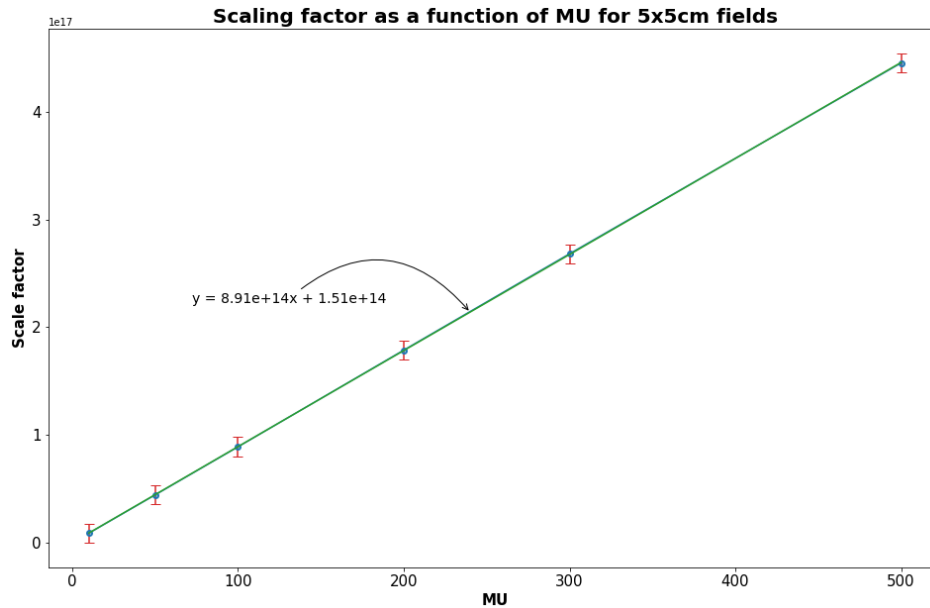
Figure 5.20: Ratio of simulated signal with measured results for 5cm square field for different MU deliveries.

5.4.3 Modelling the DAVID in MONACO

Introduction

Monaco (Elekta AB) is one of a handful of commercially available TPSs that employ Monte Carlo techniques to calculate dose. In contrast to EGSnrc, Monaco necessarily contains features common to all TPSs, namely: patient set-up tools, beam set-up tools, plan comparison graphics, plan optimising options, patient contouring tools etc. The calculation options are much more limited and geometries are constrained by what is physically possible. Furthermore, while EGSnrc is general-purpose code, designed to simulate photons and electrons in a wide range of materials over the energy range of KeV to GeV, Monaco takes advantage of the fact that it will only be used to calculate dose over the radiotherapy energy range (1-30 MeV) for low-Z, tissue-like material with a density range of 0-3gcm^3 (Kawrakow et al. 1996). Although the range is small, unlike EGSnrc where the materials are defined by the user, Monaco has to infer the properties from a CT scan and, ideally, perform the calculation quickly enough for it to be useful in a clinical setting.

The beam model

The beam model in MC-based TPSs is either a full MC calculation: Brainlab (AG) (Künzler et al. 2009) and ISOgray (DOSIsoft) (Salvat et al. 2006); or model based: Peregrine (NOMOS corporation) (Siantar et al. 2001) and Monaco (Seco & Verhaegen 2016). Monaco uses a Virtual Source Model (VSM) that models the particle / photon energy and fluence distribution at three points in the head: a photon source at the target (Fippel et al. 2003), a secondary photon source at the base of the flattening filter (Fippel et al. 2003) and an electron contamination source, also at the base of the flattening filter (Sikora & Alber 2009). These VSMs take the form of Gaussian distributions with the parameters derived and verified

through a combination of MC simulations and measurements taken during the commissioning process (Sikora et al. 2007). Depth Dose Curves (DDCs) and in-air measurements are used to determine the energy fluence of the histograms, while profile measurements dictate the width of the distributions (Fippel et al. 2003, Sikora & Alber 2009). The VSM takes care of the simulation prior to the patient-specific beam-shaping modules. This approach has several advantages over full MC modelling. The particle generation is fast, and the sampling efficiency of the patient-specific part of the simulation is not limited by the size of the phase space file. In contrast to a full MC approach, the VSM is numerically light. It can be defined by a relatively small number of free parameters that can be gathered by the user (Fippel et al. 2003, Sikora et al. 2007), making it much more practical to implement in the clinical setting than a full MC model.

While the model-based approach stops at the top of the beam modifiers in Peregrine, and MC-based particle transport takes over, Monaco use a Transmission Probability Filter (TPF) to model MLC and Jaw transmission. This analytical approach taken by Monaco results in a speed increase by about a factor of 100 compared to a full MC simulation of the beam shaping apparatus (Seco & Verhaegen 2016). The TPF transforms leaf thicknesses into an absorption map. The number of levels of the TPF a photon has to pass through to get from the source to the patient determines the cumulative transmission probability. The TPF is defined by two types of parameters: geometrical and transmission probability factors. Geometrical factors are the position and shape of the leaves and jaws. Transmission probability factors are inter-leaf leakage and leaf transmission (Sikora et al. 2007).

The dose calculation

The combination of the VSM and the TPF create a fluence and energy distribution at the patient surface that can be sampled by the Monte Carlo algorithm for the dose calculation inside the patient. The dose calculation is performed by the X-ray Voxel-based Monte Carlo (XVMC) code (Fippel 1999) which, for electron transport, uses the Voxel-based Monte Carlo (VMC) electron transport method described by Kawrakow et al. (1996). The VMC algorithm works by simulating a single electron history in a homogeneous water environment. The electron scatters through the medium, losing energy through Bremsstrahlung interactions and scattering events. The cross-sections of these interactions define the properties of the resulting particles and are discussed at length in Kawrakow et al. (1996). Electrons produced through these collisions are treated in the same way as the primary electrons. As the contribution to the final dose from Bremsstrahlung photons produced from electron interactions is small, time is saved by not tracking their individual histories. Instead, the Bremsstrahlung histories are immediately terminated and a gamma background applied at the end of the calculation (Kawrakow et al. 1996). The "water history" is then applied to the heterogeneous CT and the distances the electron travels scaled according to the density in the voxels. This process is repeated, and some histories recycled in order to achieve the accuracy required by the calculation (Fippel 1999, Kawrakow et al. 1996).

Calculations relating to photon transport only consider Compton scattering and pair production, as other contributions have a negligible effect in the energy

range of radiotherapy. In both cases, the interaction cross-sections, energies and angles of the particles produced are treated in the same way as EGSnrc (Kawrakow et al. 2000, Fippel 1999). Once produced, electrons and positrons are treated the same and are both subject to the VMC algorithm.

While a DOSXYZnrc simulation allows the user to define the materials in it from the PEGS data, all that is available in the clinical setting is the patient's CT scan. Most MC TPSs map the Hounsfield units to the mass density. This requires a HU-to-material-density conversion; errors in this can lead to significant shortcomings in the dose calculation (Seco & Verhaegen 2016). The developers of the XVMC code opted for the computationally more difficult approach of directly extracting the interaction probability from the CT number (Vanderstraeten et al. 2007). This is made possible because Equation 5.1, describing material density as a function of Hounsfield number, matches measured data (ICRU 1992) well for biological tissue in the range of energies expected in radiotherapy as shown in Figure 5.21. The mass density can then be used to determine Compton and pair production cross sections through methods discussed below.

$$\rho(h') = \begin{cases} -0.008 + 1.033h', & h' \le 0.895 \\ 0.108 + 0.904h', & 0.985 < h' \le 1.1 \\ 0.303 + 0.685h', & 1.1 < h' \le 2.381 \\ 0.580 + 0.580h', & h' > 2.381 \end{cases} \tag{5.1}$$
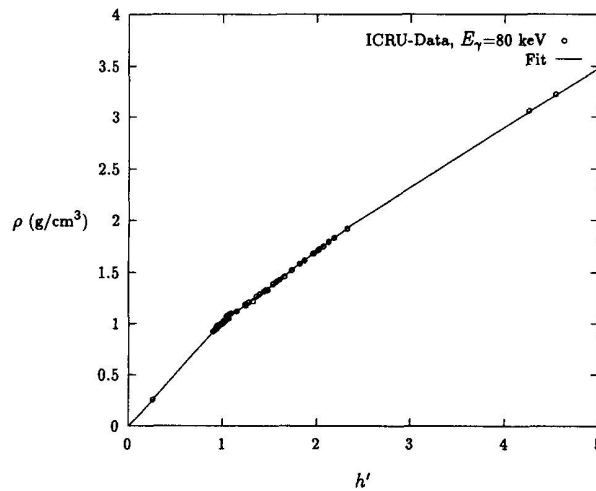
Figure 5.21: Density as a function of Hounsfield number, for measured data (ICRU 1992) and a fit defined by Equation 5.1. Taken with permission from Kawrakow et al. (1996)

Electron transport

Electron transport is dictated by the collision stopping powers, radiation stopping powers, plus the scattering power and density of a material (Kawrakow et al. 1996). Figure 5.21 shows that the density can be determined from the Hounsfield number. The scattering power is defined as: "the mean-square scattering angle per unit thickness of medium due to elastic electron-nucleus Coulomb interactions"(McParland 1989). This, in the VMC code, is determined analytically

from the density (determined from the Hounsfield number) and known constants (Kawrakow et al. 1996). Both radiative and mass stopping power can also be described by functions that match measured data well for the range of energies used in radiotherapy and biological tissues (Figure 5.22).
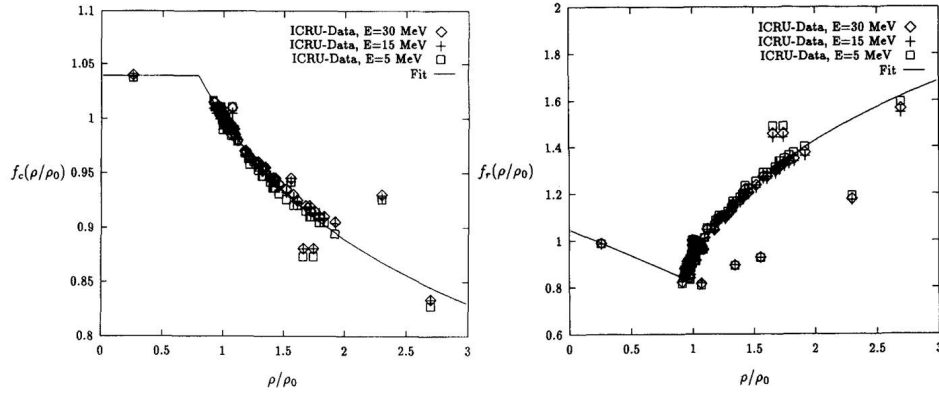
Figure 5.22: Mass collision stopping power normalised to MCS of water as a function of density (right); Mass radiation stopping power normalised to MRS of water as a function of density (left). In both cases, the function used by VMC is plotted along with empirically-derived data from ICRU (1992). Taken with permission from (Kawrakow et al. 1996)

Compton scatter

It can be shown that, over the range of energies in radiotherapy, the linear attenuation coefficient associated with Compton scatter is proportional to the electron density. The relationship between electron density and density can be described by Function 5.2 and density can be derived from Hounsfield units (Figure 5.21). This relationship (Figure 5.23) holds only for natural body tissues (Fippel 1999).

$$
\frac{\eta_e(\rho)}{\eta_e^w} = \begin{cases} \rho/\rho^w, & \rho \leq \rho^w \\ 0.85 \times \rho/\rho^w + 0.15, & \rho \geq \rho^w. \end{cases} \tag{5.2}
$$

Pair production (PP)

Similarly, for the range of energies in radiotherapy, it can be shown that the attenuation coefficient of pair production is proportional to a function $f_r(\rho)$. The cross-sections of Bremsstrahlung and electron-positron pair production can be shown to be equivalent through transformation, or substitution, of their respective Feynman diagrams (Fippel 1999). The mass radiative stopping power (Figure 5.22) depends on the Bremsstrahlung cross section; also, the PP mass attenuation coefficient depends on the PP cross section. This means that they have the same material dependence , so $f_r(\rho)$ for PP is the same as mass radiative stopping power (Fippel 1999).

The dosimetric accuracy of the XVMC code has been demonstrated by its creators. Fippel et al. (2003) demonstrate good agreement between XVMC and EGSnrc, by comparing PDDs and profiles for both homogeneous and heterogeneous phantoms. Fippel et al. (1999) take a similar approach, but find good agreement (less than 2% different) between XVMC and data measured with a pinpoint DIAMOND detec-
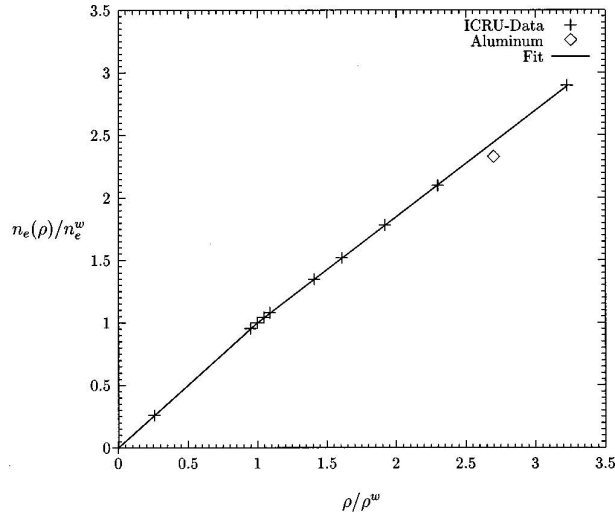
Figure 5.23: Electron density normalised to the electron density of water as a function of normalised density for tissue materials and aluminium. Tissue (defined by ICRU (1992)) is well described by Function 5.2. Aluminium is not well described by this relationship, demonstrating the need to manually enter the electron density of any metal artefacts when using the XVMC code. Taken with permission from Fippel (1999)

tor (PTW, Freiburg) (Laub et al. 1999) and Gafchromic film on a linac. Monaco which combines VSM, TPF and the XVMC dose engine (all discussed above) with the Hyperion biologically-based treatment optimiser (Alber & Nüsslin 1999), has been tested by numerous groups, but Grofsmid et al. (2010) were one of the first to test Monaco in a clinical setting. PDDs and profiles calculated in Monaco were compared to measurements made in a water tank. The PDDs were found to agree within 1% for all depths apart from the build-up region; this discrepancy was still less than 2% and attributed to a shortcoming in the VSM that was improved upon in the Monaco 2.0 release. Measured and calculated in-plane cross-plane profiles of square and rectangular fields were found to have an average discrepancy of 1%, and output measurements of offset rectangular fields were shown to agree on average 0.4 +/- 1.1%. Measurements of step-and-shoot IMRT deliveries taken using a linear array showed excellent agreement between measured and calculated differences, as did point measurements taken inside the tumour region on an anthropomorphic lung phantom. Overall, dose distributions of IMRT fields measured in water and anthropomorphic phantoms, using the Seven29 array (described in Appendix B.2.2) and Gafchromic file respectively, were within the 2%/2mm gamma index criteria and the TPS was deemed fit for clinical use.

5.4.4 Data collection and validation of the Monaco beam model[1]

Monaco beam modelling is done by the vendor (Elekta) using data measured on the specific linac in the clinic and supplied to them, using the model described in Section 5.4.3. The final model can be tweaked in discussion with local staff. What

---

[1]This work was done by many members of the physics team at Leeds Teaching Hospital (LTH). Though the author was involved in collecting beam data and verifying the beam model, the work was done for the clinical service and the burden shared throughout the department

follows is a brief description of the data required. The collection techniques were the same used in Section B.2.1.

- PDD measurements of 1, 2, 3, 4, 5, 7, 10, 15, 20, 30 and 40 cm square fields at a SSD of 90cm. Scans extended to 35cm and were in steps of 1mm.

- In-plane and cross-plane profiles of 2, 3, 5, 10, 15, 20, 30 and 40 cm were measured at depths of d_{max}, 5 , 10 and 20cm, all with a SSD of 90cm. Scans were done using 1mm steps in the penumbra and 2mm steps in the field. Where possible the data collection extended to 6cm outside the field.

- Diagonal scans of a 40cm square field were taken at 45 and 135 degree angles, at depths of 5 and 10cm. Where possible, the scans extended at least 2cm beyond the edge of the radiation field, and they were collected at increments of 3mm. Again, these scans were performed at a SSD of 90cm.

- Absolute dose measurements were taken for 10cm square fields at a depth of 10cm for SSDs of 90 and 100cm.

This data is sufficient to model the virtual sources in the beam model. The post-modelling adjustment, to refine the parameters entered in to the TPF model and better characterise the MLC, requires the measurement of 8 fields. The 8 fields are supplied as part of the Monaco installation; a description of their delivery and purpose is included in Table 5.2. These fields were measured using the PTW Seven29 array discussed in Section B.2.2; the results were returned to Elekta who, in turn, used them to supply LTH with a beam model. The beam model was tested by a variety of static and step-and-shoot deliveries. The dose for these deliveries was calculated on a phantom in Monaco and then delivered to the phantom with a linac. In the case of point measurements, both Farmer and semiflex chambers were used. For comparisons of dose distributions, the Delta4 was used. The Delta4 consists of over 1000 diodes spread over two orthogonal planes within a cylindrical plastic phantom. Dose calculated on a scan of the Delta4 in Monaco can be compared, using gamma index analysis, with the distribution measured on the linac. In this way, even after Monaco had been commissioned, the system was subject to weekly testing through pre-treatment verification. Clinical plans were subject to Delta4 and point-dose assessment until enough confidence was built up in the planning and delivery system that they were no longer required.

Field name	Field description	Field purpose
10×10	Static. 10×10 delivery	Assess calibration of matrix detector.
20×20	Static. 20×20 delivery	Assess calibration of matrix detector and assess beam symmetry.
3ABUT	Three 6 cm wide fields. Step and shoot delivery.	The fields are meant to abut, assessment of the position and width of the junctions evaluates the MLC calibration and offset.
7SEGA	Seven, 2cm wide fields. Step-and-shoot delivery	Same as 3ABUT.
FOURL	Four L-shaped fields. Step and shoot delivery	The L shapes are designed to abut. Assessment of the position and the leakage outside the field is used to help determine the leaf offset, MLC transmission and inter-leaf leakage.
DMLC	Dynamic. 10cm sweep of a 2cm leaf gap.	Tests MLC position and offset calibration as well as transmission.
HIMRT	Clinical head and neck plan. Step-and-shoot delivery.	Evaluation of MLC model in clinical delivery.
HDMLC	Clinical head and neck plan. Dynamic delivery.	Evaluation of MLC model in clinical delivery.

Table 5.2: Field descriptions and their purpose included in the ExpressQA package used for setting up the Monaco MLC model (Kinsella et al. 2016).

The Monaco DAVID model

Using a similar approach to that discussed in Section 5.4.2, the DAVID was modelled in Monaco (Semenenko et al. (2008)) by creating a thin air gap between two slabs of Perspex and placing it at an SSD of 63.4cm. A voxel spacing of 0.1cm was used and the calculation was set to have an uncertainty of 1%. The dose was extracted by exporting the dose plane from the centre of the air gap. This was subjected to the same analysis as the dose plane extracted from the DOSXYZnrc-generated 3ddose file to generate the Monaco-predicted signal.

Unlike the EGSnrc simulation, where the dose is normalised by the number of primary electrons (Kawrakow et al. 2000), Monaco (being a clinical planning system) bases its calculation on the intended number of monitor units. This means that the derived dose should only need a single correction factor that will be the same for all deliveries. Like the EGSnrc factor, this will include the loss in signal from the wires not being modelled. By comparing the total measured DAVID signal of five H&N VMAT deliveries with equivalent Monaco-generated signals,

this was found to be the case (Table 5.3). A value of 4.32×10^{-4} was used to convert the Monaco dose to a DAVID signal.

Delivery	Ratio of Monaco to measured signal($\times 10^{-4}$)
VMAT1	4.43
VMAT2	4.50
VMAT3	4.04
VMAT4	4.29
VMAT5	4.34
Mean	4.32×10^{-4}
S.Deviation	1.77×10^{-5}

Table 5.3: Ratio of total Monaco-generated signal (dose) to total measured DAVID signal (arbitrary units proportional to collected charge)

5.4.5 Calibration summary

The EGSnrc and Monaco approaches to predicting the signal generate two different sorts of 2D dose map. The EGSnrc dose map is normalised to the number of histories used in the simulation. The Monaco dose map is scaled according to the number of MU in the plan. A consequence of this difference is that the dose maps require different treatment in order to convert them to DAVID signal. Both need a scaling factor, however. The EGSnrc dose maps also need to be scaled by the delivered MU. The Monaco dose maps do not need the MU to be included.

The EGSnrc scaling factor was determined by comparing a simulated 5cm square field with the signal measured by delivering different MU through the same square field (Figure 5.20). The equation defined by this linear relationship was used to determine the dose-to-signal correction factor for the EGSnrc-generated dose maps (Equation 5.3).

The Monaco factor was determined by comparing the total Monaco-predicted dose to the total measured DAVID signal for 5 VMAT plans (Table 5.3). An average of these ratios was used to determine the dose-to-signal correction factor for the Monaco-generated dose maps (Equation 5.4).

$$Signal = EGSnrc_{dose} \times MU(8.91 + 1.51) \times 10^{14} \tag{5.3}$$

$$Signal = MONACO_{dose} \times 4.32 \times 10^{-4} \tag{5.4}$$

5.5 Comparing VMAT signals

Both EGSnrc (Section 5.4.2) and Monaco (Section 5.4.3) models were used to generate predicted DAVID signals for individual wires for 5 clinical H&N VMAT fields. H&N neck treatments were selected as these complicated sites demand complex, highly-modulated treatment plans that would test the models more rigorously than less-modulated deliveries.

For the EGSnrc signals, the BEAMnrc accelerator discussed in Section 5.3 was used to generate 5 phase-space files - one for each VMAT plan. In each case,

the plan was exported from the TPS in dicom format; software written for this work, was used to read the dicom plan and extract the leaf and jaw positions with associated cumulative monitor units (MU) for each segment. After applying the calibration factors (derived in Section 5.3) the positions of the jaws and MLCs were written, with the accumulated MU meterset, to a formatted text file designed to be read by the SYNCJAW and SYNCMLCE CMs (Lobo & Popescu 2010) included in the linac model. 1.5×10^9 histories were used in the generation of each phase space, and directional bremsstrahlung splitting was used as a variance reduction technique. The phase space was located 8.5cm from the mylar window (63.39 cm from the target) on the linac head to account for the space between the linac and the DAVID.

The DAVID model in DOSXYZnrc was as described in Section 5.4.2. The Monaco model used was as described in Section 5.4.3

DAVID signals from the two MC approaches were calculated. These were compared with measured signals; a wire-by-wire comparison can be and can be seen in Figures 5.24 to 5.28, the uncertainties that dictated the size of the error bars on these graphs were derived as follows:

EGSnrc: For every volume analysed by DOSXYZnrc both dose and estimated uncertainty are reported. These values were extracted and used for the error bars.

Monaco: Monaco produces can produce an uncertainty map for all the voxels where dose is calculated. This map was exported and the uncertainties extracted.

Measured: The plans were delivered three times on two separate days. The measured signal showed a variation of about +/- 0.5% for all wires, so this was used for the uncertainty on the measurement.

Table 5.4 summarises the differences on the whole VMAT field delivery. The biggest difference between a Monaco-generated signal for an entire delivery was 3.1% (VMAT2), whereas the for an EGSnrc this values was -1.45% (VMAT4).
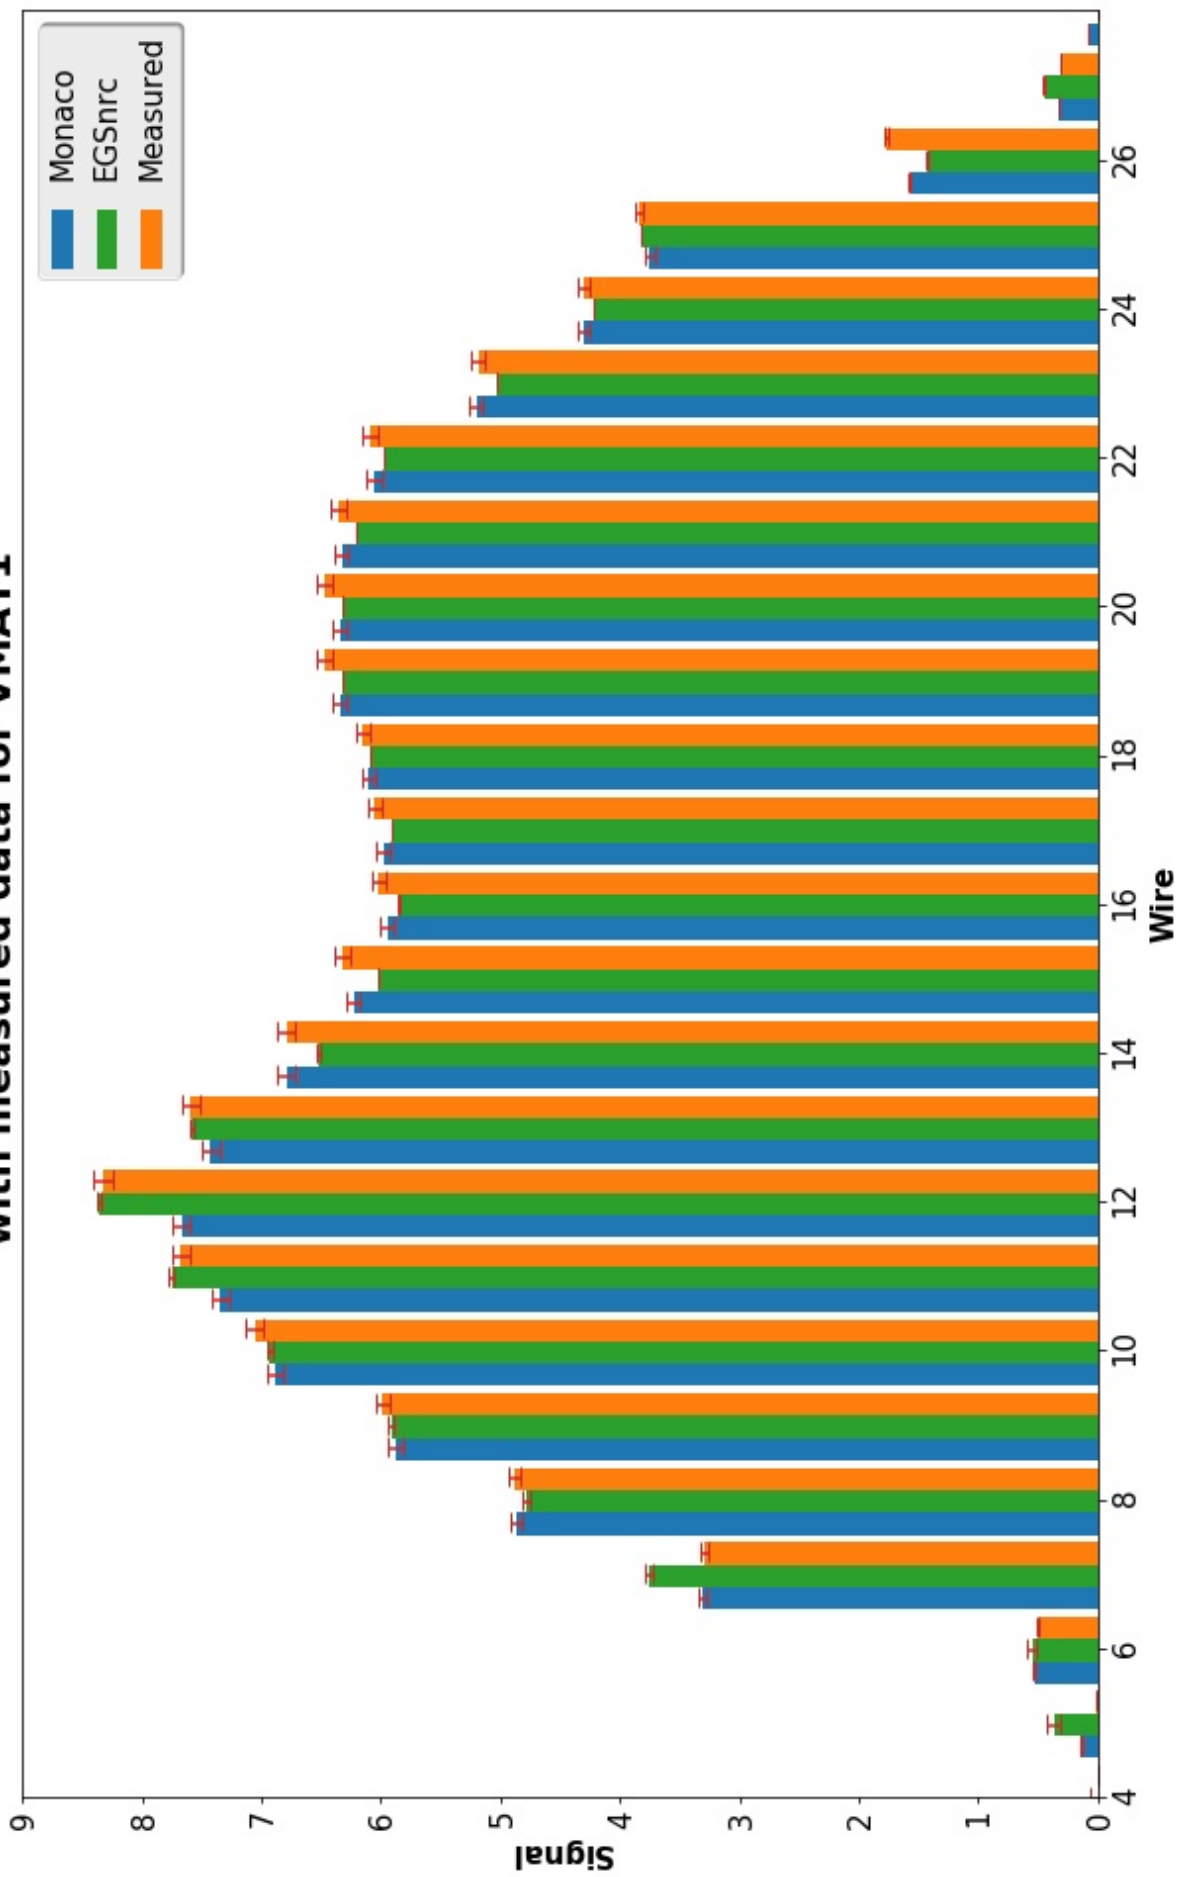
Figure 5.24: Measured and Monte-Carlo generated signals for VMAT 1
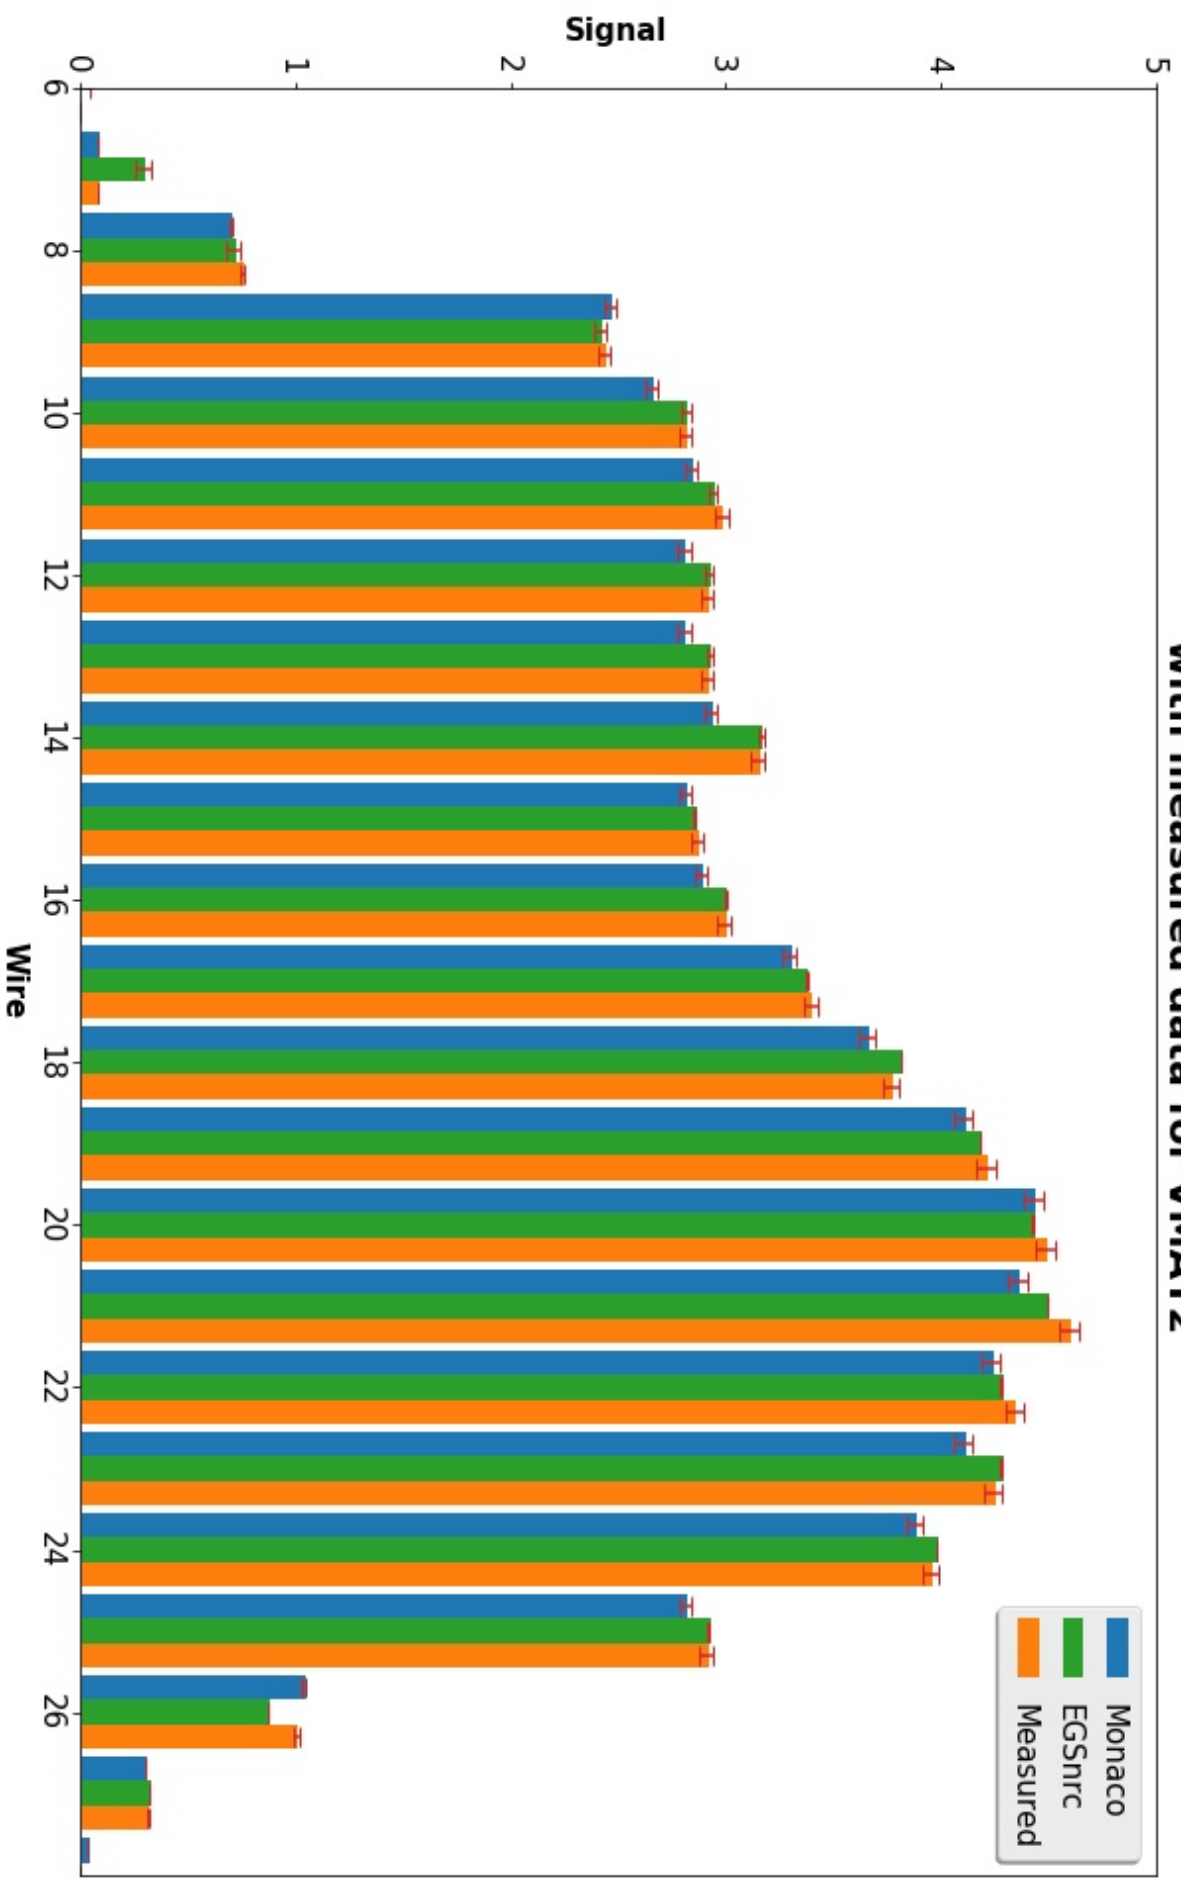
Figure 5.25: Measured and Monte-Carlo generated signals for VMAT 2
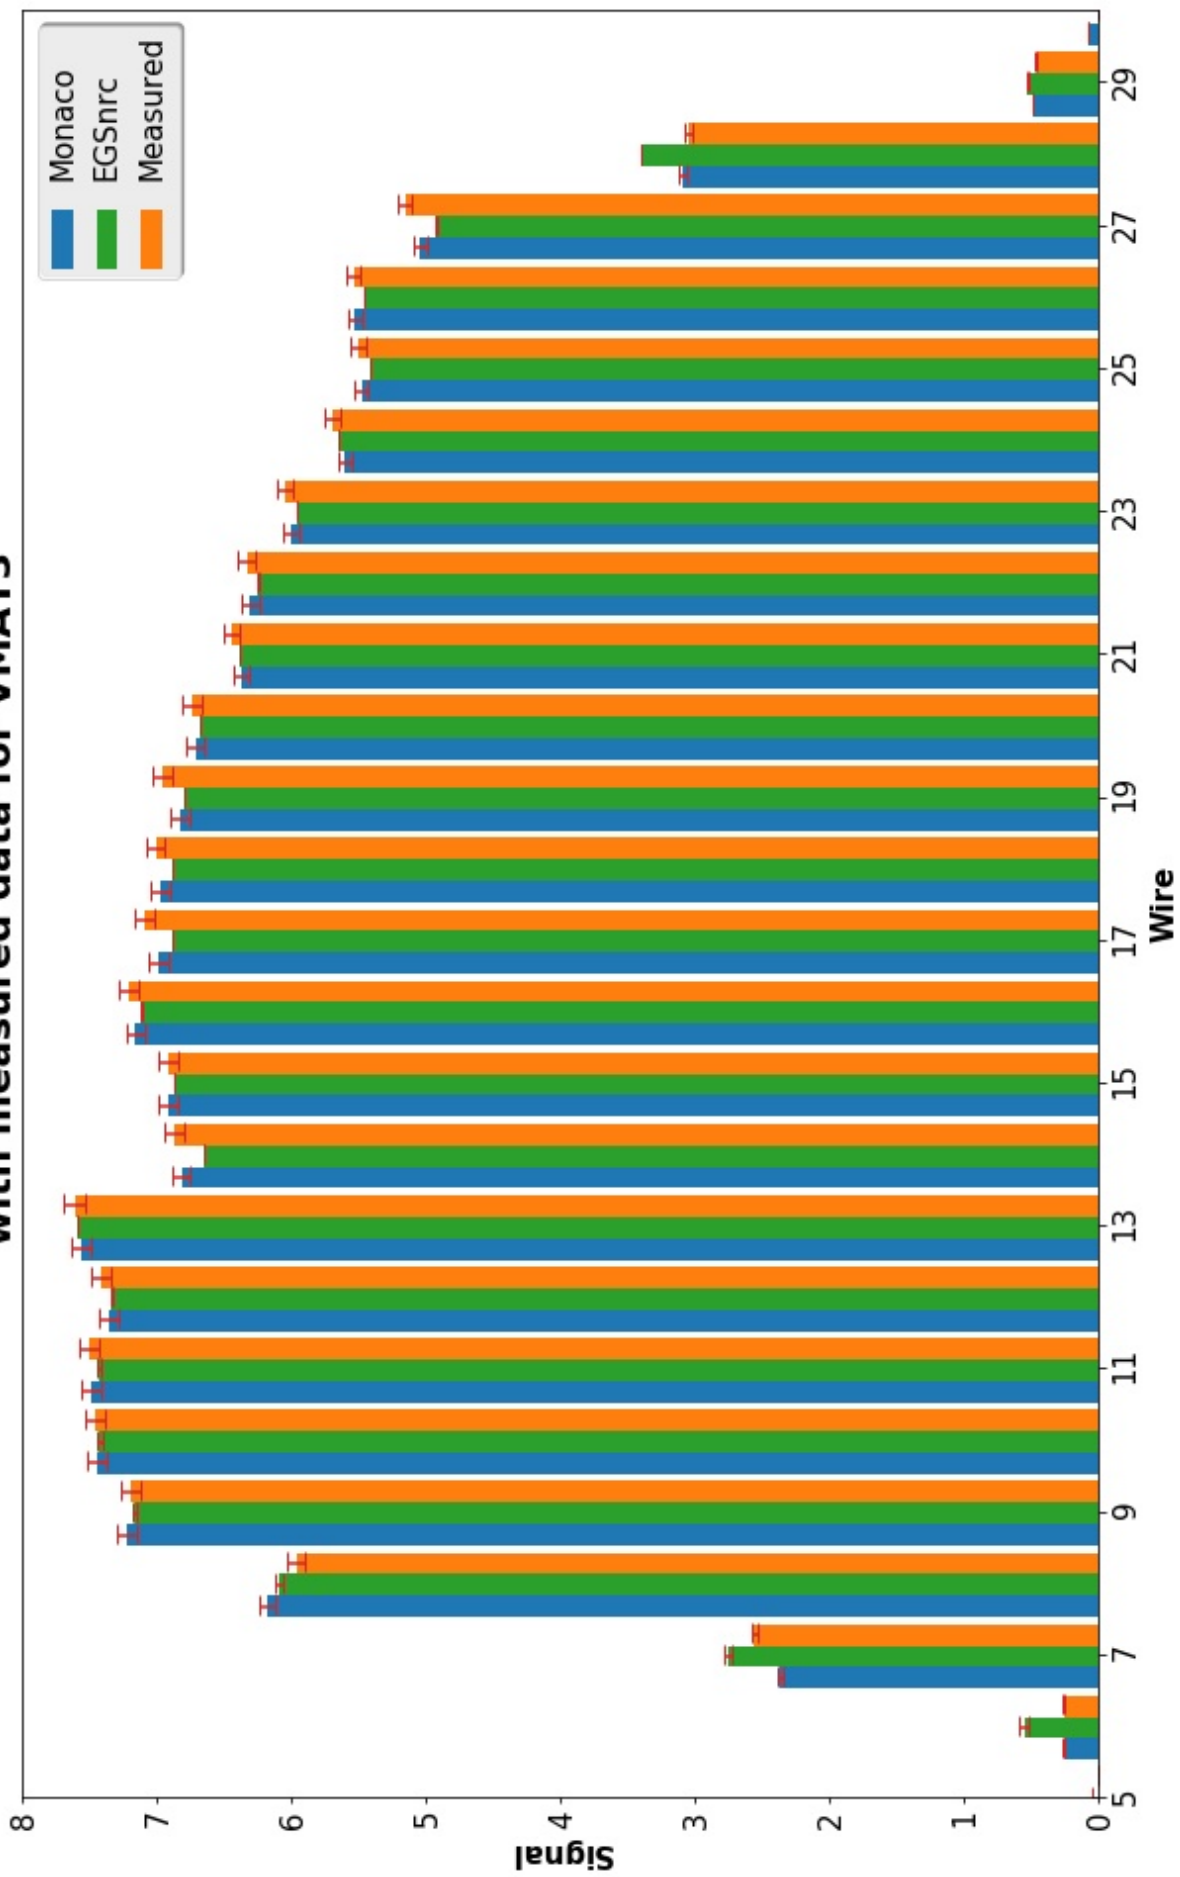
Figure 5.26: Measured and Monte-Carlo generated signals for VMAT 3
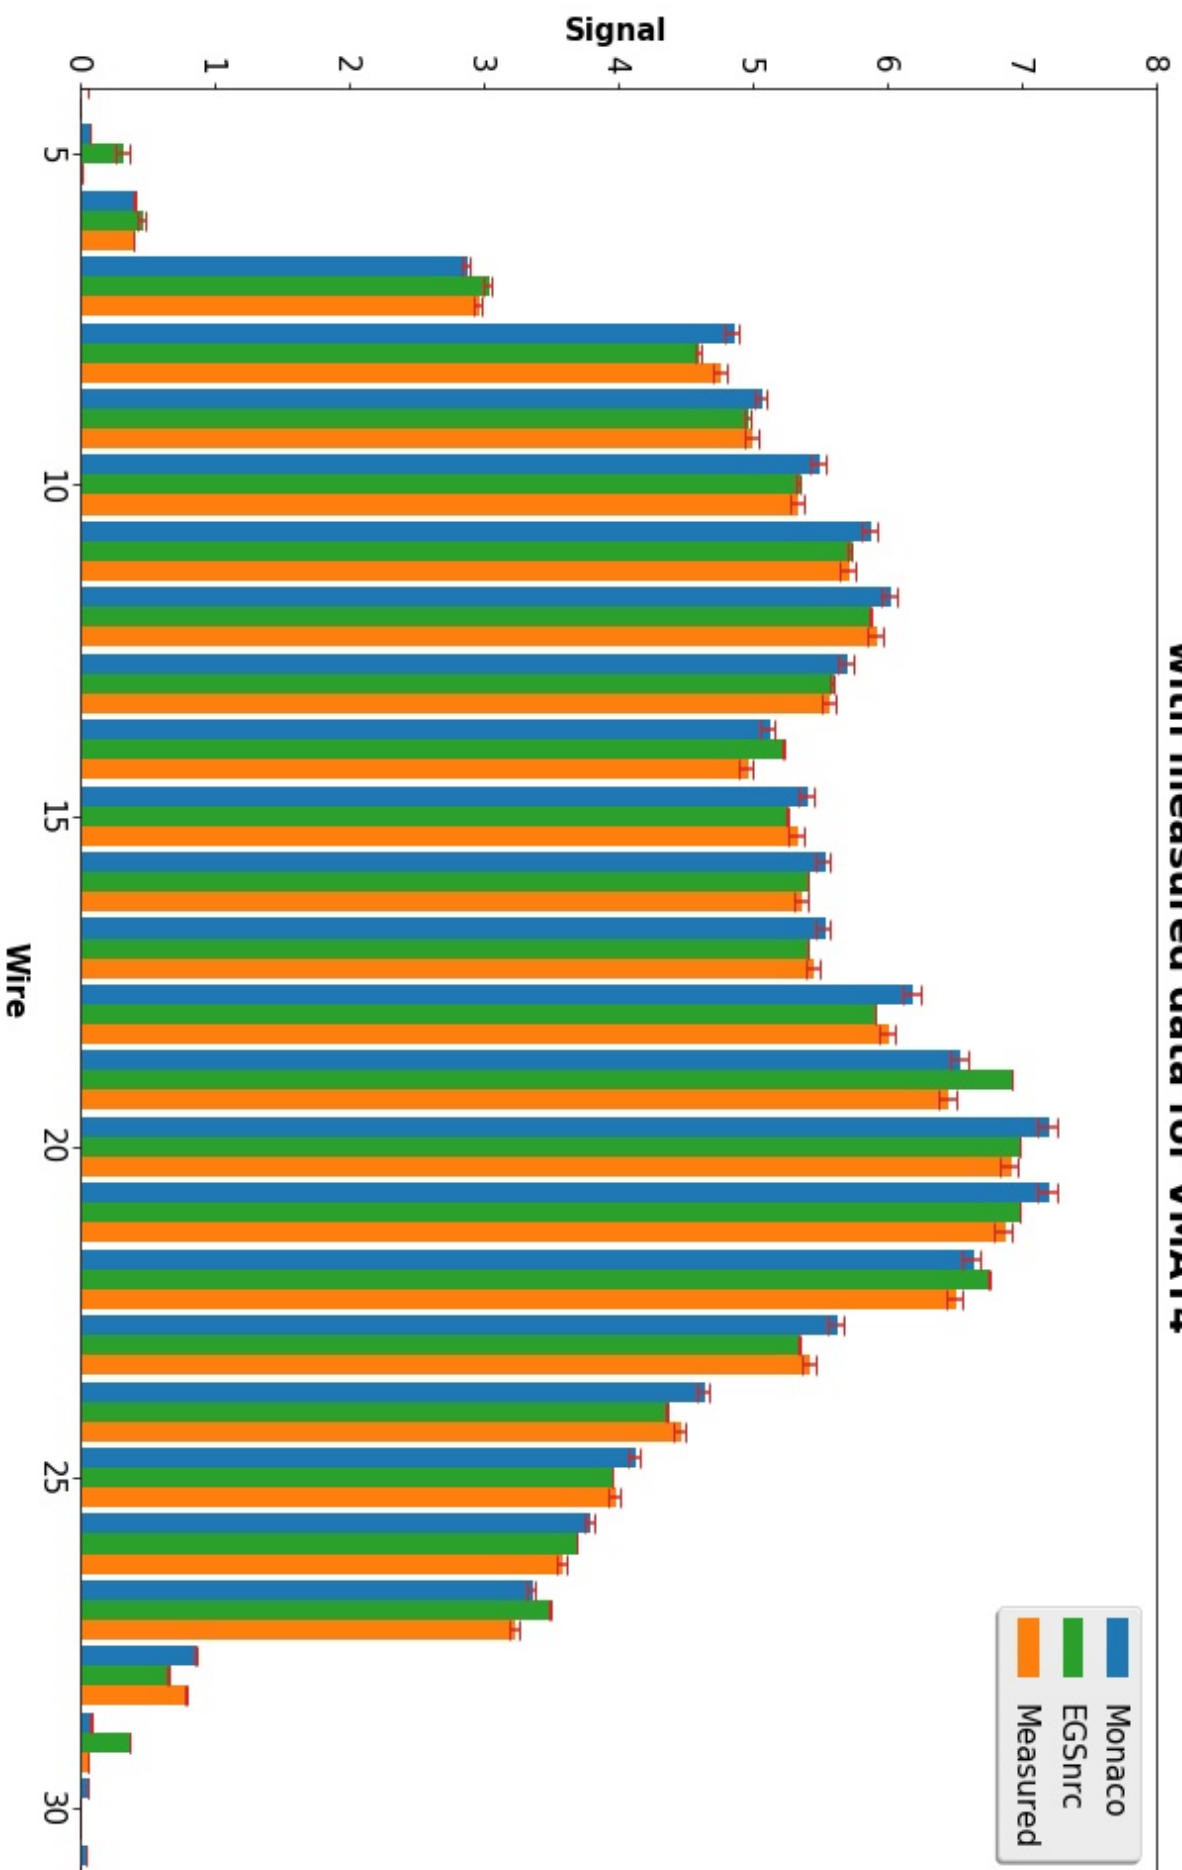
Figure 5.27: Measured and Monte-Carlo generated signals for VMAT 4
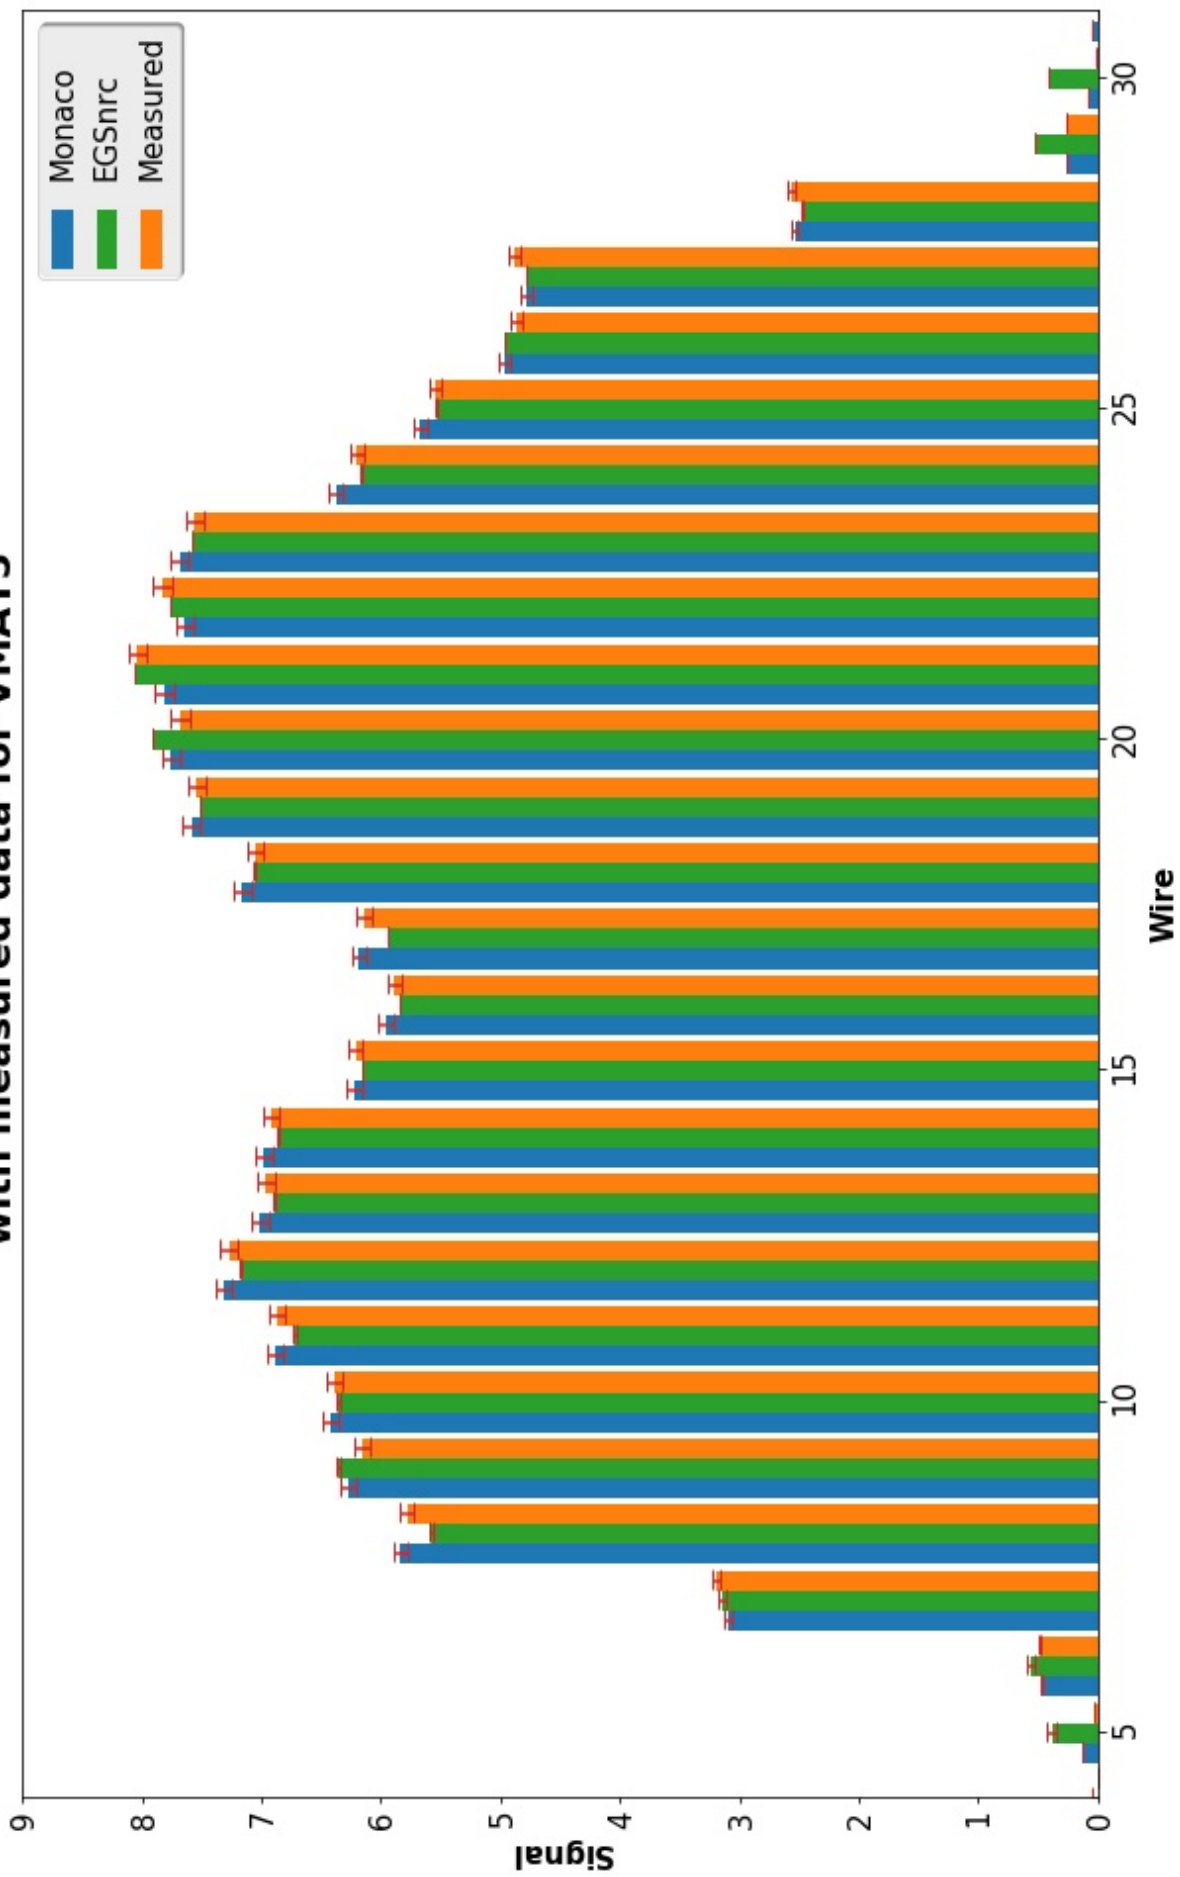
Figure 5.28: Measured and Monte-Carlo generated signals for VMAT 5

	Total Monaco	Total MC	Total Measured	Difference between Monaco and Measured	Difference between Measured and EGS
VMAT1	117.24	117.35	117.44	0.17%	0.08%
VMAT2	59.26	60.69	61.11	3.07%	0.70%
VMAT3	137.75	137.54	138.94	0.86%	1.01%
VMAT4	113.97	112.52	110.90	-2.73%	-1.45%
VMAT5	139.47	138.45	138.28	-0.86%	-0.13%

Table 5.4: The total signal for each VMAT plan for measured and predicted instances.

5.6 Discussion

5.6.1 Spectra

The energy distribution of the electron beam hitting the target shown in Figure 5.10 is reproduced in Figure 5.29. Electron beam energy distributions striking the target are normally described as Gaussian in nature, and are typically presented in literature as a nominal energy with accompanying FWHM (Sheikh-Bagheri & Rogers 2002*b*, Sheikh-Bagheri et al. 2000, Padilla-Cabal et al. 2015, Fix et al. 2004). Given this, the two low-energy components seen in Figure 5.29 are not typical. The slalom beam-bending system employed by Elekta linacs acts as an excellent energy discriminator (Greene & Williams 1997), so it is unlikely any low-energy electrons that somehow made it to the end of the accelerating waveguide manage to get to the target. Furthermore, the electron peaks in Figure 5.29 are too low to be from nuclear relaxations, and a 6MV beam is, strictly speaking, too low energy to cause nuclear excitations. While too low for nuclear excitations, they are also too high to correspond to electron relaxations caused by photoelectric excitations or Compton interactions of material in the linac head by the treatment beam. However, the excellent match of simulated to measured data for PDDs taken at both 90 and 100cm SSD, (Figures 5.10 and 5.11) as well as good penumbral matching of the profiles (Figures 5.12 and 5.13), suggest that the simulated photon beam is equivalent to the clinical beam. To further show that the electron spectrum
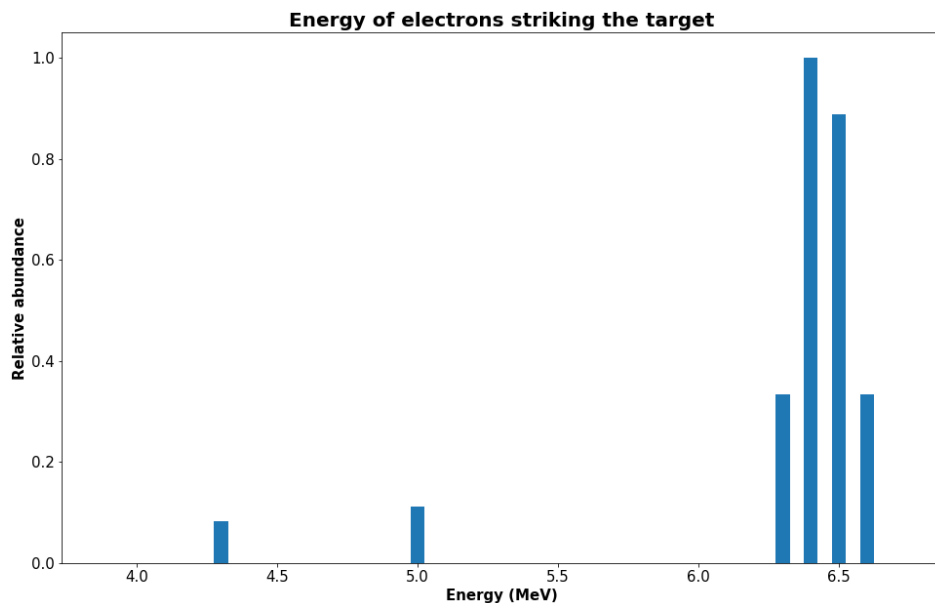
Figure 5.29: The electron spectrum hitting the target for the BEAMnrc simulations.

(Figure 5.29) used was appropriate, the BEAMnrc simulation used to generate the 10cm^2 profiles displayed in Figures 5.12-5.13 was run, but with only 1 million histories. BEAMdp (Ma & Rogers 2018) was used to extract the photon spectrum. This can be time consuming if a lot of histories are used to generate the phase space, hence the comparatively small number used for this investigation. This photon spectrum (Figure 5.30) was qualitatively compared to equivalent photon spectra (not graphed here) defined using Monte Carlo (Sheikh-Bagheri & Rogers

$2002a$) and analytical techniques (Juste et al. 2008). All three spectra were similar in character, showing a sharp build up with the peak at about 0.9MeV, followed by an exponential tail-off with a maximum cut-off at about 6.5MeV. Exact comparison is not appropriate as the linacs that the two studies used were different to each other and different to the one that the model, used in this thesis, was matched to. However, the qualitative likeness between all three suggests that even if the electron spectra used in this model does not reflect the actual electron fluence, then at least the photon spectrum it produces is equivalent to ones used in other works.
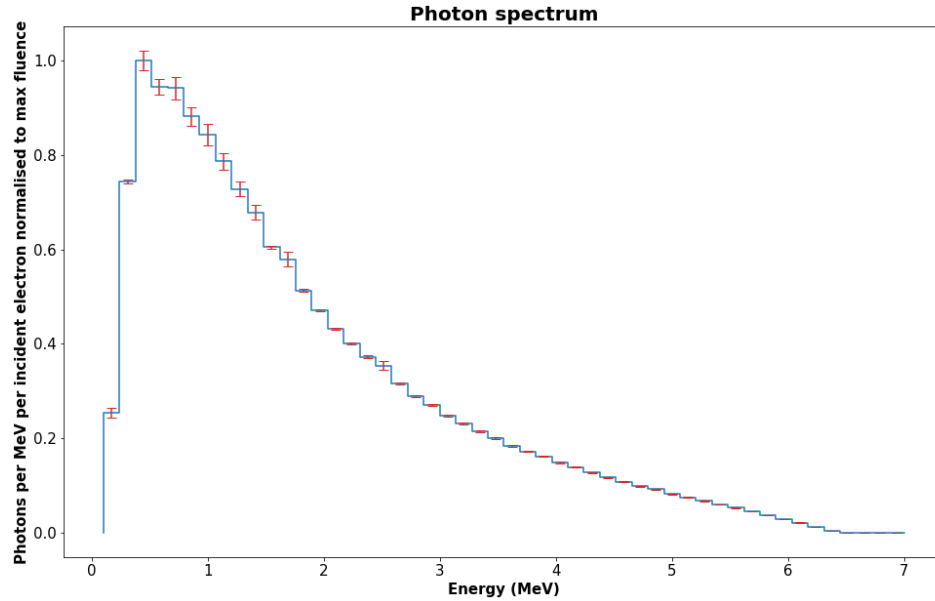
Figure 5.30: Histogram showing normalised photon fluence as a function of energy for a 10×10cm field. Fluence split in to 50 bins with a minimum energy set to 0.1MeV, maximum energy set to 7MeV. The error bars on this graph were defined using the estimated uncertainty reported in the BEAMdp output file (Ma & Rogers 2018).

5.6.2 Calibration

In Section 5.4.2, a 5×5cm field was used for the calibration of the EGSnrc model. Calibration graphs, like that seen in Figure 5.20, were derived for 10 and 20cm$_2$ square fields (Figure 5.31). The spread in results seen in Figure 5.31 could possibly be attributed to backscatter into the monitor chamber. Other groups (Popescu et al. 2005, Oborn et al. 2014) have incorporated the dose scored in the linac's monitor chamber in the MU scaling factor. This approach corrects for the observed phenomena of smaller fields counting up MUs quicker than larger ones, due to the increased backscatter from the jaws contributing to the ion chamber signal (Verhaegen et al. 2000, Popescu et al. 2005, Liu et al. 2000). Popescu et al. (2005) and (Oborn et al. 2014) both show that backscatter can contribute up to 3% more dose to the monitor chamber for smaller fields, compared to larger fields. However, in both cases only Varian linacs are modelled. This effect is investigated for Philips linacs (models and designs that are the basis for Elekta linacs) by Hounsell (1998). This work concluded that the backscatter plate, included between
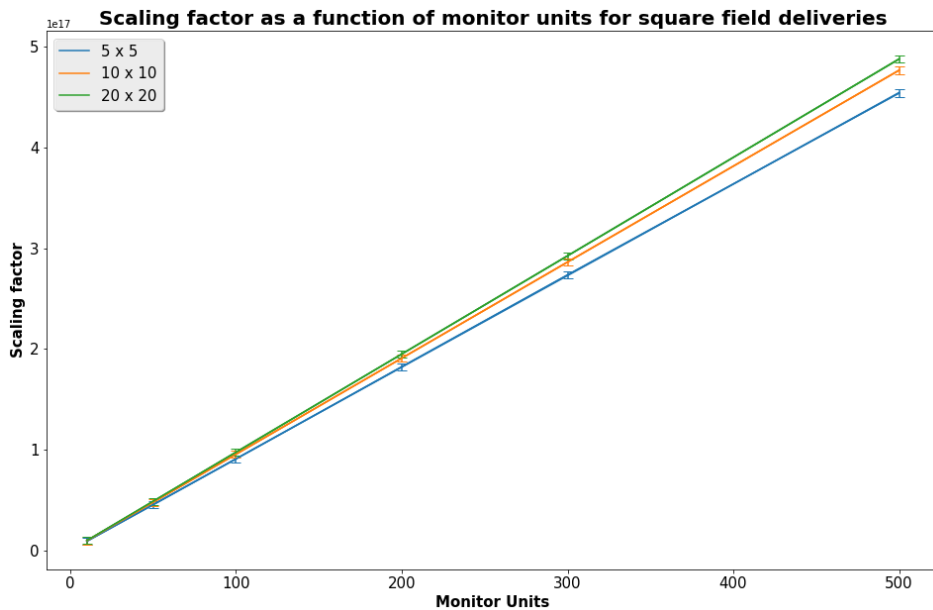
Figure 5.31: Measured and Monte-Carlo generated signals for different field sizes for a range of MUs

the monitor chamber and the beam-shaping apparatus, effectively eliminates this effect, a result backed up by Kairn et al. (2009), where the same effect is shown for a MLCi2 head. To investigate the backscatter effect in this model, the three volumes in the monitor chamber were defined as one scoring plane. Elekta linacs have three monitor chambers, one to count MU, another to verify it and a third for beam-shape analysis. Although only one of these chambers really determines the output of the linac, using all three as the same dose scoring volume improves statistics without affecting relative signal comparison (Popescu et al. 2005). Figure 5.32 shows schematically how the ion chamber was modelled to investigate this effect. The energy deposited in dose scoring volume is recorded after each BEAMnrc simulation in a .egslst file. Simulation of 2.5, 5, 10 and 20cm$_2$ square fields showed that, within error, the results were the same. However, the uncertainty on the measurements was about of the same order of magnitude as the results (Figure 5.33). In an effort to reduce the uncertainty, the number of histories was increased by a factor of 20 (from 1.5×10^{15} to 30×10^{16}) and BEAMnrc simulations run for the 20cm and 5cm$_2$ square fields. In accordance with the backscatter effect, energy deposited in the monitor chamber was seen to be higher in the case of the 5cm field (Figure 5.33). Although this result seems in disagreement with Kairn et al. (2009) (the only published work on this effect for a Elekta MLCi2 head) if the uncertainty in both works are considered, the results are just about consistent. It is possible that this effect is responsible for the spread in the calibration results seen in Figure 5.31; including this factor would have the effect of reducing this. Another possibility is that the spread seen in Figure 5.31 is a result of some other head scatter effect (Kairn et al. 2009). The reason a 5×5cm$_2$ field was used for the calibration was that it best matched the ratios of the simulated-to-measured results for the uncalibrated VMAT deliveries (Figure 5.34). This is thought to be because the smaller square field best matched the scatter / backscatter conditions of the clinical H&N VMAT segments. For a more accurate, general model that
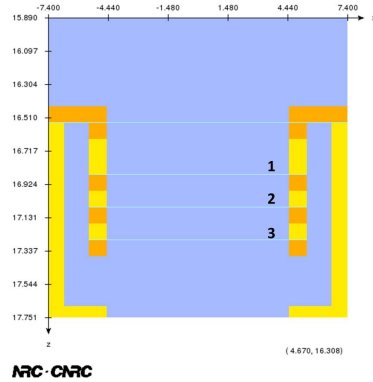
Figure 5.32: The ion chamber in BEAMnrc. Blue represents AIR700icru, yellow AL700icru and orange ALUMINA700icru. Regions 1, 2 and 3 represent the segmented ion chamber, the primary MU counter and the backup MU counter respectively; they are separated by MYLAR700icru. Strictly speaking, it is only the energy deposited in region 2 that contributes to the MU count. However, we are interested in the relative energy deposited in the chamber, so by assigning regions 1,2 and 3 to be the same scoring volume, the total energy deposited in all three is calculated and the statistics are better than if just one of the volumes were used Popescu et al. (2005).

includes larger VMAT segments or even large static fields, this effect is likely to have a greater impact. Future work could include a more in-depth look at the source of the spread in calibration factors (Figure 5.31). If the source is found to be backscatter in to the monitor chamber, then this could be accommodated for by taking a similar approach to that taken by Popescu et al. (2005). If the source turns out to be head scatter on to the DAVID, helper volumes in the DOSXYZnrc model and latch filtering could determine head scatter contributions and the effect taken into account in this way. For the current work, Figure 5.34 shows that the calibration determined by $5 \times 5 \text{cm}_2$ field was sufficient for the VMAT fields that were investigated.
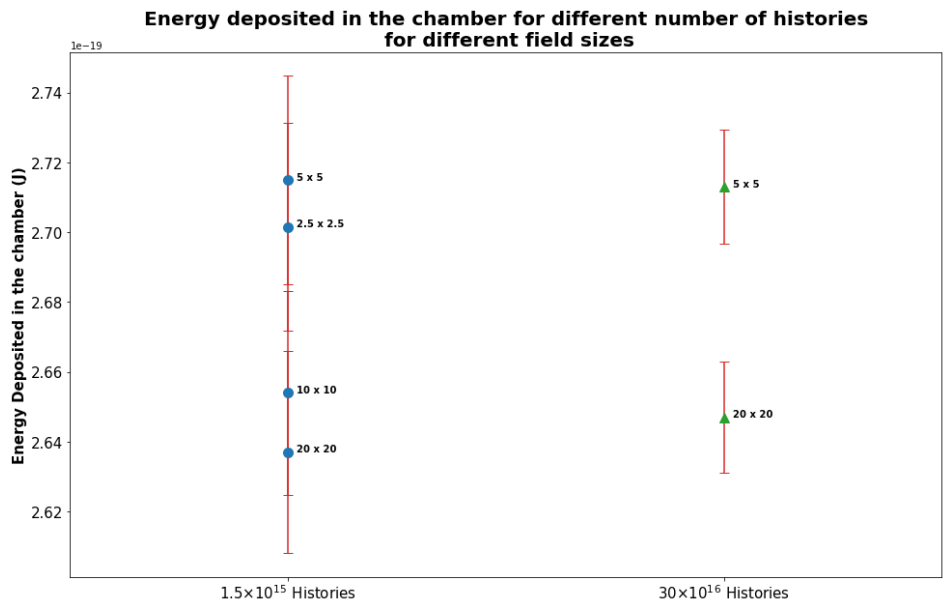
Figure 5.33: Energy deposited in the chamber volumes shown in Figure 5.32 for different numbers of simulated histories. The uncertainty on the scored dose in the initial simulations using 1.5×10^{15} histories was too high to show any effect definitively. The simulations for the 20 and 5cm square fields were rerun using 30×10^{16} histories and were seen to be distinct even when the uncertainty associated with the dose was accounted for. The fact that there is more dose deposited in the chamber for the smaller field is consistent with the backscatter effect.
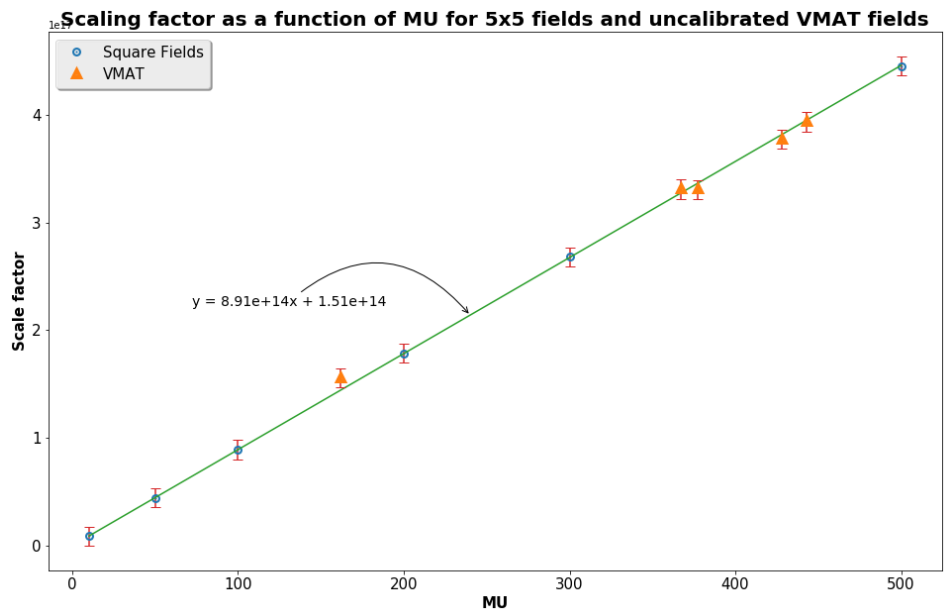
Figure 5.34: Ratio between signal and EGSnrc-simulated dose for 5cm^2 calibration fields and un-calibrated VMAT deliveries.

5.6.3 Signal prediction

Figures 5.24 to 5.28 show a good agreement between the measured signals and the signals predicted by EGSnrc and Monaco. In both cases, the discrepancies between measured and predicted signals show no particular trend, with the total difference between the signals always being less than 3.1%. This is, to the knowledge of the author, the first evidence of the Monaco dose calculation being tested in such a non-clinical set-up (combination of the position and composition of the DAVID). Although phantom-based studies have calculated dose in Perspex, it has never been done at such a low SSD. This provides good evidence that the VSM (Fippel et al. 2003), electron contamination (Sikora & Alber 2009) and leaf transmission (Sikora et al. 2007) used to define the beam model and determine energy fluence at the surface of the dose-scoring medium, work well when scaled to non-patient SSDs. It also shows that the XVMC code models the attenuation and dose deposition in Perspex and air well, despite these not being human tissue.

Signals generated by EGSnrc matched the measured data marginally better than the Monaco-generated signals. This can be attributed to assumptions made in the Monaco dose calculation. These include incomplete Bremsstrahlung modelling, ignoring the photoelectric effect and use of condensed history as a variance reduction technique, which could all result in minor errors over the small volumes used in this study. It was good to see that the clinical treatment planning system, even with a number of approximations, had a performance comparable to the fuller MC modelling done by EGSnrc.

The Monaco calculations performed in the generation of Figures 5.24 to 5.28 took less than 20 minutes using a 16-core processor. The EGSnrc calculations took upwards of two hours using 64 cores on the marc1 computer grid. The Monaco calculation used the pre-existing beam model and could be calculated inside the TPS with limited user training - calculating on the DAVID signal would be much like calculating the dose on any QA or verification phantom. The EGSnrc calculation requires the verification of a new EGSnrc model as well as considerable computer literacy and resources to run. To summarise, the calculation in Monaco is easier and quicker and has comparable accuracy.

Generation of the baseline signal, used to compare with the measured signal on the linac, using the TPS is not as resistant to error as an independent piece of software calculating the signal. In the case of the field being wrong in Monaco, the baseline DAVID signal will be wrong and will simply verify that the wrong field has been delivered. Creation of the baseline signal using the software that checks the dose is a much more resilient approach (Section 4.5.2). Unfortunately, is has been shown that the simple algorithms employed by dose-checking software (Appendix B) may lack sufficient complexity to produce accurate baselines. However, the next generation of checking software is emerging. These generally include more complex algorithms with at least one, Prosoma (Medcom, Darmstadt), using MC techniques (MedCom 2018). Prosoma uses a variation on the XVMC code used by Monaco. It is likely that this program, and ones like it currently emerging onto the market, will have sufficient complexity to calculate accurate baselines for DAVID, and can also be used to independently check the treatment MU / dose. However this will need to be thoroughly tested. If this is the case, then the work flow suggested in Section 4.5.2 could be implemented as the shortcomings discussed in

Appendix B.3 would not be present.

5.7 Conclusion

In this chapter, Monte Carlo algorithms were utilised to better model the DAVID and its interaction with clinical treatment beams for the purpose of signal prediction. EGSnrc and Monaco were used to model the DAVID; both use Monte Carlo-based algorithms to calculate dose. In order to use EGSnrc, an accurate model of the linac had to be constructed. BEAMnrc was used to compile the vendor-supplied geometry of the linac. Using an iterative approach, minor adjustments were made to the position of the MLCs and jaws so that EGSnrc-generated profiles matched measured data (Figures 5.12-5.13). A similar approach was taken with the beam energy (Figures 5.8-5.10).

It was then shown that the presence of the collection wires increases the signal by a constant 3% (Figure 5.17) and that the source of this increased signal was extra electrons generated in the wire passing through the collection volume (Figure 5.19). This effect could be accounted for by multiplying any predicted signal by 1.03, though in practice this amplification becomes part of the factor that converts the Monte Carlo-generated dose to measured signal based on the delivered MU (Figure 5.20)[2]. The details of the EGSnrc model were discussed in Section 5.4.2.

Having shown that the effect of the wires can be accounted for with a 3% correction factor, the DAVID could be modelled in Monaco as two slabs of Perspex, separated by an appropriate air gap. The details of the Monaco model were discussed in Section 5.4.3.

Having built two Monte Carlo models, a comparison between the two was made by using them to predict the DAVID response to five clinical H&N VMAT plans. The predicted results were, in both cases, compared to measured results. The measured and predicted results are shown in Figures 5.24 to 5.28. A summary of the results, showing the difference in total signal between the two predictive approaches and the measured signals, is shown in Table 5.4. The differences between predicted and measured signals are all within 3.1%, with the full EGSnrc model performing marginally better than the Monaco TPS-based model, but the latter performing acceptably well for clinical use and with significantly less effort to set up and use.

There are a number of EGSnrc-based models of the Elekta MLCi2 head dealing specifically with electron beams (Pitcher et al. 2017, 2016); the Elekta MLC head has been modelled in EGSnrc (De Vlamynck et al. 1999, Van de Walle et al. 2003), as has the Beam Modulator (Asnaashari et al. 2013, Herwiningsih & Fielding 2016, Sikora et al. 2007) and the MLCi2 head has been modelled in Geant4 (Fleckenstein et al. 2013). However, to the authors knowledge there are, at the time of writing, no MC models of photon beams for the Elekta MLCi2 treatment head using the EGSnrc code. This makes work outlined in Section 5.3 novel. Furthermore, there are no MC models of the DAVID and the method of calibrating the model - that could be generalised to any transmission detector - is not discussed in literature. It is the intention of the author to publish this work in the near future. It is

---

[2]This work was undertaken on MARC1, part of the High Performance Computing and Leeds Institute for Data Analytics (LIDA) facilities at the University of Leeds, UK.

intended that the published work will discuss the EGSnrc model of the MLCi2 head as well as the DAVID model and its calibration. The validation of each model will be reported through comparison of the predicted DAVID signals of the 5 VMAT plans with measured deliveries. Modelling a transmission detector in Monaco is also novel and the intention is to publish this separately alongside the work discussed in the next chapter.

Chapter 4 used an analytical approach to predict the DAVID response given a 1D matrix of leaf separations. The model could predict the response to the extent that 5% delivery errors were detected. The key shortcomings in that model were the lack of information in the leaf separation matrix about relative separation positions and penumbral modelling. It was stated at the end of the chapter that a 2D fluence or dose map could offer a solution provided the method of generating it was sufficiently complex to accurately model these effects. MU-check software was commissioned and used to generate fluence maps, however, these lacked the complexity to produce accurate signal predictions. It was decided to use Monte Carlo techniques to generate dose maps as these are regarded as the most accurate approach to in-silica modelling of the beams generated by theraputic linacs (AAPM 2007, Seco & Verhaegen 2016, Rogers 2006). The subject of this chapter was to use two MC codes for the purposes of 2D dose mapping at the level of the collection volume in the DAVID. The EGSnrc approach had the advantage of being a complete MC model with the disadvantage of being slower and requiring more advanced computing skills / equipment. The Monaco dose algorithm makes a number of approximations but, in practice, generating a DAVID signal would be quicker and easier to implement than in EGSnrc. Both the MC approaches performed well and were more accurate than the analytical approach. The three approaches (analytic, Monaco MC, EGSnrc MC) increase in accuracy, at the same time as they increase in complexity to both use and implement. The analytical approach uses a handful of variables and can calculate a predicted response to a VMAT delivery in less than 30 seconds on a standard PC. Both the MC approaches require considerable hardware to implement and even then take 15 minutes for a Monaco calculation and over 2 hours for a EGSnrc calculation. For implementation in the clinic it is desirable to have the quickest approach necessary to detect clinically relevant errors with sufficient accuracy. The question posed in the next chapter is how sensitive are these approaches to clinically relevant errors and which, if any, offer a viable solution to routinely predicting an accurate DAVID response.

Chapter 6

The performance of the models in detecting MLC errors using the DAVID

6.1   Introduction

With traditional 3DCRT deliveries, minor ($\sim$1mm) MLC miscalibrations are only likely to affect the dose around the periphery of the target. IMRT uses the MLC to build up non-uniform fluences in order to deliver complex dose distributions. As a result, the MLC separations can be quite small ($\sim$1cm) and, as is the case for plans optimised with the Hyperion system used by Monaco, these apertures can sweep across the target volume during the delivery. As the leaf penumbra determines the dose at various points in the field, there is a high demand on the modelling of the leaves in the TPS and the accuracy with which they are moved to their position by the linac control system (Figure 6.1). It has been shown that for intensity-modulated deliveries (step-and-shoot IMRT, dynamic IMRT, and VMAT), even sub-millimetre errors in the MLC position can have clinically-relevant consequences on the dose distribution across the whole of the irradiated volume (Rangel & Dunscombe 2009, Heilemann et al. 2013, Oliver et al. 2010, Nithiyanantham et al. 2015, Bai et al. 2013).

Rangel & Dunscombe (2009), Mu et al. (2007), Bai et al. (2013) and Oliver et al. (2010) all investigate the impact of systematic and random errors in MLC position. This was done by applying these errors, with various magnitudes, to the MLC positions of clinical treatment plans (prostate and H&N), recalculating the modified plans using the TPS, then looking at the effect they have on target and OAR Equivalent Uniform Dose (EUD) compared with the un-modified plans. EUD is a way of reporting dose that accounts for the asymmetries of the TCP and NTCP curves (Section 1.1) when applied to the non-uniform doses across target and OAR structures inevitably encountered in practical radiotherapy. It is thought to be a better single predictor of clinical outcome than mean dose. In the case of mean dose, a hot spot will mask the effect of a cold spot provided that they are of equivalent volume and equal displacement from the mean dose (Figure 6.2), whereas in reality, this is not the case (Niemierko 1997).
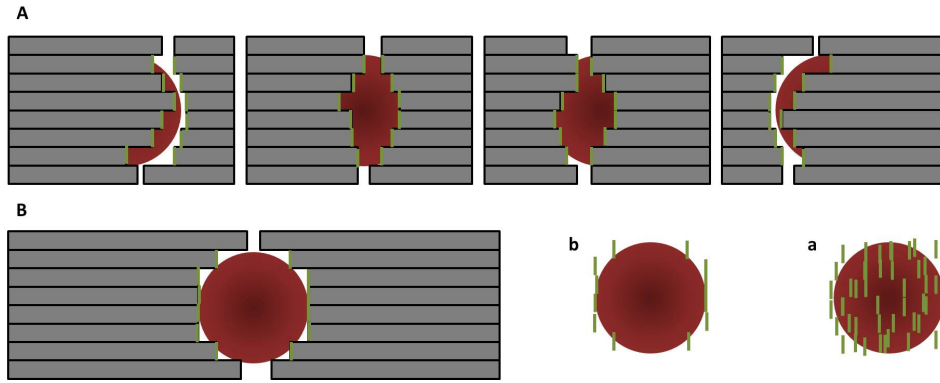
Figure 6.1: An IMRT field split into four segments (A) and a 3DCRT field conformed to the same target (B). The leaf penumbra clearly makes up more of the field in the case of the IMRT delivery (a) than it does with (b)

Random errors

LoSasso et al. (2001) cite motor error or fatigue as likely sources of MLC position error. The effects of motor fatigue/error on leaf velocity could introduce trailing/leading leaves that may contribute to random errors during dynamic deliveries. Although daily variation in MLC position is typically small (0.1 - 0.2mm (Budgell et al. 2000)), the random errors (Figure 6.3 E) could, in principle, be as large as the sum of the precision of the control system and accuracy of the calibration (these being 0.1 and 1mm, respectively, for Elekta systems) (Budgell et al. 2000). Oliver et al. (2010), Bai et al. (2013) and Rangel & Dunscombe (2009) simulated random errors by modifying the original plan's MLC positions by values sampled from a Gaussian distribution centred on 0 with a standard deviation equal to the magnitude of the error being investigated. Mu et al. (2007) randomly sampled shifts from discrete values of -2.0 mm, -1.0 mm, 0, +1.0 mm and +2.0 mm. In each case it was shown that random errors of up to 2mm in MLC position had a negligible effect on the calculated EUDs when compared to the original plan (Figure 6.4). Given this, the effect of random errors is not investigated in this chapter.

Systematic errors

Systematic errors can be introduced through miscalibration (Figure 6.3 B, C and D) and gravitational effects (Figure 6.3 D) (Budgell et al. 2000, Oliver et al. 2010). Again, these errors were investigated by Zhen et al. (2013), Oliver et al. (2010) and Rangel & Dunscombe (2009) who conducted TPS-based studies. In the studies conducted by Oliver et al. (2010) and Rangel & Dunscombe (2009), leaf positions were modified in the DICOM RT files. These were recalculated in the TPS and compared to the unmodified plan doses. Zhen et al. (2013) planned the treatments using a "golden beam model", then recalculated the dose using a modified beam model. All three showed that type B and type C shifts (Figure 6.3) produced errors of equal magnitude and opposite sign. A type B shift of $x$mm will increase the dose to the target volume by d% and a type C shift of $x$mm will decrease the dose to the target volume by d%. Given this, the results were presented in % change per mm shift. For RapidArc H&N plans, Oliver et al. (2010) reported error gradient of 3.2%mm^{-1}, Zhen et al. (2013) reported 1.2%mm^{-1} and Rangel & Dunscombe (2009) reported 5.6 %mm^{-1} for CTVs. Oliver et al. (2010) and Rangel & Dun-
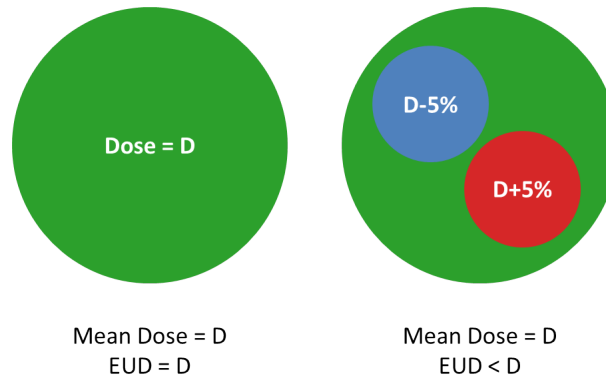
Figure 6.2: For the case where the the dose is perfectly homogeneous, the EUD matches the mean dose, introducing inhomogeneities moves away from this. For example: if two equally-sized volumes of equivalently high and low dose are introduced to an otherwise homogeneous distribution (right), the mean dose will be the same as the mean dose of the homogeneous distribution (left). However, the loss in TCP from the low-dose region will not be fully compensated for by the increase in TCP gained by the high-dose region. This is due to the asymmetry in the TCP curve around the treatment dose. This is why, traditionally, uniform and precise dose delivery is one of the corner stones of accurate radiation therapy (Brahme 1984). EUD is an attempt to account for inhomogeneities like this when reporting tumour doses.

scombe (2009) used Eclipse (Varian Medical Systems, Palo Alto, USA), whereas Zhen et al. (2013) used Pinnacle[3] (Philips Radiation Oncology Systems, Fitchburg, WI), all three use collapsed cone algorithms to calculate dose, but the range of results could be due to the different planning systems in each study generating plans with different modulations, though there is no clear evidence to support this. Typically, OARs are equally (Zhen et al. 2013), or less (Oliver et al. 2010, Rangel & Dunscombe 2009) sensitive and Rangel & Dunscombe (2009) showed that H&N plans were more sensitive to leaf errors than prostate plans, attributing this to the higher degree of modulation (Figure 6.5). Zhen et al. (2013) noted little difference in sensitivity between prostate and H&N cases.

The findings reported by Rangel & Dunscombe (2009) and Heilemann et al. (2013) showed that current approaches to pre-treatment verification (Alber et al. 2008, Stathakis et al. 2013, Ezzell et al. 2003, AAPM 2018) using either the Octavius and (Appendix B.2.3) or the Delta 4[1] (Scaniddos), and applying 3%/3mm gamma index analysis (Appendix A) were not sufficient to detect clinically-relevant leaf bank errors. In order to detect clinically-relevant leaf bank errors, a more advanced approach is needed. A similar conclusion was reached by Zhen et al. (2013), whose in-silica approach used gamma index analysis to compare plans made with and without a faulty beam model. By doing this, they demonstrated that gamma index analysis pass rates were not a good surrogate for TCP/NTCP changes for identifying TPS or machine errors. Instead, Zhen et al. (2013) propose metrics based upon DVHs. This concept was used by (Nithiyanantham et al. 2015), who demonstrated that the MatrixX (Sun Nuclear) device (a head-mounted detector array) could be used to record dose, then back-project this through a CT and structure data set to determine the delivered dose. This approach was shown to

---

[1]Another widely-used, pre-treatment verification device consisting of two orthogonal diode arrays (Feygelman et al. 2010)
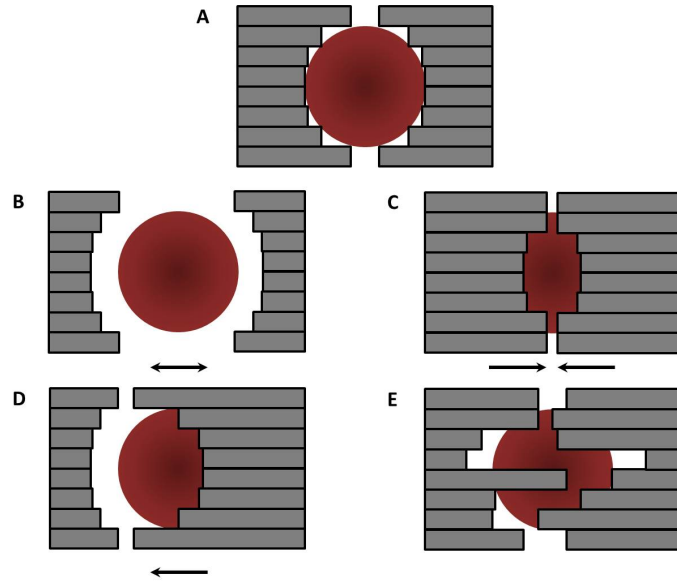
Figure 6.3: MLC errors. A. Baseline B. Systematic growth C. Systematic shrink (if leaves overlap due to the shrink, then they are set to minimal separation) D. Systematic shift E. Random errors. (Oliver et al. 2010). The errors drawn here are much larger than would be expected in reality, the effect has been exaggerated for illustrative purposes.

be sensitive to minor leaf errors.

The DAVID does not contain an inclinometer, so this sort of analysis is not possible. However, it has been shown to be sensitive to small ($\sim$ 1mm) changes in leaf position when compared to a baseline measured on the linac (Poppe et al. 2006). In this chapter, it will be shown how sensitive the predictive algorithms developed in previous chapters are to leaf errors; also, given the uncertainty in the predictions, whether they are capable of detecting these errors.

6.2   Materials and method

Five clinical H&N plans were selected and were delivered by an Elekta Synergy linac with a MLCi2 head (Crawley, UK) with the DAVID mounted on the head. The measured signals were compared to the signals predicted by the Basic Algorithm (Chapter 4, (Johnson et al. 2014)), the Monaco model (Section 5.4.3) and the EGSnrc model (Section 5.4.2). The DICOM RT files were then modified for each plan so that the leaves were moved +/- 2, +/- 1 and +/- 0.5 mm, where positive displacements indicate a move of opposing leaf banks away from the central axis (Figure 6.3, B) and negative displacements indicate a movement of opposing leaf banks towards the central axis (Figure 6.3, C) creating larger and smaller fields respectively. Predicted signals were generated for each of the 6 modified plans using each of the three models for all 5 VMAT plans. In the case of the basic algorithm, the same variables were used to calculate the predicted response as those determined in Section 4.3. In the case of the the Monaco and EGSnrc predictions, the same simulation parameters were used as those outlined in Section 5.4.
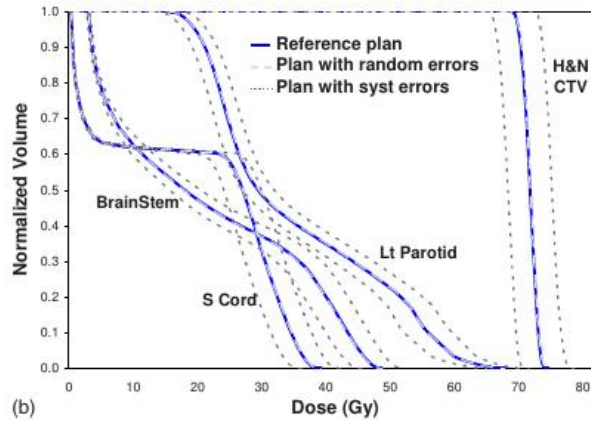
Figure 6.4: Effect of random and systematic MLC errors on cumulative DVHs for a H&N plan. Taken with permission from Rangel & Dunscombe (2009)
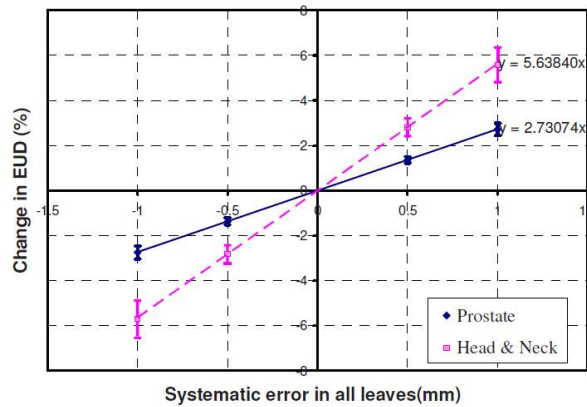
Figure 6.5: Change in dose as a function of MLC error for prostate and H&N plans. Taken with permission from Rangel & Dunscombe (2009)

6.3   Results

Figures 6.7 to 6.11 each show the results for a VMAT plan. Each figure contains three graphs, one for each approach - Basic Algorithm (Chapter 4), Monaco (Section 5.4.3) and EGSnrc (Section 5.4.2). Each of the three graphs show the difference between the signals from the 6 modified plans and the unmodified plan in addition to the difference between the unmodified plan and the measured signal. The measured signal was taken as an average of the signals taken from repeating the deliveries on 3 different days using two different linacs. The process of generating Figures 6.7 to 6.11 was as follows:

1. The predicted signal for the VMAT plan was generated using one of the three approaches.

2. 6 RT plan files were generated with leaf errors of +/- 0.5, +/- 1 and +/- 2mm (boxed legends in Figures 6.7 to 6.11).

3. Using the same approach as step 1, a predictive signal was generated for each of these 6 plans.

4. The signal generated in step 1 was deducted from each of the 6 plans gener-

ated in step 3.

5. This difference in signal calculated in step 3 was plotted as a function of position on the DAVID for all 6 modified plans.

6. The difference between the measured signal and the signal generated in step 1 was calculated and displayed on the same axis.

A graphic outlining the process of creating these graphs is shown in Figure 6.6.

The total signal for all the measured, modified and unmodified plans was calculated by summing the response for each wire for each plan. The difference between the total measured signal and the un-modified predicted signal are shown in Tables 6.1 to 6.3. Additionally the difference between the unmodified predicted signal and the modified plans are also displayed.

Figure 6.6: Explanation of how the graphs are developed in this section
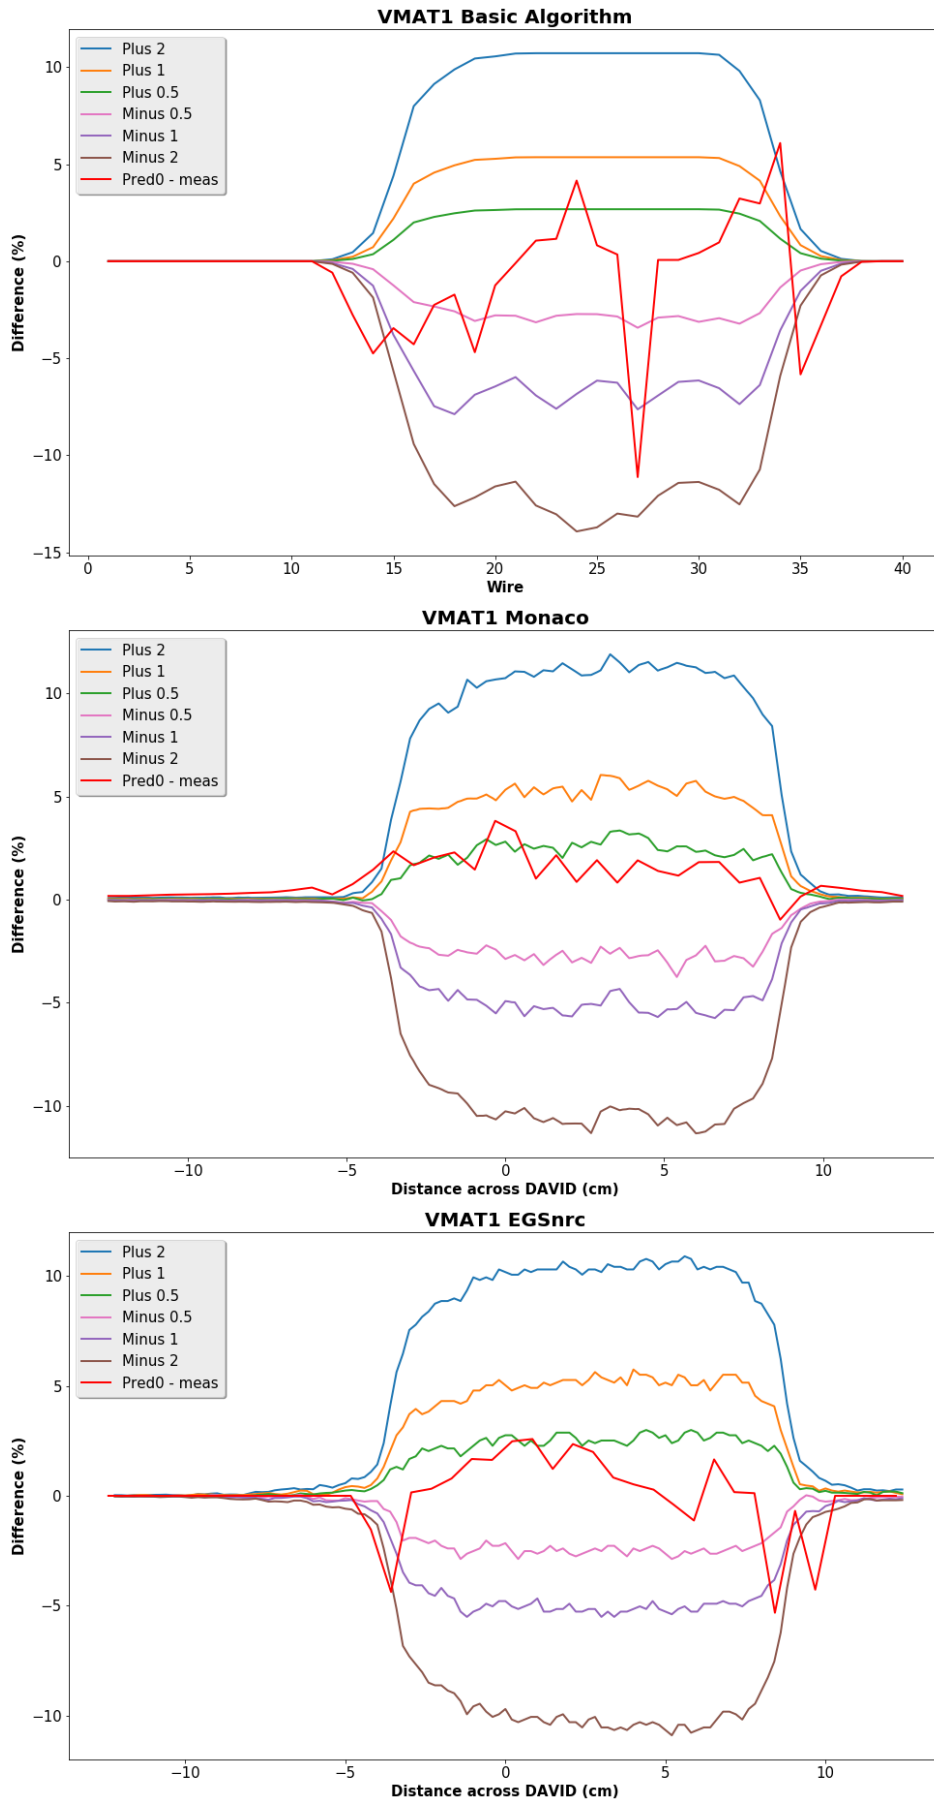
Figure 6.7: Results for all three approaches for VMAT1
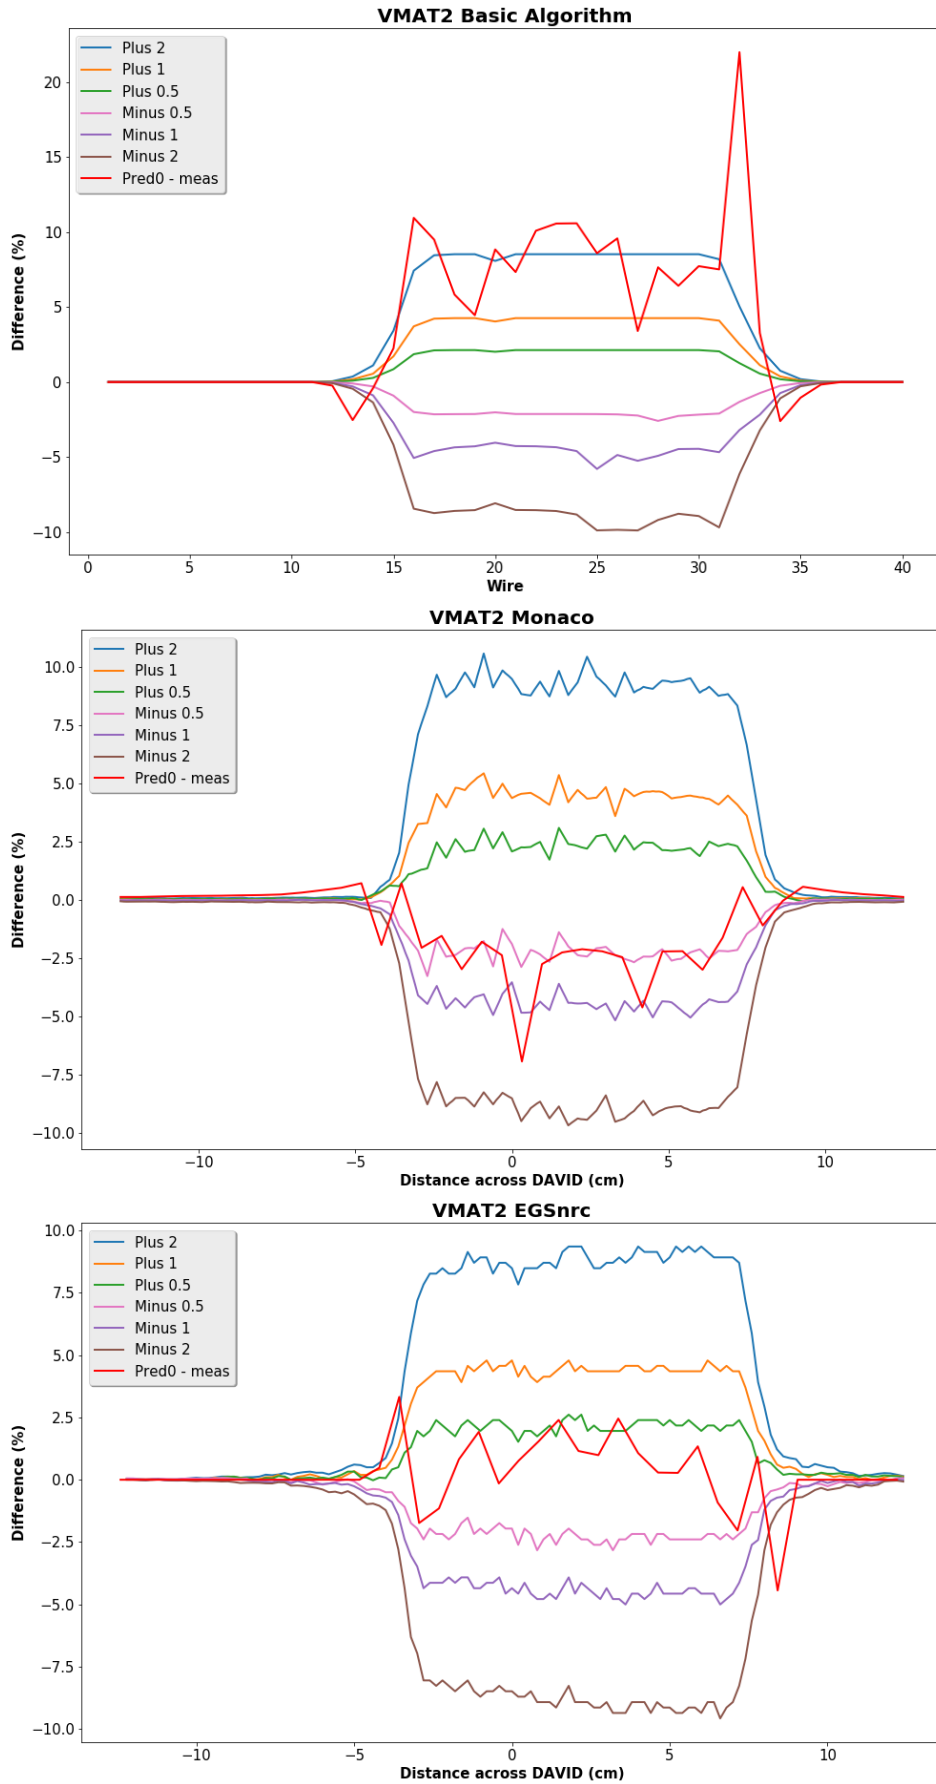
Figure 6.8: Results for all three approaches for VMAT2
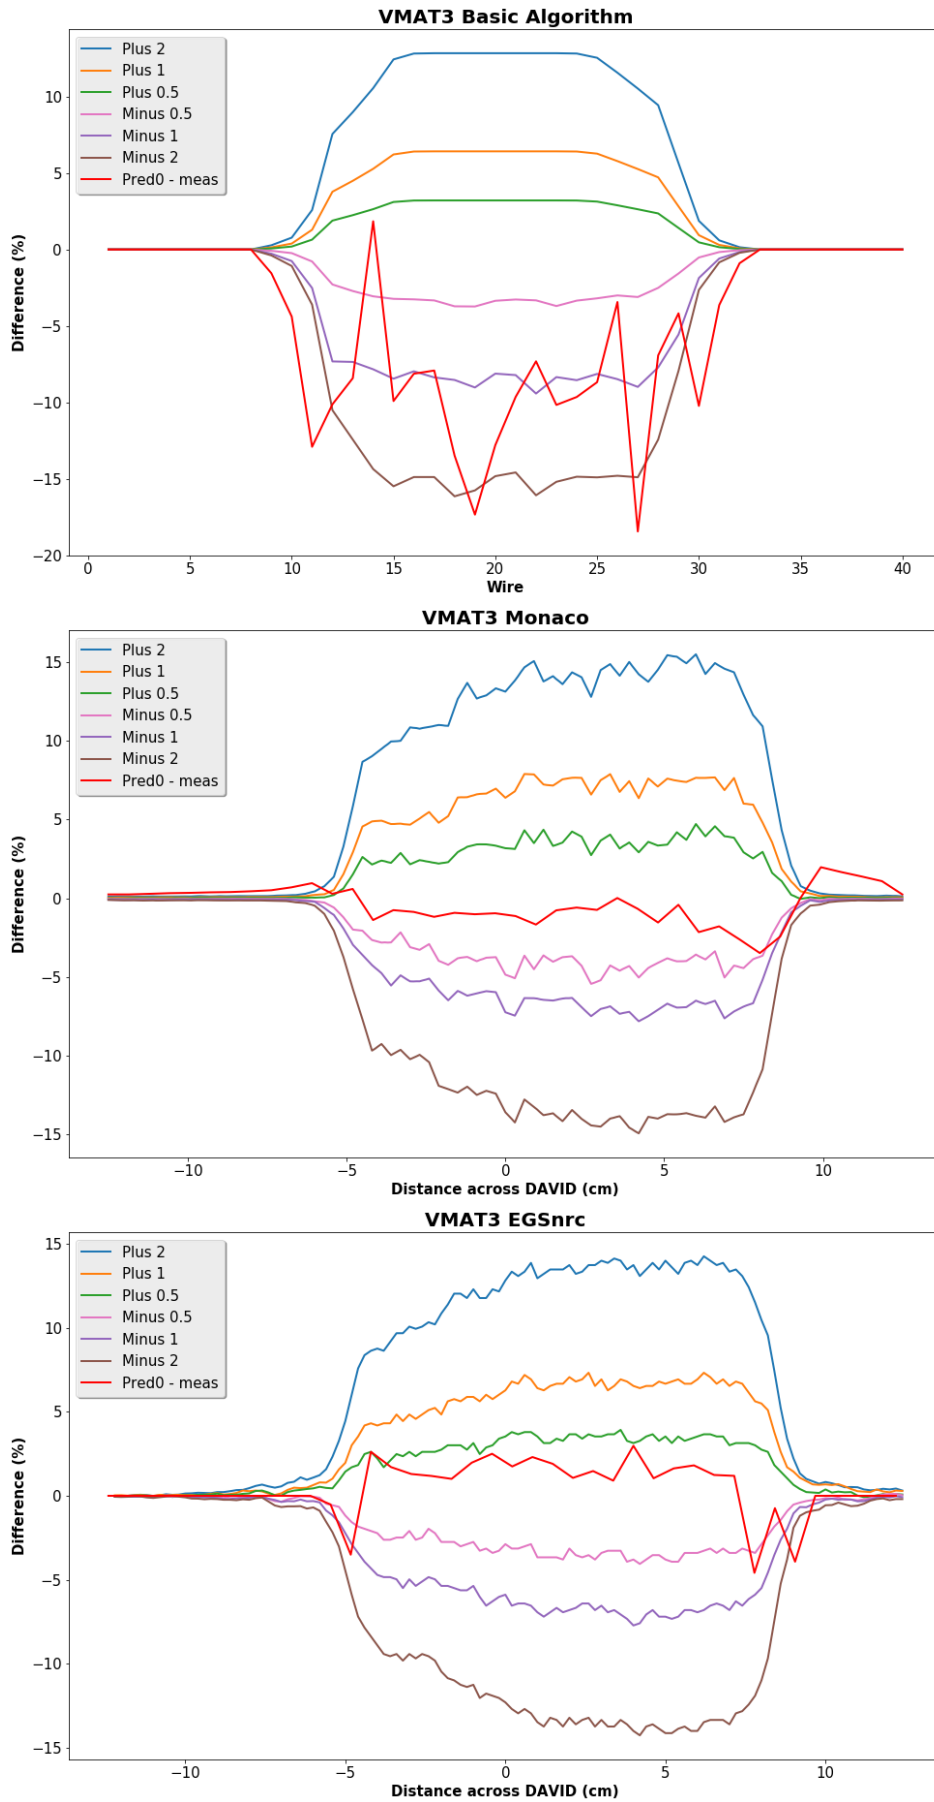
Figure 6.9: Results for all three approaches for VMAT3
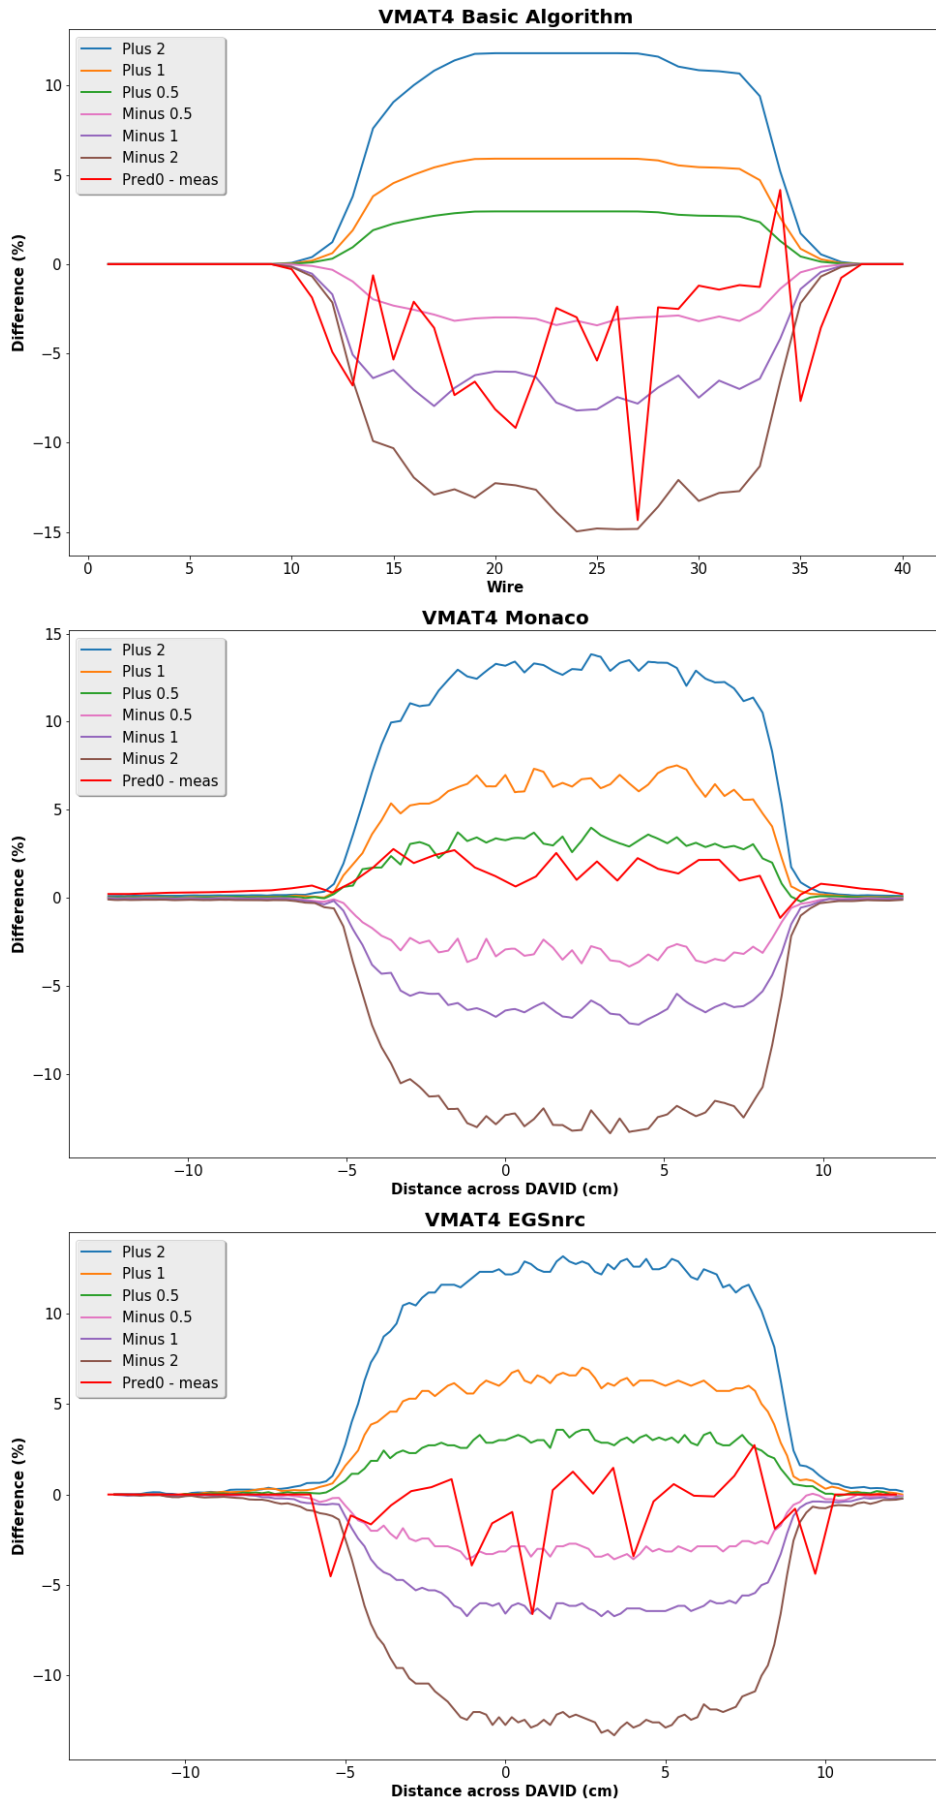
Figure 6.10: Results for all three approaches for VMAT4
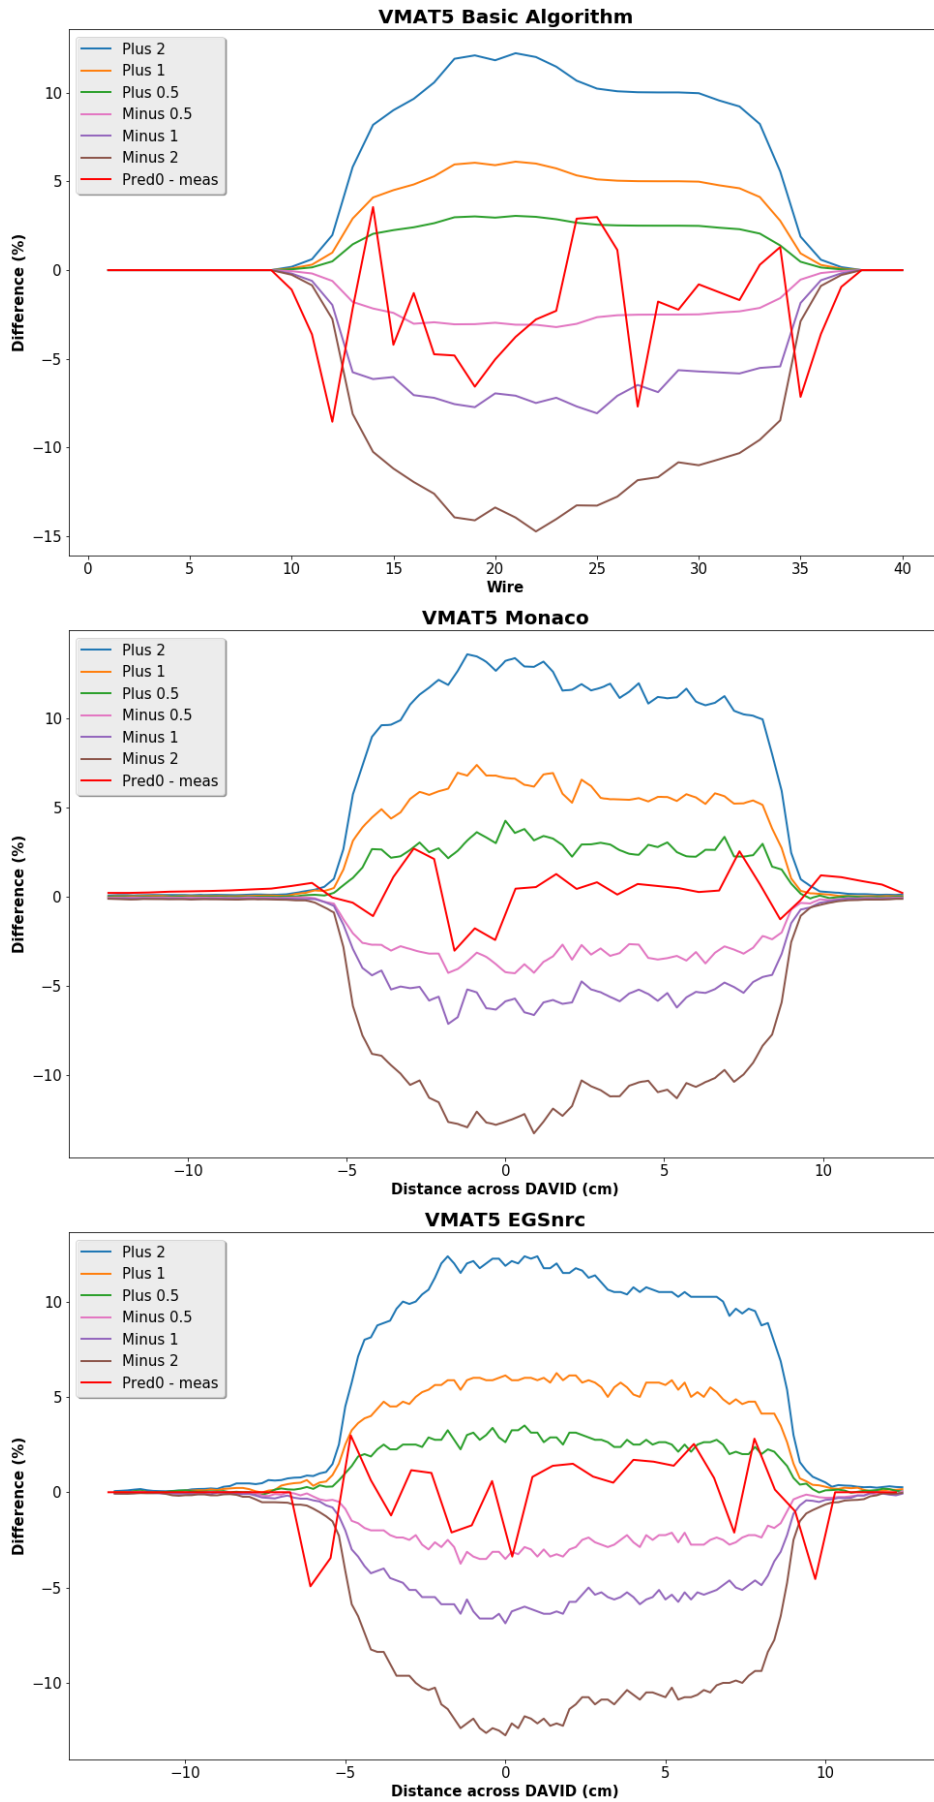
Figure 6.11: Results for all three approaches for VMAT5

	Meas - Pred	Difference between predicted with no shifts and with shifts					
		+2mm	+1mm	+0.5mm	-0.5mm	-1mm	-2mm
VMAT1	-2%	13%	7%	3%	-4%	-10%	-18%
VMAT2	11%	11%	6%	3%	-3%	-7%	-13%
VMAT3	-16%	14%	7%	4%	-4%	-12%	-21%
VMAT4	-7%	14%	7%	4%	-4%	-10%	-20%
VMAT5	-4%	12%	6%	3%	-4%	-9%	-17%

Table 6.1: Basic Algorithm results

	Meas - Pred	Difference between predicted with no shifts and with shifts					
		+2mm	+1mm	+0.5mm	-0.5mm	-1mm	-2mm
VMAT1	0.2%	14.0%	7.0%	3.4%	-3.9%	-7.4%	-15.7%
VMAT2	3.1%	12.0%	6.0%	3.0%	-3.0%	-6.0%	-13.0%
VMAT3	0.9%	14.7%	7.6%	3.9%	-4.9%	-8.1%	-16.7%
VMAT4	-2.7%	15.2%	7.9%	3.9%	-4.2%	-8.5%	-17.5%
VMAT5	-0.9%	13.6%	7.1%	3.5%	-4.3%	-7.3%	-15.0%

Table 6.2: Monaco results

	Meas - Pred	Difference between predicted with no shifts and with shifts					
		+2mm	+1mm	+0.5mm	-0.5mm	-1mm	-2mm
VMAT1	0.1%	14.6%	7.1%	3.4%	-3.4%	-6.5%	-12.7%
VMAT2	0.7%	12.3%	6.0%	3.0%	-2.8%	-5.6%	-10.9%
VMAT3	1.0%	15.4%	7.5%	3.8%	-3.7%	-7.0%	-13.6%
VMAT4	-1.5%	16.2%	7.9%	3.8%	-3.6%	-7.3%	-14.0%
VMAT5	-0.1%	13.9%	6.9%	3.3%	-3.2%	-6.4%	-12.3%

Table 6.3: EGSnrc results

6.4 Discussion

6.4.1 General observations

It is interesting to see that the same shifts had different effects on different plans. In the plus 2mm column in Table 6.3 the total signal from VMAT4 has a 16.2% change from the no-shifts signal compared to 12.3% for VMAT2. As discussed later in this section, the EGSnrc model has the least difference between measured and predicted signals - so for the purposes of this discussion, it will be the only one of the three models considered. To try and explain the varying sensitivity of the different plans to equivalent leaf-bank shifts, the plans were delivered to a phantom in Monaco. The total dose deposited in the phantom was calculated for the no-shifts plans and the +2mm plans. The difference in total DAVID signal using the EGSnrc model (Table 6.3, column +2mm), was plotted as a function of the difference in the total dose delivered to the phantom; these results are shown in Figure 6.12. It can be seen that the difference in DAVID signal correlates with the difference in beam output - not an unexpected result as the DAVID signal has been shown to correlate to delivered dose (Section 3.1.1). However, this still does not explain what it is about the plans that causes this difference.
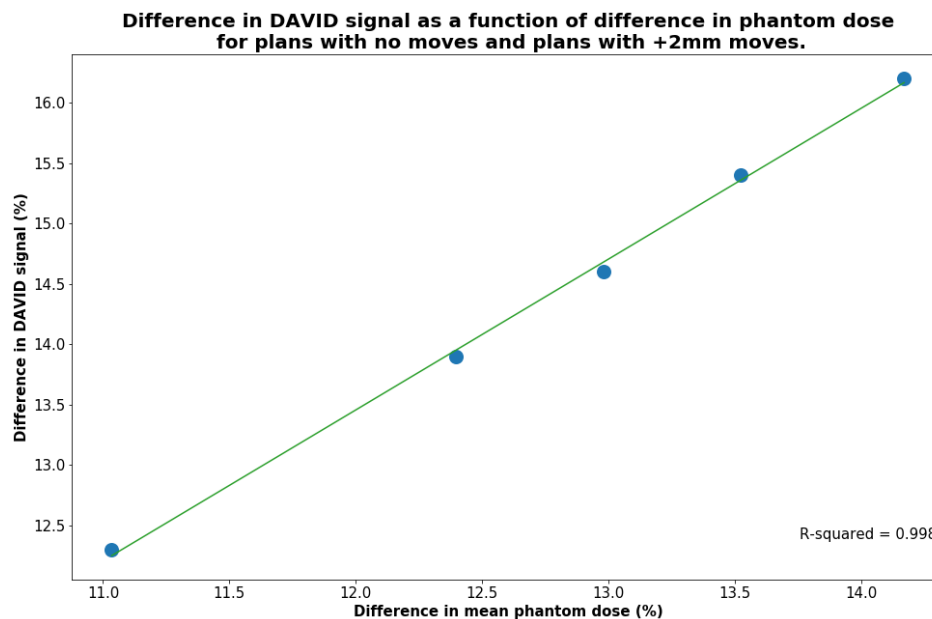
Figure 6.12: Change in EGSnrc-simulated total DAVID signal between original plans and plans with a +2mm leaf bank shift plotted as a function of the change in Monaco phantom dose.

To try and determine why different plans responded differently to equivalent leaf bank shifts, software was written that calculated the area of each segment in the isocentric plane. The average segment area was calculated for each unmodified plan and plotted against the difference in signal seen from a 2mm shift (Figure 6.13). It can be seen that the smaller the average segment area, the greater the effect a leaf bank shift had on the DAVID signal (and, given the results in Figure 6.12, the delivered dose). This result makes sense as the same leaf bank shift will have a proportionally higher change on smaller segments in agreement with LoSasso et al. (1998).

140
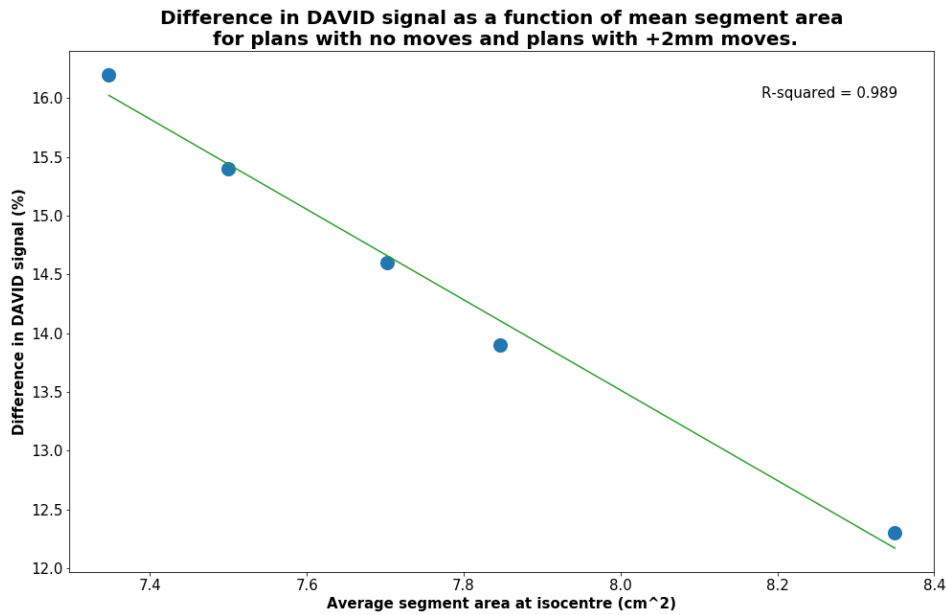
Figure 6.13: Change in EGSnrc-simulated total DAVID signal between original plans and plans with a +2mm leaf bank shift plotted as a function of average segment area.

The graphs shown in Figures 6.7 to 6.11 are an indication of how good each approach is for picking up errors for that specific plan. The smaller the deviations of the red line from zero, the better the prediction matches the actual results. Where the red line crosses a difference line it demonstrates that, in this region, the prediction is not accurate enough to pick up the shift associated with the difference line. Qualitatively it is evident that the Basic Algorithm is the worst of the three as it has the largest spread of the red line in each plan and crosses the most difference lines. Both the Monaco and EGSnrc approaches are considerably better - a narrower spread on the red line, crossing fewer difference lines. A more quantitative assessment is shown Tables 6.1 to 6.3 where the accuracy of the approach is shown in the first column; the lower the magnitude of these values, the better the prediction. The difference between the values in the first column and the other values in the row demonstrates the error-catching ability of the approach - in the case where the error caused by a leaf movement is greater than the difference between the measured and unmodified plans, the uncertainty on the prediction is greater than the difference caused by the simulated error. It has been shown that leaf bank shifts $\sim$ 1mm can have clinically relevant effects on the delivered dose distribution (Section 6.1, (Rangel & Dunscombe 2009, Nithiyanantham et al. 2015, Oliver et al. 2010)). With the large errors in the first column of Table 6.1 it is evident that the basic algorithm lacks the accuracy to pick up clinically-relevant leaf errors.

6.4.2 Determining tolerances for routine use

Figure 6.14 is an example of two deliveries. The green lines represent the difference between measured and predicted signal, the red lines represent the pass / fail tolerance. The delivery in Case A is inside the red lines, so would pass; the delivery in Case B goes outside the red lines, so would fail. The results seen in Case B could

141

be brought about either through an output miscalibration or a leaf bank offset. Determining the position of the tolerance lines is crucial for appropriate use of the device as outlined in Figure 6.16. If the tolerances are too tight, there will be too many false positives and the physics staff will spend time investigating alerts that do not correspond to an actual fault, or faults with no clinical impact. If the tolerances are too slack then errors that may affect the outcome of the treatment may be missed. The other key factor in determining the position of the tolerance line is the accuracy of the predictive technique. The less accurate the predictive technique, the wider the tolerances need to be set so that false positives do not occur too regularly.

The differences between the measured and predicted signals displayed in Figures 6.7 - 6.11 were collected for each predictive approach and are displayed in Figure 6.15. It can be seen that the Basic Algorithm would need considerably higher tolerances than the two MC approaches. Setting the tolerance limits at +/- 2 standard deviations should ensure that 95% of predicted - measured points pass for normal deliveries. By doing this for the results displayed in Figure 6.15 the tolerance limits for the Basic Algorithm were found to be +/- 10% whereas the two MC approaches allowed for a tighter tolerance of +/- 3%.

Looking at the results displayed in Figures 6.7 - 6.11 it is apparent that a tolerance of +/- 10% is not appropriate as leaf bank errors up to 2mm would not trigger an alert. Therefore, while the Basic Algorithm can be used to detect gross errors, it is not accurate enough in its prediction of the signal to be used to detect leaf errors of 2mm. A tolerance of +/- 3% is sufficient to detect leaf bank errors of 1mm and even some 0.5mm errors. Consequently, both MC approaches could be used routinely for the assessment of VMAT deliveries to give confidence that the treatments are being delivered with sufficient accuracy to ensure that clinical objectives are maintained.
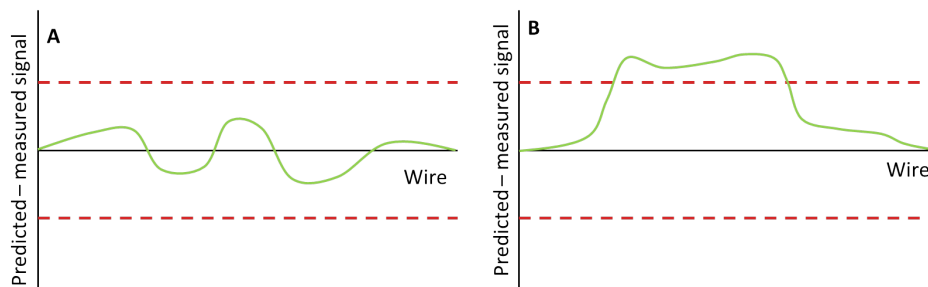
Figure 6.14: An example of two deliveries where the difference between the predicted signal and the measured signal is shown in green; the pass/fail tolerances are shown in red.
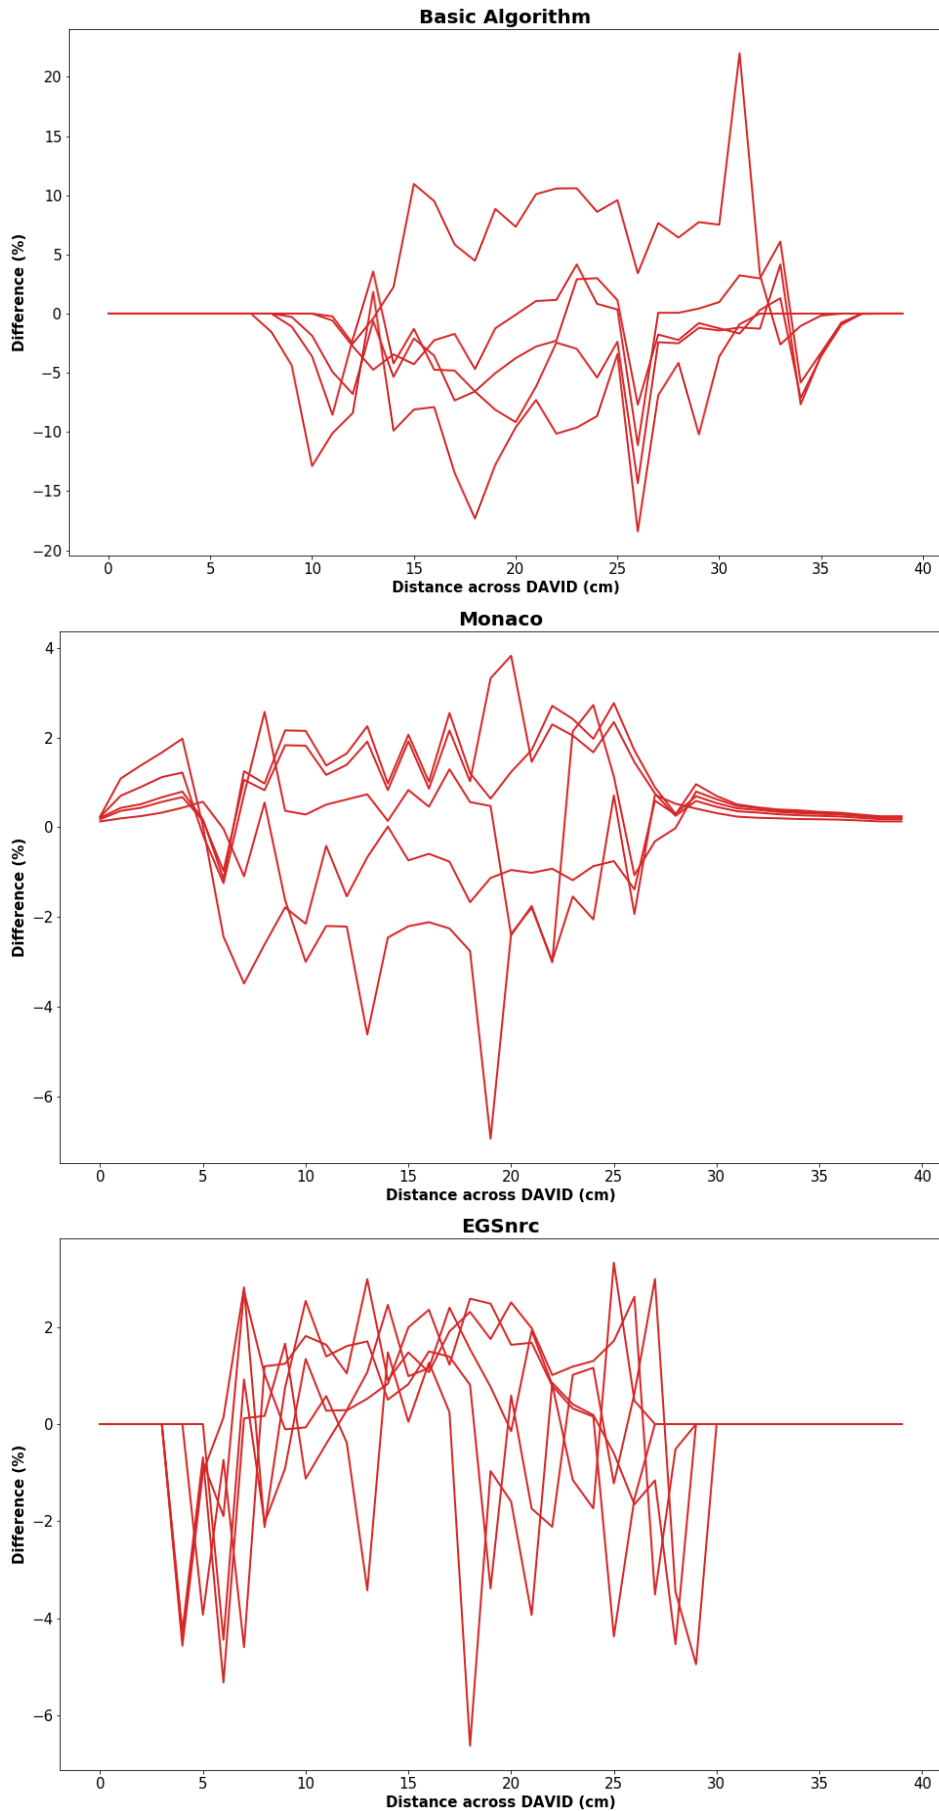
Figure 6.15: Difference between measured and predicted signals for 5 VMAT deliverers for each predictive approach.
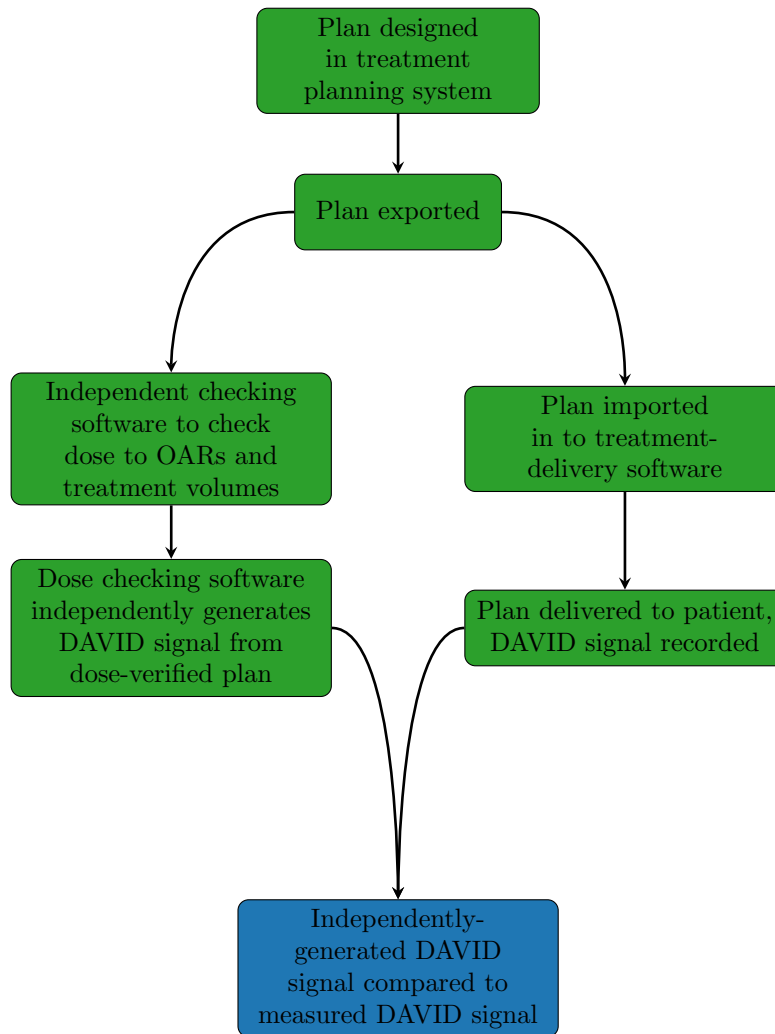
Figure 6.16: Suggested paradigm, reproduced from Chapter 4

6.5   Conclusion

The work in this chapter has demonstrated that the Basic Algorithm is not accurate enough to produce a baseline with sufficiently small uncertainty to detect leaf bank errors that could have an adverse clinical effect. The EGSnrc approach produces the predicted signals with the lowest levels of uncertainty, however the calculation is CPU intensive. Without a computing grid of at least 64 cores, the calculation will take more than two hours and is not really conducive to the efficient day-to-day running of a clinical department.

The Monaco model is easy to set up and quick to run. However, using the TPS to generate its own checking baseline is not a particularly robust approach as it lacks independence and is susceptible to generating false negatives. Not only is it not sensitive to systematic errors in data collection, but it is conceivable that the TPS could generate an erroneous plan and simply generate a correspondingly erroneous baseline. As outlined in Section 4.5.2 the best way of generating a baseline is the use of third-party software that also checks the dose using independently-collected beam data (Figure 6.16).

While EGSnrc could potentially fulfil this role, it is too complex and CPU /

staff-time demanding. Basic checking programs could easily have simple baseline calculations added on, however, this approach lacks the complexity to generate sufficiently accurate signals (Appendix B). It is the opinion of the author that the solution lies in the new generation of checking software. These use more advanced algorithms to check the dose generated by the TPS, in the case of Prosoma a Monte Carlo approach is used that, like Monaco, implements the XVMC (Section 5.4.3) code. Having shown that Monaco can generate accurate baselines, it is suspected that Prosoma could too as, fundamentally, they use the same calculation algorithm. Implementing Prosoma, or an equivalently complex checking program, to check the dose and generate the DAVID baseline signal would retain the independence of third party software. If set up with appropriate diligence, the accuracy of the baseline signal should allow tolerances to be placed on it that would detect clinically relevant MLC errors without too many false positives.

Chapter 7

Conclusions and future work

An outline of the key points from each chapter and how they led to work in subsequent chapters is given in Figure 7.1.

Chapter 3 - *Preliminary work for predictive algorithms.* The DAVID is shown to be a stable device; two algorithms are devised that predict the DAVID response for limited cases. The linear-fit algorithm comfortably outperforms the look-up table approach.

Chapter 4 - *An analytical model for predicting the DAVID signal.* The linear-fit algorithm is developed so that it can predict the response to any field shape for any number of MUs, this allows it to be used to predict the DAVID response to IMRT deliveries. The algorithm is shown to be able to predict the response accurately enough so that output errors >5% are detected and gross errors (delivering the wrong field) are easily detectable. Ideally the device should be able to pick up errors of 3%. It is postulated that a 2D dose map could provide sufficient information to do this.

Chapter 5 - *Monte Carlo Modelling of the DAVID.* An EGSnrc Monte Carlo model of a MLCi2 head is developed and tested then used to show that the effect of the collection wires can be taken into account in a multiplicative correction factor. The DAVID is modelled as two, 4mm slabs of perspex separated by an air gap of 2mm in both EGSnrc and Monaco. The 2D dose map produced in the air gap is exported and used for signal prediction. Both models are used to predict the response to 5 clinical VMAT deliveries. When compared to measured signals both approaches perform well.

Chapter 6 - *The performance of the models in detecting MLC errors using the DAVID.* All three approaches to predicting the signal are tested in the context of clinically relevant leaf errors. The simple algorithm is found to lack the necessary accuracy in predicting the response to impose working tolerances that could detect these errors. Both the MC approaches are shown to produce predicted responses accurately enough so that a working tolerance would detect 1mm leaf errors.

Figure 7.1: Flow of the thesis

7.1   Overview

Chapter 3 demonstrates that the DAVID has a stable and predictable response. This work was built on in chapter 4, where a simple predictive algorithm was developed based on the lateral response function (Figure 4.4) and penumbral modelling (Figure 4.7). The simple algorithm was used to predict the DAVID response for IMRT (Section 4.5.1) and VMAT (Section 6.3) plans. The algorithm was accurate enough to detect gross output errors (Section 4.5.1). However, it lacked the complexity to detect leaf bank offsets ∼ 2mm (Table 6.1) that have been shown to have clinically significant effects on delivered dose distributions for IMRT deliveries as discussed in Sections 1.4 and 6.1.

Figure 4.21 demonstrated that relative leaf position, not just leaf separation, was a factor in the DAVID response to a field. It was postulated that the extra information available in an accurate fluence / dose map would model this effect as well as improving the penumbral modelling (thought to be another shortcoming of the model). An attempt was made to use MU check software to produce fluence maps at the level of the DAVID, these lacked sufficient complexity to accurately predict the DAVID response (Appendix B). So the decision was made to use Monte Carlo methods to produce 2D dose maps for the purposes of signal prediction.

Monte Carlo techniques given sufficient time / computing resources can accurately model coupled photon-electron transport (Rogers 2006, Seco & Verhaegen 2016) and were employed in Chapter 5 to generate a DAVID response. Two models were created, one using the EGSnrc code (Section 5.4.2) and the other using the Monaco treatment planning system (Section 5.4.3). In order to model the DAVID using the EGSnrc code, it was first necessary to build a working model of the MLCi2 linac head (Section 5.3). Once this had been completed, it was shown that the collection wires in the DAVID could be ignored in specific modelling, since their effect was demonstrated to change the response by a constant percentage correction factor. The multiplicative correction, to account for the wires, was built in to the correction factor applied to the simulated dose values to convert them to a predicted signal (Section 5.4.1). This allowed the DAVID model, in both EGSnrc and Monaco, to be modelled as two slabs of Perspex with the signal proportional to the dose measured in the air gap. The two models were verified by comparing the simulated response to five H&N VMAT fields with measured signals (Section 5.5).

Chapter 6 uses the three methods (Basic algorithm, EGSnrc and Monaco) to predict the DAVID response to five complex H&N VMAT plans. The difference between the measured and predicted signal is compared to the predicted signal generated by each method for the same plans with known errors introduced. The Monte Carlo techniques are shown to be capable of predicting the signal accurately enough to distinguish between plans with no errors and and ones with leaf bank errors <1mm. If the signal measured by the DAVID during a delivery matches that predicted by one of the Monte Carlo techniques, then it is implicit that leaf separation is accurate to within 1mm, ensuring appropriate dosimetric accuracy as recommended by IPEM (2018), AAPM (2009), Thwaites (2013), Smith et al. (2017).

7.2 Vendor and beam specificity

The practical work for this thesis was conducted exclusively on an Elekta Synergy linac using a MLCi2 head to deliver flattened 6MV treatment beams. All the parameters in all the models were tailored specifically for this beam, being delivered by this treatment unit. During the time when the work for this thesis was conducted, there were no other treatment units available for research and only the 6MV beams were being used to deliver advanced treatments. Both 6MV flattened and 6MV unflattened beams were being used to deliver IMRT, however, the decision was made to limit the work in this thesis to investigations concerning flattened beams. The models and techniques outlined in this thesis are expected to be broadly applicable to to all modern linacs provided that there is due consideration of a few key points. The following is a qualitative discussion on this topic.

Unflattened beams

The flattening filter (First discussed in Section 1.2.2, EGSnrc component module displayed in Figure 5.4) is a conical beam attenuater, made of a medium or high Z material, located in the treatment head below the primary collimator. The purpose of the flattening filter is to correct for the forward bias of the photon beam that is produced from the target, so that the beam is approximately flat at some depth (normally 10cm). Historically this feature of the beam was useful as it made dose calculations considerably easier when computers were not available or when only simple dose algorithms could be implemented (Georg et al. 2011, Budgell et al. 2016). Removing the flattening filter produces a non-uniform forward-peaked (conical) treatment beam, however, the more advanced algorithms that are currently employed in radiotherapy planning can produce plans with unflattened beams that are equivalent to flattened-beam plans. Removing the flattening filter brings additional advantages:

- Increased dose rate (normally 2 to 3 times higher) resulting in shorter treatment times

- Reduced extra-focal scattered radiation

- Reduced electron contamination in the primary beam

- Reduced leakage from the treatment head

- Reduced susceptibility to beam-steering errors.

Consequently, the number of Flattening Filter Free (FFF) treatments is rapidly increasing (Budgell et al. 2016). What are the consequences for the application of the work in this thesis? The first one is positive: one of the criticisms of the DAVID is that it is insensitive to unilateral leaf bank offsets. If both leaf banks were displaced by the same amount, the apertures would be the same size, so the DAVID response would be approximately the same. Although unilateral shifts have been shown to have limited clinical impact (See section 6.3), the ability to detect them would offer a more complete solution. For beams having a conical dose profile, the DAVID becomes sensitive to these errors as a unilateral offset would result

in a significant change in radiation fluence incident on the detector. Investigation of the DAVID's sensitivity to unilateral offsets in FFF treatment beams could be an interesting topic for future work. The predictive models discussed in this work could potentially take this in to account too.

The basic algorithm discussed in Chapter 4 and published by (Johnson et al. 2014) could be adapted to include a beam profile factor, this would be a function that matched the beam profile and would be applied to wires according to their position in the beam. FFF beam models exist for both EGSnrc (Kajaria et al. 2017, Mishra et al. 2018) and Monaco (Valdenaire et al. 2016, Georg et al. 2011). Once the beams have been set up, including a DAVID model would be the same process as that discussed in Section 5.4.

Different units, vendors and energies

Different treatment units to the ones discussed in this work, either from Elekta or Varian will obviously require different EGSnrc models in order for them to be used to accurately predict the DAVID signal through the methodology discussed in Section 5.4.2. In the case of Elekta units, the deviation from the model architecture discussed in Section 5.3 will be less than modelling Varian linacs. There is published work on EGSnrc models for the flagship linacs from both Varian (Cheng et al. 2016, Song et al. 2012) and Elekta (Gholampourkashi et al. 2019, Oderinde & du Plessis 2016), demonstrating that both are possible. Once a linac unit has been verified at one energy, creating a model for another energy on the same unit will require a new electron spectrum that can be validated through comparison of PDDs, similar to the methodology in Section 5.3.1. If the MLCs and jaws have energy-dependent settings it will also be necessary to determine new beam-shaping scaling factors that can be validated through profile measurements (See Section 5.3.2).

Provided the EGSnrc MC model of the linac is well matched to the treatment beam then the DAVID model discussed in Section 5.4 should work well. A new calibration would need to be determined following the process described in Section 5.4.2. Having done this, extracting the dose and converting it to signal would be the same, the only difference being the location of the collection volumes would need to match those of the DAVID designed for the linac associated with the beam model.

The same can be applied to the Monaco MC model. If the TPS's beam model for the new linac / beam is suitably validated, and a new calibration factor determined (Section 5.4.4) then the DAVID model, and the conversion of dose to signal should work well.

The analytical algorithm discussed in Chapter 4 uses a small number of free parameters. For any new treatment unit / beam these would need to be reestablished. Provided work is done to get sufficiently good estimates of these values (Following a similar workflow to that outlined in Sections 4.2.2 to 4.2.5) then the values could be tuned as described in Section 4.3 to improve the accuracy of the prediction.

Differences in delivery technique, for example Varian RapidArc treatments can either be delivered with static or dynamic jaws (Feng et al. 2015, Wu et al. 2016, Mani et al. 2017), should be accounted for in all models as they base their

prediction on the DICOM plan file. The position of the beam-shaping apparatus is read for each segment and included in the model. In the case of the EGSnrc approach, the leaf and jaw positions are converted into sequence files for the SYNC modules to read (Lobo & Popescu 2010). The Monaco model retains the leaf and jaw positions and uses them when the plan is recalculated on the DAVID model. The analytical approach uses the DICOM values to determine the leaf separation matrix and where to apply the leaf-leakage factor (Section 4.2.4). In the case of static jaws, this factor will then be applied at the same wires for each segment; for dynamic jaws, it will be applied to different wires.

7.3   Basic Algorithm

The only other published work on predicting the signal of a transmission detector with a simple, analytical algorithm was published by Islam et al. (2009) and discussed in Section 2.4.2. Before the work in Chapter 4 there was no published literature on predicting the DAVID signal, nor was there any published work predicting the signal of a head-mounted transmission detector. This makes the content of Chapter 4, the posters displayed at ESTRO 2011 (Appendix D), ESTRO 2012 (Appendix E), ESTRO 2013 (Appendix G) and the publication that this contributed to (Johnson et al. 2014) original. Ultimately the algorithm devised in Chapter 4 is accurate enough to detect gross errors of the magnitude of those discussed in Section 2.2.1, but as demonstrated in Chapter 6 the approach is not accurate enough to predict small ($\sim$ 1mm) leaf errors that have been shown to have a clinical impact on IMRT deliveries (Section 6.1).

In the follow up work to Islam et al. (2009), Pasler et al. (2017) made no mention of predicting the response of the device. It may be that they too found modelling the response of a device to VMAT fields requires the inclusion of more parameters than either model contained.

It is the opinion of the author that an analytical model is capable of predicting the response of a transmission detector to a VMAT field. Inclusion of more parameters could be the subject of future work. It is likely that the algorithm presented in Chapter 4 would benefit from improved penumbral modelling and an approach that accounted for the relative position of the leaves (not just the separation) therefore reducing the discrepancy between predicted and measured signal seen in Figure 4.19. The advantage of such an algorithm would be that the signal could be predicted quickly, with minimal computing power and would require the entry of only a small number of free parameters specific to the linac the device would be used on.

7.4   Monte Carlo models

As stated in Section 5.7, there are a number of EGSnrc-based models of the Elekta MLCi2 head dealing specifically with electron beams (Pitcher et al. 2017, 2016); the Elekta MLC head has been modelled in EGSnrc (De Vlamynck et al. 1999, Van de Walle et al. 2003), as has the Beam Modulator (Asnaashari et al. 2013, Herwiningsih & Fielding 2016, Sikora et al. 2007) and the MLCi2 head has been modelled in Geant4 (Fleckenstein et al. 2013). However, to the author's knowledge

there are, at the time of writing, no published MC models of photon beams for the Elekta MLCi2 treatment head using the EGSnrc code. This makes work outlined in Section 5.3 novel.

While EGSnrc has been used to model transmission detectors (Asuni et al. 2011), this was only to investigate containment electrons. To the author's knowledge a Monte Carlo model of the DAVID has not been produced and there is currently no published literature on predicting the response of a head-mounted transmission detector to a VMAT plan using Monte Carlo techniques. This makes the DAVID models presented in Chapter 5 and displayed as a poster at ESTRO 2016 (Appendix H) completely novel.

In showing that the effect of the collection wires in the DAVID would be accounted for in a general dose-to-signal conversion factor (Section 5.4.1), it allowed a simple two-slab approach to be used in both EGSnrc and Monaco DAVID models. EGSnrc is a full Monte Carlo approach so an accurate model was anticipated. Monaco, however, uses the XVMC code that implements a number of approximations to speed up the simulations for calculating dose in tissue (Section 5.4.3). It is interesting to see that these approximations hold for simulating dose in the distinctly non-clinical conditions (positioning, geometry and materials) of a transmission detector.

While there have been several publications that discuss calibrated Monte Carlo models of linacs (Popescu et al. 2005, Oborn et al. 2014), they tend to focus on correlating the dose scored in the monitor chamber to the linac's output. It appears that the relatively simple approach of generating a calibration curve outlined in Section 5.4.2 has not been published previously and although there are shortcomings to it (Section 5.6.2), the results it yields are accurate. Not only is this novel and useful for future work on Monte Carlo modelling of transmission detectors, it has the potential to be useful whenever a calibrated Monte Carlo linac model is required. It is the author's intention to publish this work in the near future.

By showing that Monaco can predict the DAVID response to VMAT deliveries, it is implicit that the XVMC dose engine is also capable of calculating an accurate baseline signal for the DAVID. The XVMC code is freely available; the subject of future work could be the packaging of software that uses the XVMC code to calculate patient dose and DAVID response to a treatment plan. The software would require a beam model, the data for this would have to be collected independently of the TPS data. Provided the beam model was accurate, if the software was implemented in the way outlined in Figure 7.2 it would provide a robust approach to ensuring that VMAT treatments were delivered with the precision required to deliver the dose accurately (discussed in Section 4.5.1) without the need to acquire a pre-treatment DAVID baseline on the linac.

7.5   Summary

In summary, the work was intended to investigate the characteristics of the novel 2-D transmission detector, the DAVID, with the intention of being able to predict its response to radiotherapy deliveries. The aim was to eliminate the need to collect a pre-treatment baseline signal for the device on the linac. The approach was to directly predict signal for individual treatment deliveries. This would make

Figure 7.2: Suggested paradigm, reproduced from Chapter 4

individual patient QA and treatment verification faster and more efficient, as well as independent and therefore safer. By developing and testing predictive models of increasing complexity and creating an appropriate methodology to use this practically, this aim was clearly met. A robust approach to ensuring that advanced and complex treatments, e.g. VMAT, are delivered with the precision required to deliver the dose accurately without the need to acquire a pre-treatment DAVID baseline on the linac was demonstrated.The Monte-Carlo based models perform well in identifying leaf errors within tolerances of $\sim$ 1mm; one of these techniques works with the relatively widely-used MC-based TPS, Monaco making it easy to implement in the clinic. In addition the work presents a suggested clinical workflow for application in daily use to monitor patient treatment. This work can provide the basis of fast and efficient intra-fraction monitoring of complex radiotherapy delivery, that can be further developed for appropriate workflows to support adaptive radiotherapy verification and QA.

Appendix A

Gamma analysis

A.1 Introduction

Gamma analysis is a popular tool that allows the comparison between two, 3D or 2D dose distributions; the original paper describing the technique (Low et al. 1998) has, at the time of writing, been cited over 2000 times. Although the approach has been criticised for not picking up some clinically relevant dose discrepancies (Zhen et al. 2011, Kim et al. 2014, Ohira et al. 2017, Caivano et al. 2014), some form of gamma analysis a feature on most commercial verification analysis software platforms (Hussein et al. 2013b) and it remains in widespread use throughout radiotherapy departments (Hussein et al. 2017).

A.2 Methodology

Gamma analysis compares a reference dose distribution with an evaluated dose distribution in terms of both geometric distance and dose difference between equivalent points. Taking the example of 2D gamma analysis, the reference dose point can be thought of as being at the centre of an "acceptance ellipsoid". Gamma criteria is normally given in terms of $\delta D\%/\delta S mm$ where δs is the distance between reference and evaluated dose points, δD is the dose difference as a percentage of some value. In the case of global analysis this value is normally the maximum dose, if it is local analysis it will be the reference point. The acceptance ellipsoid is centred on a Cartesian coordinate system, its origin at the reference dose point with spatial dimensions for the x and y axis and dose on the z axis. The gamma criteria determines the shape of the acceptance ellipsoid (Figure A.1), when comparing a reference and evaluated dose point, a gamma index of <1 is assigned if the evaluated point is inside the acceptance ellipsoid and >1 if it is outside. This is repeated for all points in the evaluated distribution and the result is normally quoted as the percentage of points with a gamma index <1 along with the gamma criteria, for example, at James Cook university Hospital the passing criteria for pre-treatment verification is 98% at 3%/2mm.
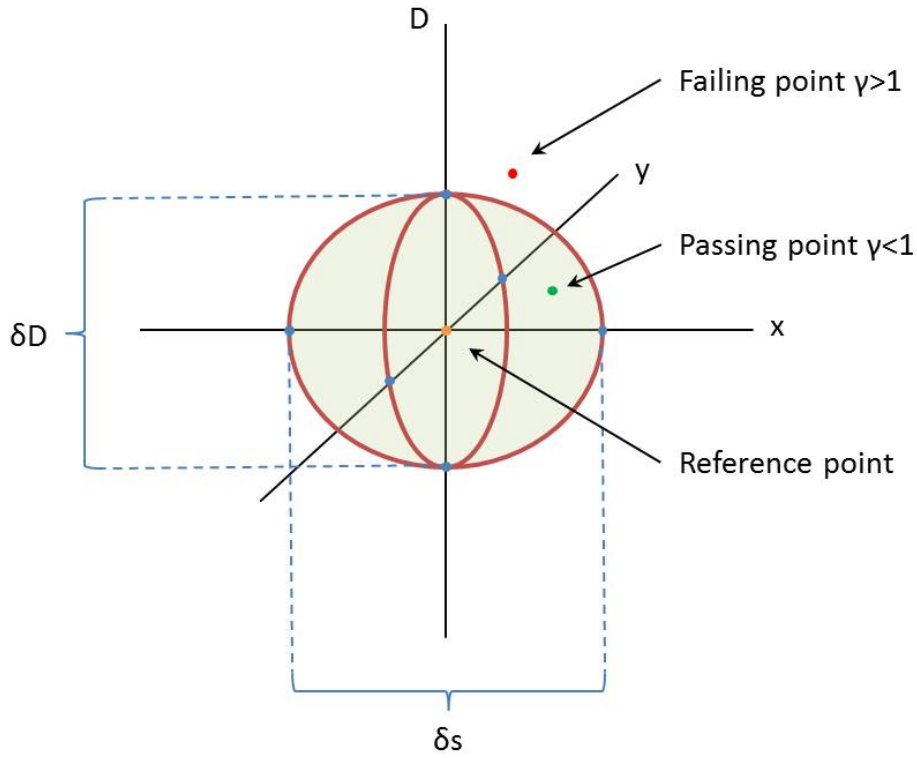
Figure A.1: Acceptance gamma ellipsoid around a reference point. 2D gamma analysis compares the two dose distributions in the xy plane.

More quantitatively, the gamma index between a reference point (r_R) and an evaluation point (r_E) is given by

$$\gamma(r_R, r_E) = \sqrt{\frac{\Delta r^2(r_R, r_E)}{\delta r^2} + \frac{\Delta D^2(r_R, r_E)}{\delta D^2}} \tag{A.1}$$

Where δr and δD are the distance and dose criterion respectively. ΔD is can be defined for local gamma analysis (Equation A.1) or global gamma analysis (Equation A.2).

$$\Delta D(r_R, r_E) = D_E(r_E) - D_R(r_R) \tag{A.2}$$

$$\Delta D(r_R, r_E) = \frac{D_E(r_E) - D_R(r_R)}{D_{m}ax} \tag{A.3}$$

(Low et al. 1998, Hussein et al. 2017)

154

Appendix B

DIAMOND software evaluation

B.1 MU checking software

There have been a number of documented radiotherapy errors, these have ranged in severity and consequence. Although these are terrible events and often, in retrospect, better risk management could have avoided them, they give us the opportunity to learn. From such errors, we can develop better strategies for avoiding risk in the department. For example, in 2000 physicists in Panama calculated the MUs for several treatments using a TPS. At the time, the centre was using apertures created from drilled holes in lead. Unfortunately, due to a glitch, the software interpreted what was meant to be an aperture as lead, assigning it the wrong density. Consequently, the number of monitor units calculated was much higher than required. Twenty-eight people were over-exposed. At the time the report was published, five people had died and it was expected that 15 more would "develop serious complications, which in some cases may ultimately prove fatal" (Vatnitsky et al. 2001).

The report into the incident made three recommendations. The first was: "ensuring that the procedures require independent (manual) verification of the monitor units (or irradiation time) and doses to the prescription points as calculated by the TPS for each individual patient, before the first treatment starts" (Vatnitsky et al. 2001).

In the UK the practice of checking monitor units is now recommended, "For plans generated by treatment planning computers (producing isodose distributions resulting from the combination of two or more beams), calculation of the monitor units and dose to the reference point must be independently checked either by hand using tabulated data or by using another computer program" (BIR 2007). The check is in place to ensure that the planning system has calculated the dose correctly. Traditionally for 3DCRT where, typically, <7 beams were used this, could be done by hand. The dose checker - normally a medical physicist – would use information like the field size, beam energy and the depth of the treatment to determine a number of correction factors from look-up tables, based on independently measured linac data. The correction factors would be used in conjunction with the beam MU to determine the dose in the patient for that beam. Due to

the different calculation methods, some difference between the TPS and checking method of dose calculation is to be expected (AAPM 2011, BIR 2007). It is recommended that in the event of the difference in dose being >2%, the data is rechecked. In the case of the difference being >5%, an alert is raised (BIR 2007).

In the cases of more advanced treatments – IMRT and VMAT – where a single beam could be made up of a number of small segments, hand calculations became infeasible due to the time it would take to perform them. Predictably, with radiotherapy moving increasingly towards advanced treatments, a number of commercially available checking programs started to appear (Radcalc (Lifeline Software Inc., Tyler TX), Diamond (PTW-Freiburg), IMSure (Standard Imaging, Middleton, WI), and Mucheck (Oncology Data Systems, Inc), to name a few). These programs typically consist of a method to import dicom plans and a simple dose calculation algorithm (AAPM 2011, 2014, Tuazon et al. 2018). The user is able to import a plan and patient structure set, and where necessary, assign densities to the structures and calculate the dose at a point of their choosing. They can then compare the dose at this point to the dose calculated by the TPS. While designed to make dose checks on advanced plans feasible in the radiotherapy department, these programs can also be applied to 3DP and can save a lot of time. On account of this, it is hardly surprising that some sort of checking program is now a standard feature of UK radiotherapy centres (IPEM 2008) and recommended by international bodies too ICRP (2009).

DIAMOND is a checking algorithm owned by PTW and supplied for the purposes of this work. DIAMOND is not only capable of calculating the dose at a point, but it can also generate and export fluence maps of the treatment. It was hoped that, after commissioning the software, the fluence map could be used to determine the DAVID signal, thereby improving the accuracy of the algorithm discussed in Chapter 4 by addressing the problems outlined in Section 4.6.

B.2 Commissioning DIAMOND

B.2.1 Data collection

Introduction

DIAMOND calculates dose based on a modified Clarkson algorithm; this approach determines the scatter at a point using sector integration. Figure B.1 shows a irregular radiation field (A), this can be split up in to n segments; one of these is shown in red. The dose due to scatter at point P can be said to be equal to the scatter contribution from the n segments. The scatter dose from one segment can be calculated as follows:

The scatter-air ratio (d, r_d) at depth d and radius r_d can be described by the tissue-air ratio $T(d, r_d)$, with contributions from both scatter and primary radiation minus the primary contribution $T(d, 0)$[1] (Equation B.1). Values of (d, r_d) are tabulated and stored for reference.

$$S(d, r_{\mathrm{d}}) = T(d, r_{\mathrm{d}}) - T(d, 0) \tag{B.1}$$

---

[1] primary contribution in this instance is defined at zero field size - where there is no scatter, this can be determined by taking measurements with increasing smaller field sizes and extrapolating back to zero.

The contribution of segment r_1, θ in Figure B.1 is simply the ratio of its area to the area of the circle (r_1), multiplied by the $S(d, r_d)$ of circle (r_1) (Equation B.2). This can be read, or interpolated, from the tabulated data.

$$S(d, r_i, \theta_i)_{\text{segment}} = \frac{\theta_i}{2\pi} \times S(d, r_i) \tag{B.2}$$

The scattered dose at point P ($D_{scatter}$) from the sector is simply the scatter-air ratio multiplied by the dose at that point in air D_A (Equation B.3) that can be determined from the SAD, air output factor and collimator scatter factors.

$$D_{\text{scatter}}(\theta_i) = D_A(d, r_i, \theta) \times S(d, r_i, \theta) \tag{B.3}$$

Making total scatter dose:

$$D_{\text{scatter}} = D_A(d, r_i, \theta) \times \sum_i \frac{\theta_i}{2\pi} \times S(d, r_i) \tag{B.4}$$

Primary dose is calculated using Equation B.5, where D_A is, as above, the dose in
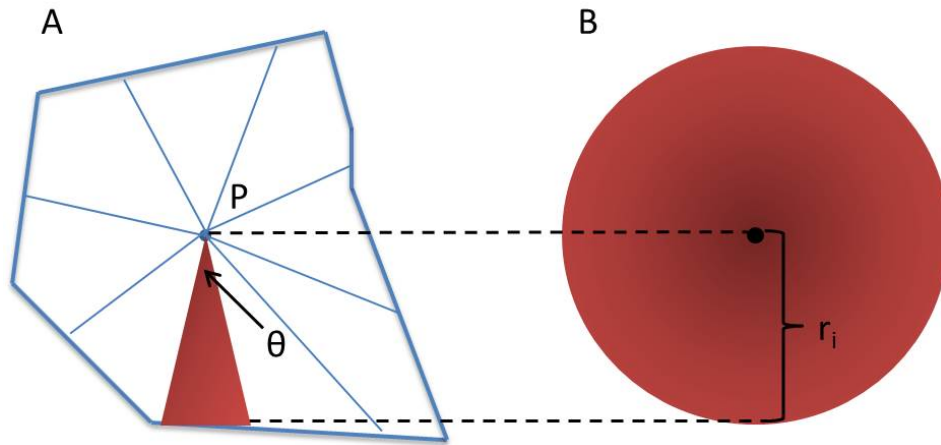

Figure B.1: The scatter-air-ratio of the segment at point P is equal to the scatter-air ratio of a circle of equal radius multiplied my the area ratio of the segment and circle (Cunningham et al. 1972).

air, $T(d,0)$ is the tissue-air ratio for the case where there is no scatter and $f(x,y)$ is a off-axis factor that corrects for changes in the beam profile (Khan 2010, Mayles et al. 2007, Cunningham et al. 1972).

$$D_{\text{primary}} = D_A) \times T(d, 0) \times f(x, y) \tag{B.5}$$

This approach of primary-scatter separation provides a computationally lite approach to dose calculation. It was developed further to be used with 3D data sets and cope with tissue inhomogeneities (Sontag & Cunningham 1978, Redpath & Thwaites 1991), providing a useful tool for calculating the dose at a point for irregular fields. The advent of IMRT made meant that the assumption made with sector-integration, that the dose was homogeneous was no longer valid. However, Kung et al. (2000) developed a modified Clarkson algorithm relying on annular integration. Three equally-sized apertures, with equal MU, at different points on

a circle around a central axis will contribute equally to the dose measured at that point. This concept can be extended further: instead of considering an IMRT field as a complicated fluence distribution on a Cartesian axis, it can can be though of as a set of concentric intensities centred on the CAX (Figure B.2). The intensity of each concentric ring is the azimuthal average of the intensity at each radius. The dose is calculated in a equation analogous to Equation B.4. However, instead of summing over 2π sectors, the sum is between 0 and R annular sectors. Varying surface contours can be accommodated for by applying attenuation corrections to incident fluences, and tissue inhomogeneities can be accounted for by using radiological depths in the place of geometrical depths (Kung et al. 2000).

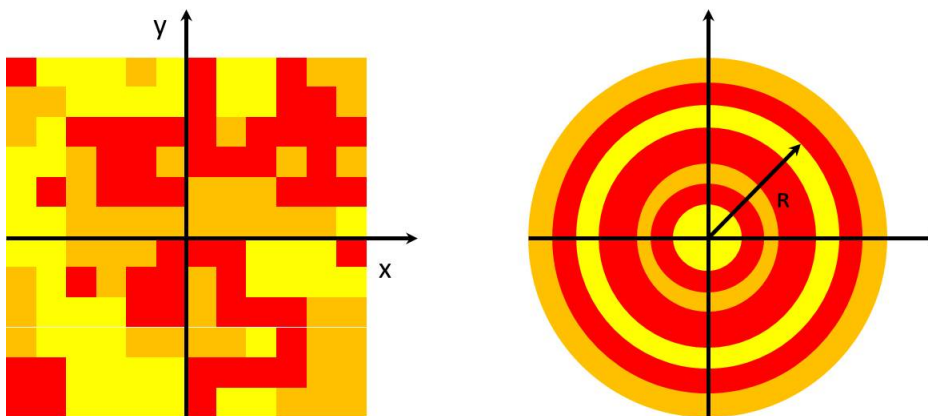This algorithm, employed by DIAMOND and RadCalc, uses patient data, the

Figure B.2: An IMRT fluence distribution on the left can be converted to an annular distribution (right) with the intensity of each ring equalling the average intensity at distance R in Cartesian space(Kung et al. 2000)

treatment plan (MU, MLC positions, gantry angle etc.) and beam data to calculate the dose at a point. The patient data is in the form of structures that are outlined on the CT images in the TPS and imported with the treatment plan. The beam data needs to be collected from the linac. To ensure that the dose check is completely independent, the data was, in accordance with recommendations (National Cancer Peer Review Programme 2014, BIR 2007, AAPM 2011), collected separately from the TPS data.

Method

Depth dose data Figure B.3 shows the water tank; this was filled with water, levelled and moved to 100cm SSD. The semi flex chamber (31010, PTW-Freiburg, Germany) was mounted, in the horizontal position (Figure B.3), on the arm and centred on the beam axis. PDDs were taken for a range of field sizes between 1×1cm and 40×40cm. For field sizes 5cm^2 and below, the data was re-collected with the photon diode (T60012, PTW-Freiburg, Germany). Below field sizes of 4cm^2, the larger semi flex chamber suffers from volume-averaging effects, making the much-smaller diode a better device for small fields (AAMP (2008), Griessbach et al. (2005), Das et al. (2008)). By taking 4cm^2 and 5cm^2 PDDs with both chambers, where volume-averaging effects suffered by the semi flex would not have an effect and the fields were not big enough to adversely affect the diode's

performance, it was shown that the PDDs were equivalent. Having shown the equivalence of the diode and semi flex, the diode was used to collect the data for field sizes less than ≤ 4cm2. The PDDs used in the software are shown in Figure B.4
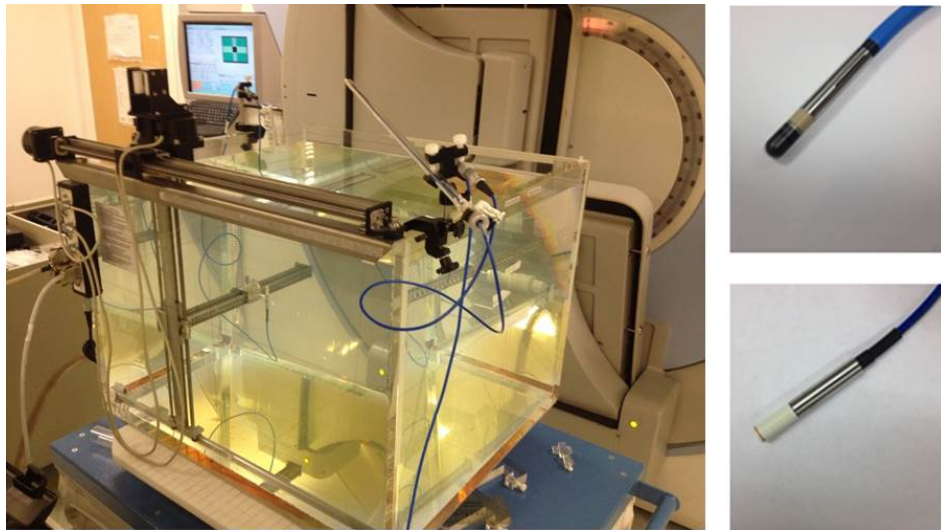

Figure B.3: The PTW water tank set up for data collection with a semi flex being used as the reference chamber and the diode being used as the field chamber. Top right: semi flex chamber, bottom right photon diode.

Off axis data. This data was taken using the plotting tank at an SSD of 100cm. A semi flex chamber, in the vertical position, was used to measure a 40×40cm profile at depths of 1.6cm, 5cm, 10cm, 20cm and 30cm in 1mm steps (Figure B.5).

MLC data. This was collected with a semi flex chamber in the vertical position, the chamber was moved in 1mm steps across a 4×4cm field in the direction of the MLC movement (Figure B.6).

Head scatter. The SCP and SC values were collected in a water tank and in a mini phantom, respectively (IPEM 2018). All values were collected with an isocentric (SAD) set up with the detector at a depth of 10cm (SSD = 90cm)(Figure B.7)

Leaf transmission. A value of 2% was used – the same as the planning system (AAMP (2008)).
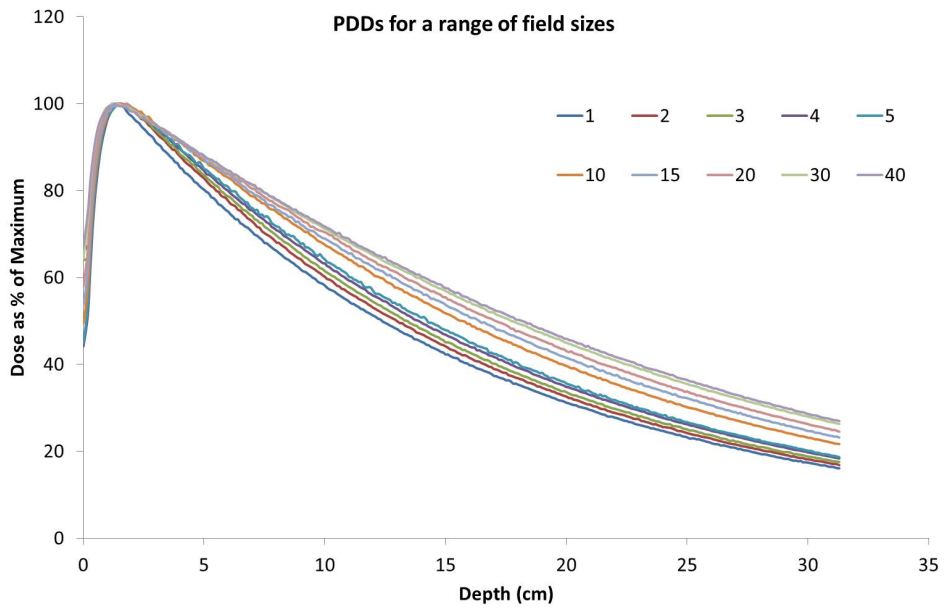
Results

Figure B.4: PDDs for a range of field sizes.
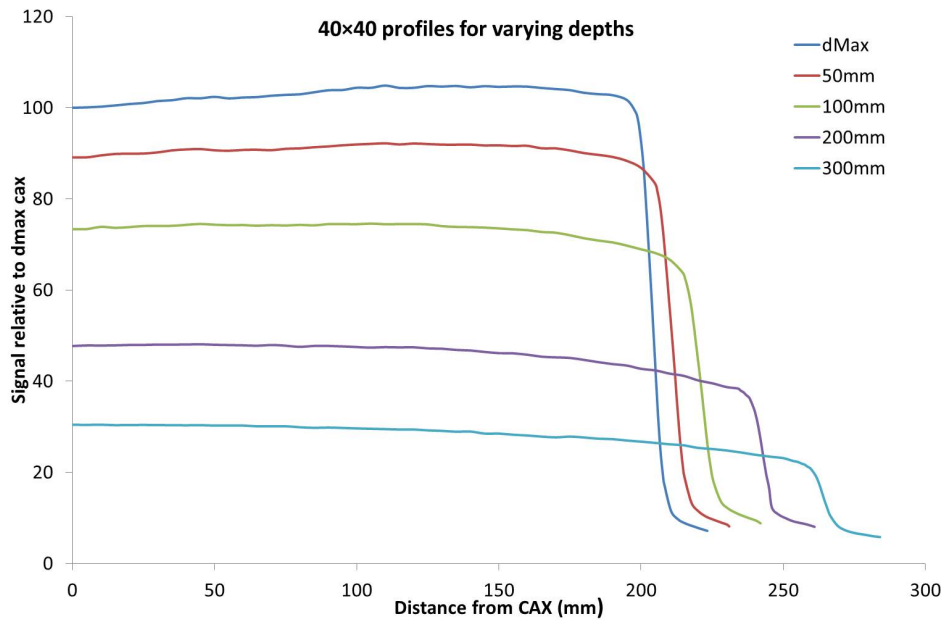

Figure B.5: 40×40cm profiles in measured in the direction perpendicular to leaf travel for a range of depths.
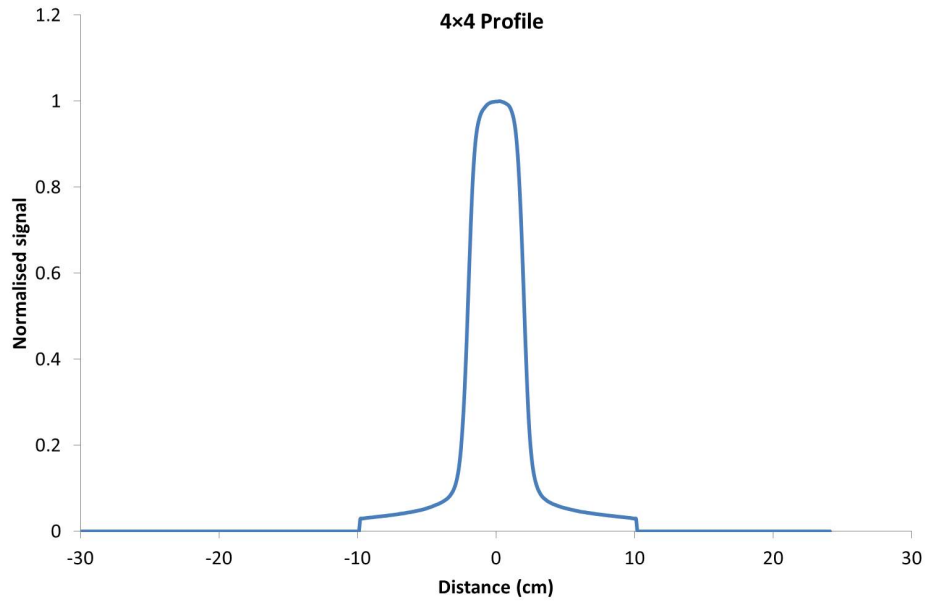
160

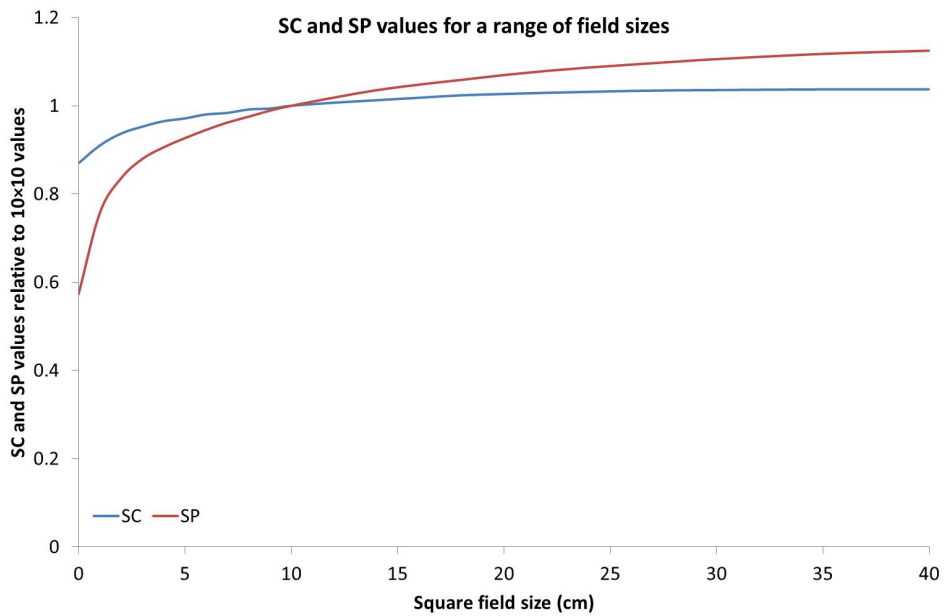Figure B.6: Profile of a 4×4cm field in the direction of MLC travel.

Figure B.7: SC and SP values for a range of field sizes extrapolated back to 0cm^2.

B.2.2   Array Calibration

Introduction

To benchmark the DIAMOND software against measured dose (Section B.2.3), planer measurements of plans delivered on the linac were taken. This was done using the OCTAVIUS phantom (PTW Fireburg-Germany) with the Seven29 2D array inserted into it. The Seven29 array consists of a plane matrix of 27×27 cubic, air-filled ion chambers. The vented plane-parallel ion chambers are $5 \times 5 \times 5 \text{mm}^3$ in size with a centre-to-centre spacing of 10mm. The upper electrode layer sits below a 0.5cm a PMMA build-up layer, whereas the lower electrode layer lies on top of a 0.2cm thick electrode plate, which itself is mounted on a 1cm PMMA base plate (Bohsung et al. 2004, PTW 2010, Hussein et al. 2013a, Van Esch et al. 2007). The Seven29 array has been shown to give reproducible results in the short, medium and long term with a response that is linearly proportional to energy and dose, showing that it can be considered an accurate and sensitive QA tool (Spezi et al. 2005). The Seven29 array has a varying angular sensitivity - this is because of the increased attenuation from the electronics behind the chambers that is not present on the front of the device. When used in conjunction with the Octavius phantom, this is effect accommodated for by an open void in the PMMA behind the array (Figure B.8). The drop in attenuation from the void compensates for the blocking of the beam by the electronics. The combination of the Octavius phantom and the Seven29 array has been shown to be a useful and appropriate tool for performing pre-treatment verifications on arc-style deliveries (Van Esch et al. 2007).
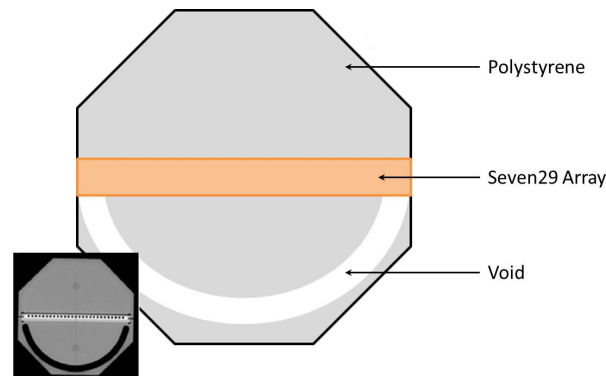
Figure B.8: Schematic of the Octavius phantom with transaxial CT slice as inset. Polystyrene = physical density 1.04 g/cm^3, relative electron density 1.00, width and length: 32cm (Van Esch et al. 2007)

Method

In an approach similar to Van Esch et al. (2007), the array was placed on the couch and centred on the beam axis. Increasing amounts of solid water were placed on top of it, dropping the couch each time to maintain a constant SSD. In this way, a PDD was measured with the central chamber. This procedure was repeated, but in place of the array, a farmer chamber was used. The farmer chamber was a field chamber with a traceable calibration (Lillicrap et al. 1990), so the measured dose could be calculated. The procedure was repeated three times for both the array and the chamber. The resulting PDDs were compared, and a factor applied to shift

the array measurements so that they matched the chamber measurements (Figure B.9). This factor was applied to subsequent array measurements, converting them to measured dose.

As both the farmer chamber and the 2D array chambers are vented, temperature and pressure corrections were applied to all readings.

Results

After applying the calibration factor, the chamber and array PDDs were seen to match well (Figure B.9), showing the process to be appropriate. The calibration factor was found to be:

2D array calibration factor = 1.0105

Although the Seven29 array is issued by the manufacturer in a calibrated state, meaning that the calibration factor should be unity, some drift is expected. A change in calibration factor of about 1% is consistent with the evaluation performed by Spezi et al. (2005).
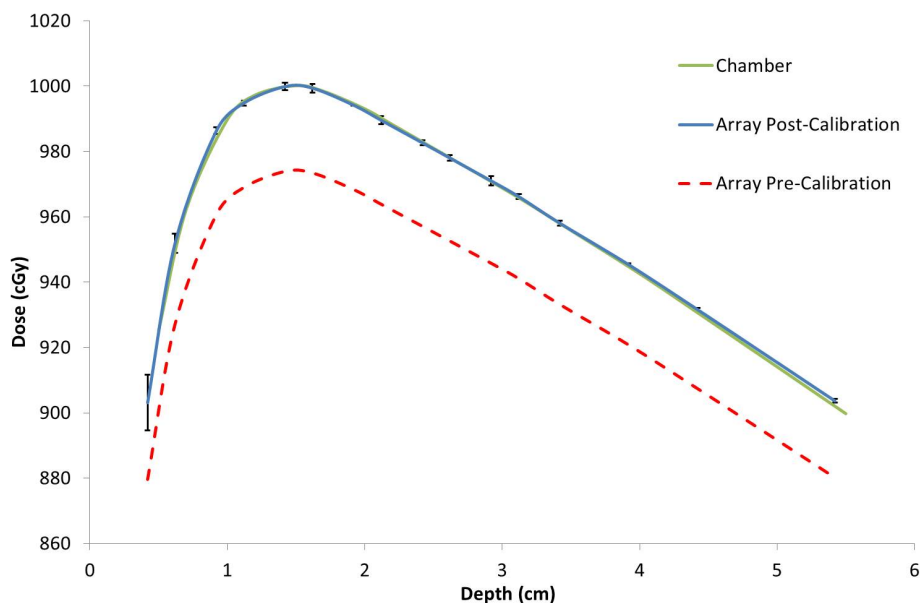
Figure B.9: PDDs from a chamber, post-calibration array PDD and pre-calibration array PDD – the discrepancy between the pre- and post-calibration PDDS has been exaggerated to illustrate the process.

B.2.3   Testing

Introduction

To test the DIAMOND software, the dose it calculated was compared with measured dose and dose calculated by two other systems. Having access to Monaco as a TPS provided the opportunity to compare the DIAMOND calculation to a Monte Carlo algorithm (Section 5.2). RadCalc is the checking software currently used in the department. Like DIAMOND, it uses a Clarkson-based algorithm to calculate dose; comparison between the two will benchmark DIAMOND against a

piece of software with an equivalent role and approach. Finally, the dose will be measured on the linac using the Seven29 array (Section B.2.2).

Method

To compare measured and calculated doses, a similar process to pre-treatment patient-specific QA (Alber et al. 2008, Stathakis et al. 2013, Ezzell et al. 2003, AAPM 2018) was followed. The OCTAVIUS phantom with the Seven29 array inserted (Figure B.10), was imaged on the CT scanner. This image set was imported into the TPS, an external contour was drawn around the image of the phantom and the electron density forced to 1. A treatment plan was then imported and the dose calculated on the phantom. Typically, for pre-treatment verification, commercial software is used to compare the dose plane at the level of the detectors calculated in the TPS to the dose measured by the detectors, normally using a 2D gamma analysis with a locally decided passing criteria (Alber et al. 2008). However, in this instance, the software being tested can only calculate dose at a point. So, in the dose calculated in the TPS, four points in the plane of the detectors were recorded - the position of these points is shown in Figure B.11. The plan and structures were exported to DIAMOND and RadCalc, and the dose was calculated at the same four points in each.

The plan was then delivered to the phantom / array on the linac and the dose measured at the four points was recorded. The TPS, RadCalc, DIAMOND and measured doses for all four points were compared for 10, clinical, step-and-shoot IMRT plans.
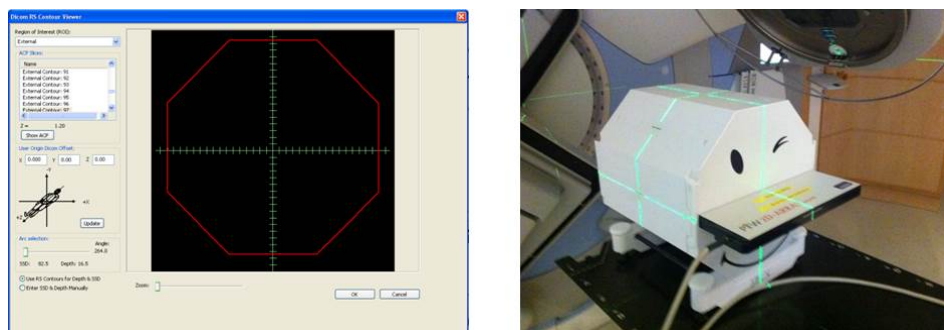

Figure B.10: The OCTAVIUS structure set in DIAMOND and The OCATAVIUS phantom with the 2darray inserted on the linac couch.
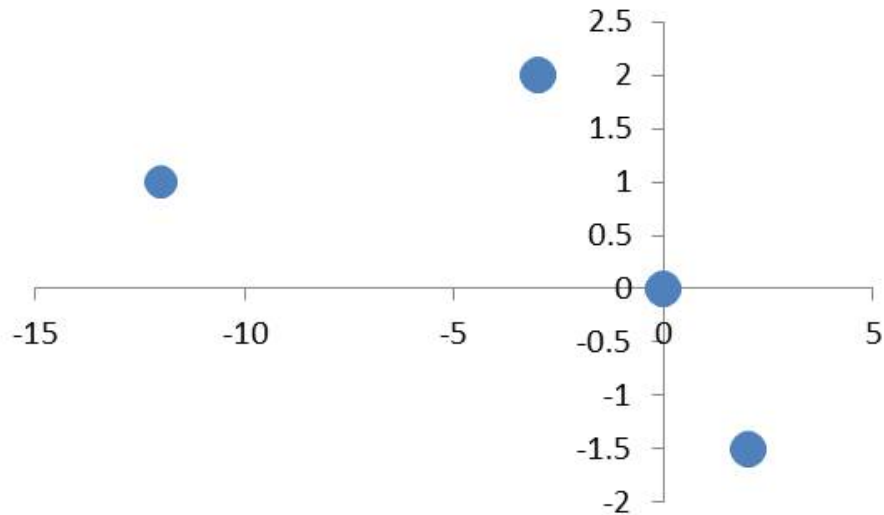
Figure B.11: Dose measurement points in the plane of the detector / coronal plane of the plan. Distances are in cm.

B.2.4 Results

Figure B.12 shows the absolute values determined by the four techniques. Figure B.13 shows the difference between the dose calculated by Monaco and the three other techniques. The differences are typically within 5% and the DIAMOND results are equivalent the clinically used RadCalc. This is consistent with the data published by Tuazon et al. (2018) and can be attributed to the simplistic nature of the Clarkson dose-calculation algorithm. In some cases, the differences are greater then 5%. This can be ascribed to the arbitrary selection of measurement point in fields with high dose gradients. The points were the same for all plans; in the event of performing these calculations in the clinic, the checker would choose a point with a low dose gradient eliminating the errors that this causes.
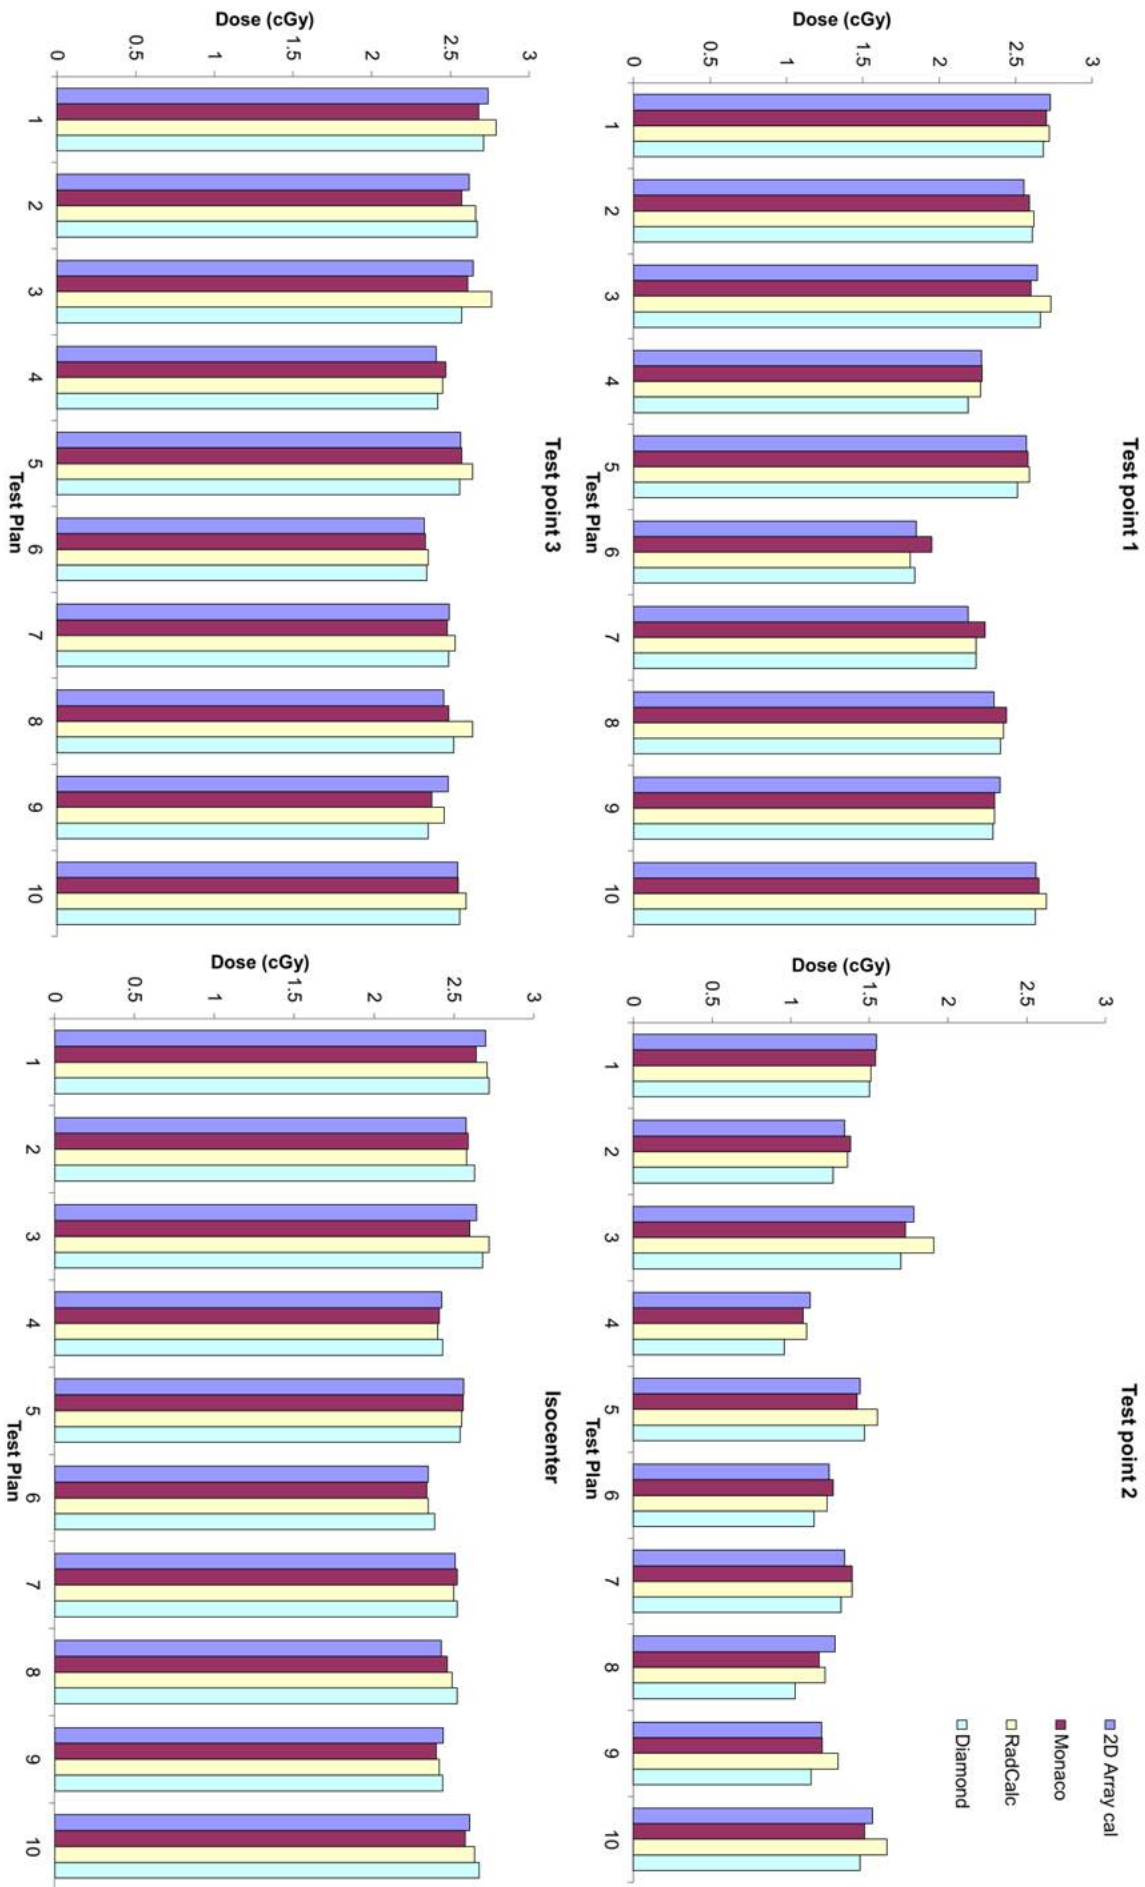
Figure B.12: Absolute dose measured / calculated at the four measurement points (B.11) for all 10 plans.
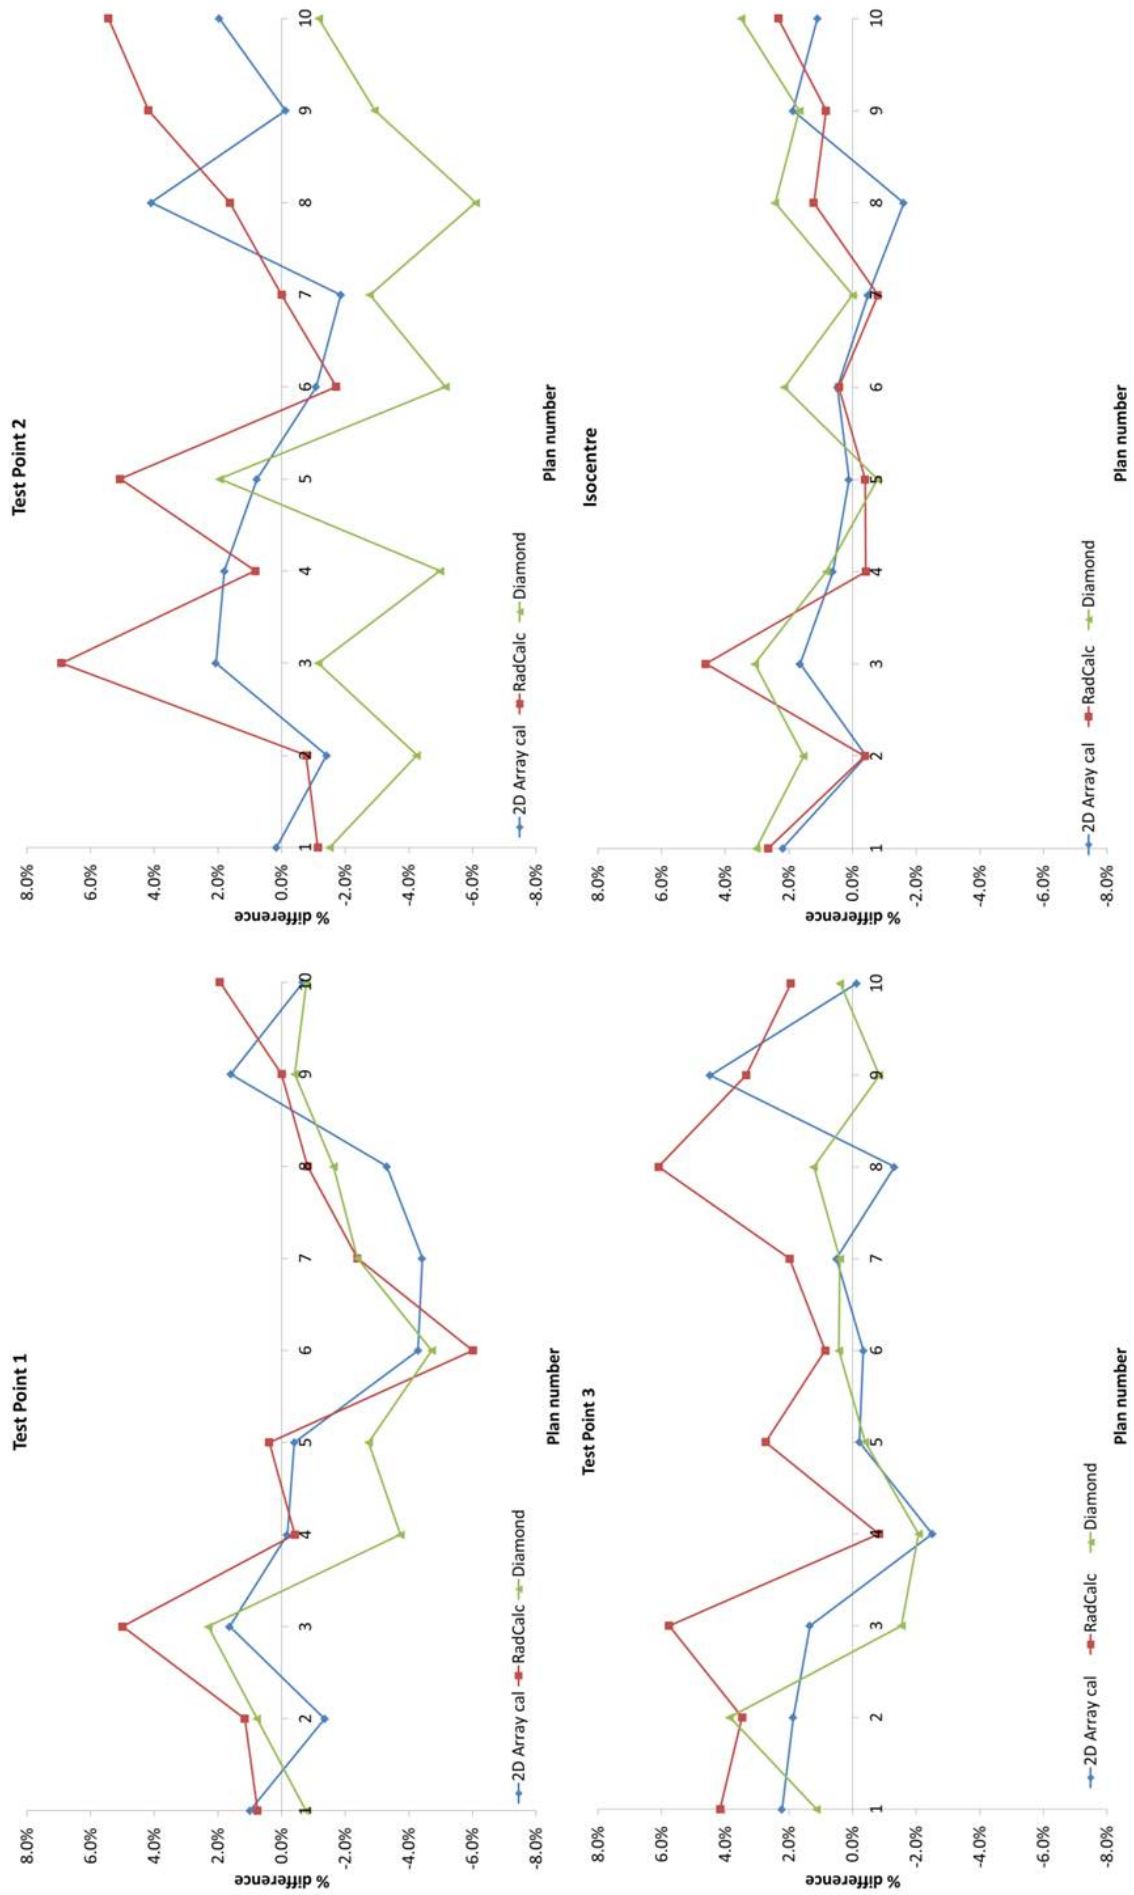
Figure B.13: Comparison of MONACO with RadCalc, DIAMOND and the measured data at the four measurement points (B.11). for all 10 plans.

B.3   Fluence from DIAMOND

Showing that the DIAMOND software calculated point doses with the same level of accuracy as RadCalc demonstrated that the software was working and that the data that had been entered was correct. To predict the DAVID signal using DIAMOND, its function to generate normalised fluence maps was utilised. Ten fluence maps of IMRT deliveries were generated and exported as matrices in .txt format. Software was written to import the fluence maps, extract predicted signals (Figure B.14) and compare them to the normalised DAVID responses measured on the linac (Figure B.15).

The results shown in Figure B.15 are typical of the results derived by the DIAMOND-generated fluence technique – notably poor. Using processes similar to those discussed in Section 4.3 efforts were made to enhance the predicted signal with iteratively derived factors, but little improvement was made and the derived factors were of the same order of magnitude as the factors themselves, so not deemed appropriate.
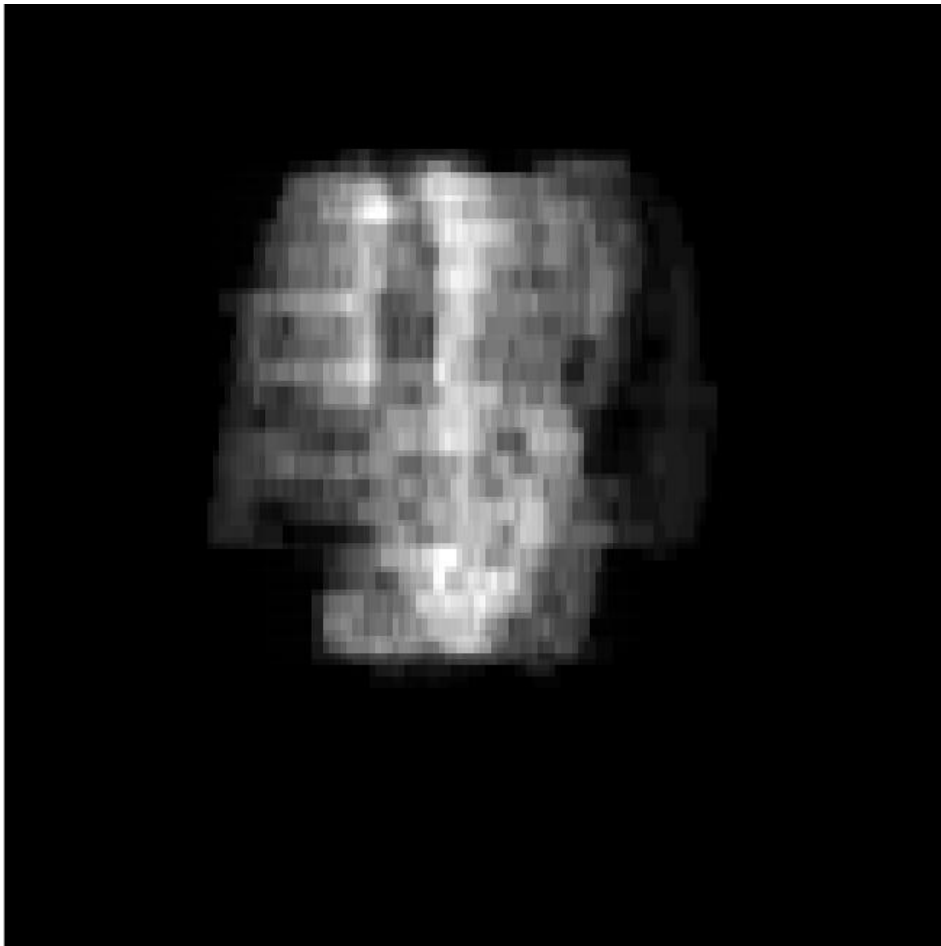
Figure B.14: Normalised fluence map of VMAT delivery.

Further investigation showed the fluence maps to be of quite a poor standard. Figure B.16 shows a $10 \times 10\text{cm}^2$ fluence map with the accompanying cross-plane profile. With the in-field normalised fluence equalling one and the penumbral region (in both in-plane and cross-plane directions) equalling 0.5, it is evident that specific and accurate penumbral modelling was not a feature of the algorithm.
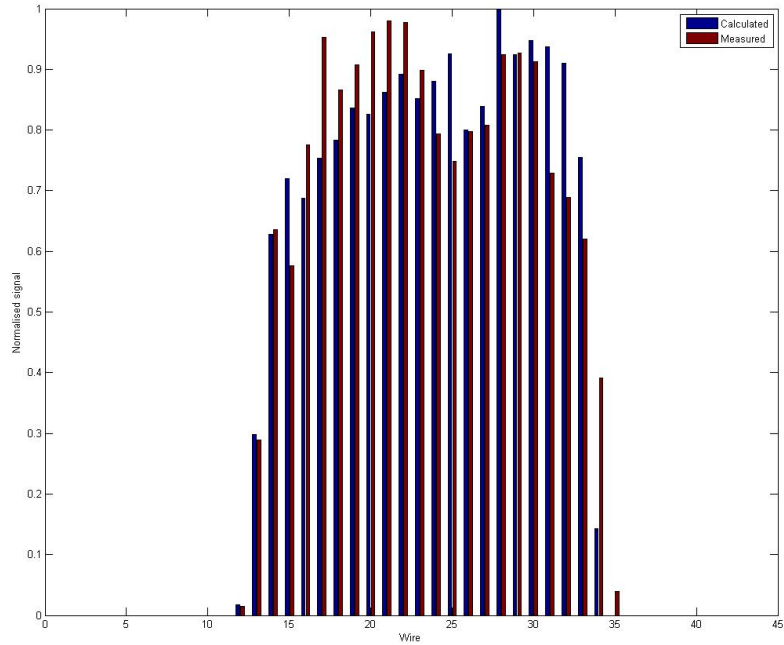
Figure B.15: Comparison of signal generated by DIAMOND's normalised fluence map and measured data.

Also with the out-of-field area equalling zero, it is clear that scatter, leaf leakage and inter-leaf leakage were not considered in the calculation. Attempts were made to modify the fluence maps and the interaction with the DAVID in an approach similar to that discussed in Greer et al. (2009), but these were unsuccessful.

With this in mind, it was decided to abandon DIAMOND as a potential source of usable fluence maps and move on to a different method – dose maps generated from Monte-Carlo modelling of the linear accelerator and DAVID. This is the subject of the next chapter.

Figure B.16: DIAMOND-generated square field fluence and profile

Appendix C

Efficiency enhancing methods in EGSnrc

C.1 Introduction

Monte Carlo programming is regarded as the most accurate way of calculating dose delivered by linear accelerators Fragoso et al. (2009), the cost of this accuracy is computing time. Even with multi-core computers the sheer number of calculations th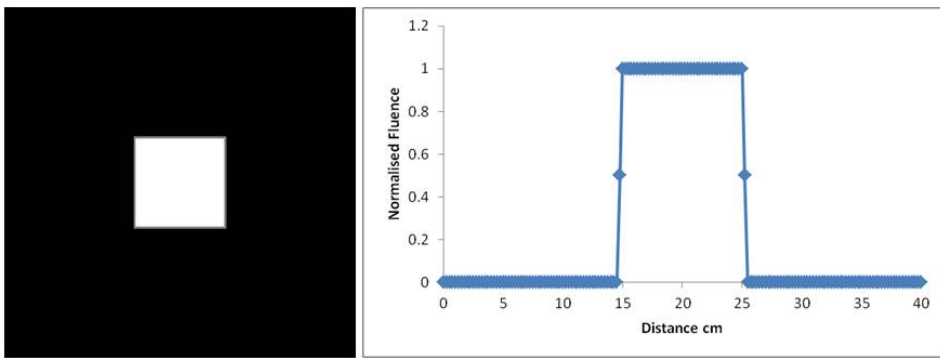at need to be made to individually simulate the histories of millions of particles means that computation time needed to provide accurate doses can often be hours. To try and reduce the computation time, a number of efficiency enhancing have been developed; a number of these are included in the EGSnrc distribution and are discussed in this section.

Broadly speaking, efficiency enhancing techniques can be split in to two types: Approximate Efficiency Improving Techniques (AEITs and Variance Reduction Techniques (VRTs). AEITS make approximations about particles and their history and ignore particles that are not going to make a meaningful contribution to the regions of interest. When using AEITs it is important to select ones with appropriate parameters that will not bias the results by ignoring important particles. AEITs included in the EGS distribution include charged particle Range Rejection (RR), photon and electron transport cut-off energies and the condensed history technique. VRTs do not change the the physics so should not bias the results. VRTs available with the EGS distribution are bremsstrahlung photon splitting, Russian roulette and photon interaction forcing.

C.2 Description of efficiency enhancing options

C.2.1 Range Rejection and energy cut off

This technique uses the MXRNGE subroutine compute the residual ranges in each medium as a function of electron energy up until the user-defined electron-energy cutoff before the simulation begins. If and electron enters a medium without sufficient energy to escape, its history is terminated Rogers et al. (1995). Range rejection means that time is not wasted on interactions that are not going to affect the total energy deposition and has been shown save considerable time in electron

calculations Rogers et al. (1990).

Range rejection was turned on for most of the calculations in the thesis as the highly accurate determination energy deposition by low energy electrons was not necessary and would be time consuming . However, in section 5.4.1 when the effect of the wire was being determined and the energy deposition in small volumes was being determined it was turned off so because the small volume meant that it was necessary to know the particles position with more precision.

ECUT is the electron transport cut off energy, a user-defined variable below which electron interactions are no longer calculated and the dose is energy is assumed deposited in the medium its in. For most detailed work ECUT = 0.7MeV is acceptable (Rogers et al. (2001)) however, in cases where the energy deposition in small volumes of interest is important, smaller values can be used. ECUT = 0.511 was used in Section 5.4.1.

PCUT is the photon transport cut off energy, in the event of a photon's energy dropping below this value, the particle tracking is terminated and it is assumed that the energy is deposited in the current medium. Having a low value of PCUT does not incur a big calculation-time penalty and it is recommend that a value of 0.01MeV is generally used (Rogers et al. (2001)).

C.2.2   Photon Forcing

For thin or low density materials, the likelihood of a photon interacting can be quite low, simulation statistics can be improved by forcing photons to interact. Photons that are forced to interact are split in to a scattered photon and a photon continuing as if the interaction had not occurred. The weight of the scattered photon is equal to the probability of the interaction taking place, the unscattered photon take the rest of the weight Rogers et al. (2001). The unscattered photon cannot be forced to interact again in the forcing CM, the scattered photon can, however, be forced to interact as many times as the user specifies. Photon forcing is useful for improving Bremsstrahlung statistics and can be particularly useful when used with Bremsstrahlung splitting Rogers et al. (2001), Mohammed et al. (2016)

C.2.3   Russian Roulette

In the case of Bremsstrahlung splitting, Russian Roulette (RR) can be turned on, this essentially eliminates some higher order charged particles and distributes their weight to the surviving charged particles. This removes charged particles, and their associated computational-time penalty, from the calculation and retains the bremsstrahlung photons. Rogers et al. (2001) Kawrakow et al. (2004).

C.2.4   Bremsstrahlung Splitting

Bremsstrahlung-generated photons are often often the most relevant consequence of an interaction, unfortunately the physical process of generating these photons is inherently inefficient. The variance associated with the resulting Bremsstrahlung spectrum can be reduced by artificially enhancing the number of photons produced at each Bremsstrahlung event. The electron path and energy is normally, for first

generation particles, unimportant, so splitting algorithms do not generally modify path of the photon-generating electron.In order to avoid biasing the results, the statical weight of each of these particles is reduced by the reciprocal of the Number of Bremsstrahlung SPLitting photons (NBSPL) produced Meurant (2012) Bielajew et al. (1989), this is a user-defined value normally between 20 and 100 Kawrakow et al. (2004). There are three types of Bremsstrahlung splitting associated with the EGSnrc code[1]:

- Uniform Bremsstrahlung Splitting (UBS) where the the NBSPL is applied to all Bremsstrahlung photons

- Selective Bremsstrahlung Splitting (SRS) where the NBRSPL is modified to enhance the Bremsstrahlung spectrum in the direction of the field, this saves time calculating histories for photons that will not reach the field of interest.

- Directional Bremsstrahlung Splitting (DBS) which is an improvement on the SBS approach. This was used for all BEAMnrc simulations and is discussed below.

Statistics in a field of interest generated with SBS turned on are compromised by the variable splitting number producing a large range of photon weights. Also, some photons aimed away from the field of view can, through "chance events" still create particles that will get to the field of interest, as the progenitor photons for these particles wont have been split, the particles and photons generated from them will have a high weight and adversely affect statistics (Figure C.2). The effect of these chance-event particles was ameliorated in the SBS code by a "background-splitting" subroutine, but this is at the expense of CPU time. DBS eliminates the need for the background-splitting and ensures that all the photons in the field of view have the same weight, the overall effect of this being a factor 6 and 20 increase in efficiency when DBS is used instead of SRS and UBS respectively Kawrakow et al. (2004).

In general, DBS handles Bremsstrahlung events initiated by fat[2] particles by splitting them NBRSPL times, the resultant photons will all have weight NBRSPL^{-1}. The resulting photons that are aimed at the field of view are kept, the ones aimed away from it are subject to RR, those surviving RR have their weight increased and making them "fat". Different interactions are handled slightly differently, but the result is that all photons inside the target field will be non-fat, each having a weight NBRSPL^{-1}, all those outside the target field will be fat, having a weight of 1 Kawrakow et al. (2004). DBS is largely concerned with photons, the electrons that are involved in the production of these photons are subject to RR meaning that the small amount (small compared to the DBS-amplified number of photons) reaching the target field will all be fat. So, although photon statistics will be good, containment electron statistics will be poor, this problem is overcome using the electron-splitting technique Kawrakow et al. (2004) Rogers et al. (2001).

Electron splitting, when used in conjunction with DBS requires a splitting plane, in addition to the RR plane needed for DBS. Both planes are typically

---

[1]UBS and DBS are available with the current distribution, SBS has been discontinued Rogers et al. (2001)

[2]fat particles are ones with high weight, all primary particles will be fat, fat particles can be left over from splitting algorithms and, in some cases generated by them

in the flattening filter, with the splitting plane being below the RR plane. The idea behind this being that electrons in the upper part of the head are unlikely to get to the target plane, whereas electrons in the lower part of the head are. Fat electrons crossing the splitting plane are split NBRSPL times each having their weight reduced by NBRSPL^{-1}. Non-fat charged particles can go on to interact and produce non-fat photons, of these photons, the ones not directed at the field of view will be subject to RR as described DBS process above Kawrakow et al. (2004)

There is a finite possibility that fat electrons will get to the target plan. In Figure C.1 the dose deposition should be symmetric about the centre, the increased dose to the right of centre in the top plot is likely to be from a fat photon as a result of DBS. While the fat photon increases the dose deposited in a voxel, it also increases the uncertainty associated with the dose - this is reflected in the error bar associated with the volume. The uncertainty used to determine the size of the error bars on the graph was extracted the BEAMdp output file - these values calculated from the uncertainties contained with in the phase-space file. For the phase-space uncertainties to be calculated correctly the parameters in the model need to be entered correctly (Rogers et al. 2001, Ma & Rogers 2018, Walters et al. 2002). The distribution in this figure should be symmetric and is if the uncertainties are taken in to account. This demonstrates that the parameters in the model are appropriate and the code written by the author to extract them worked as intended. The phase space was regenerated using different random number seeds. The result of the same analysis is seen in the bottom of the figure. Due to the random nature of MC calculations in general and fat photon contamination in particular the fat photon was not present in this run.

In the BEAMnrc simulations performed for this work, electron splitting plane was put just below the RR plane and both were in the flattening filter as recommend by Kawrakow et al. (2004) Rogers et al. (2001). the behaviour efficiency improvements relation to NBRSPL was has been shown to be independent of placement and the optimal number has been shown to be 1000 Kawrakow et al. (2004), so this was used throughout.

C.2.5 ISMOOTH

ISMOOTH is as DOSXYZnrc option to redistribute particles in a phase-space source file symmetrically about the central axes. This option is only used when the phase-space file is recycled multiple times. It has the effect of reducing statistical uncertainties (mainly in surface dose) associated with repeatedly using the same phase-space file during a single DOSXYZnrc simulation. This option should only be used for symmetrical fields, in the case of asymmetric fields, particle fluence could potentially be created where there is none and vice versa Figure C.3
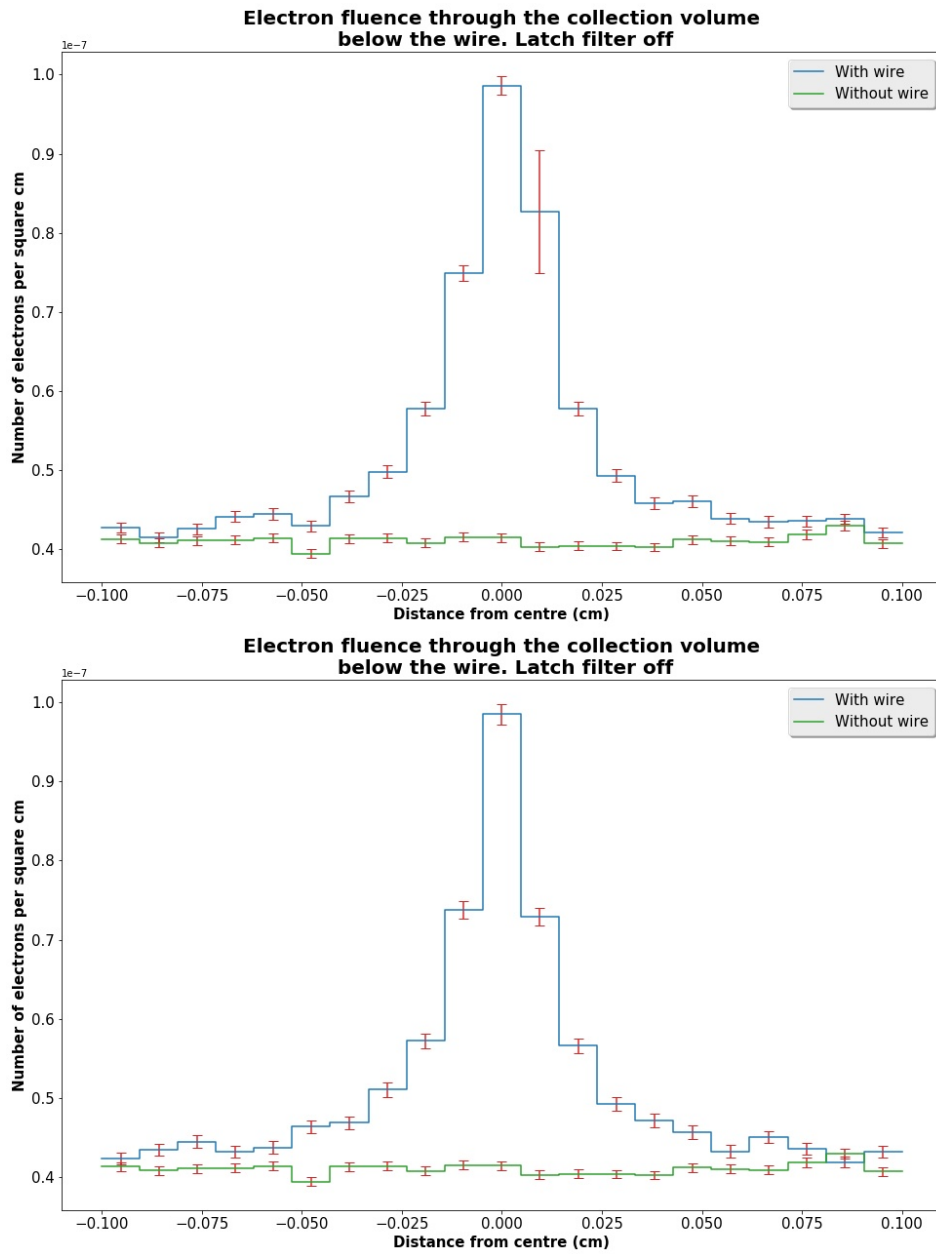
Figure C.1: Taken from Chapter 5, here Electron fluence passing through the bottom collection volume (7–9 on Figure 5.16) extracted using beamdp phase-space analysis tool (Ma & Rogers 2018).
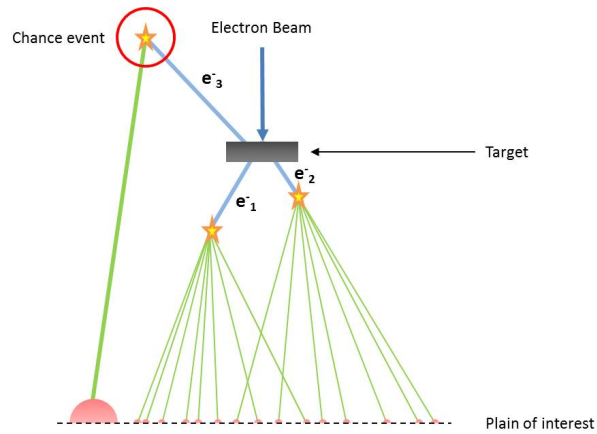
Figure C.2: The figure shows an electron beam striking a target, three electrons are produced, e^-_1 and e^-_2 produce Bremsstrahlung photons each with a weight NBRSPL^{-1}, e^-_3 is not thought to produce particles or photons that will hit the field of interest, however, an unlikely chance event causes a photon to be directed to the plain of interest, without the back-ground splitting subroutine, this photon, having not come from a splitting event, will have considerably higher statistic le weight than the other photons adversely affecting the photon statistics in the plain of interest
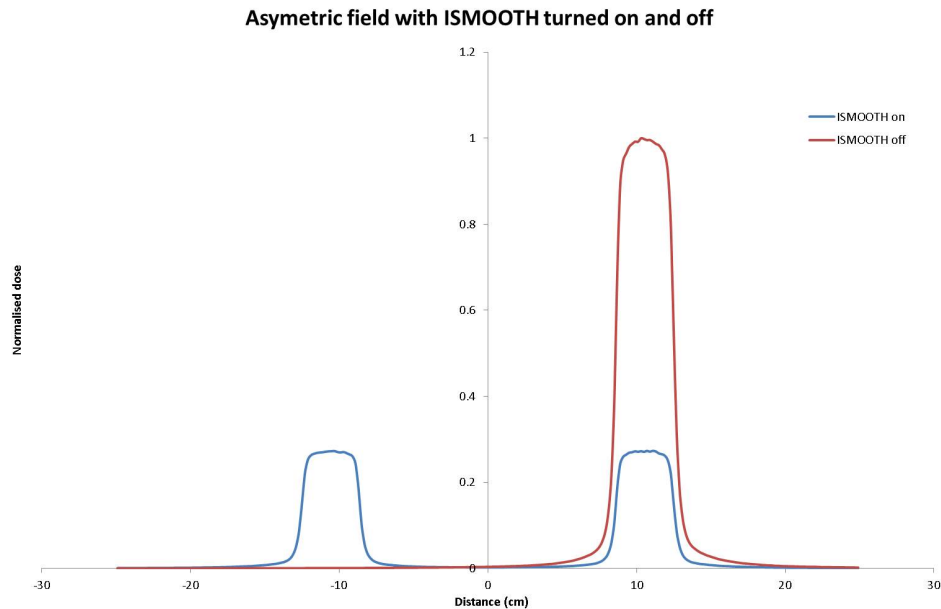

Figure C.3: 5×5 field offset in the xdirection. Only difference between the two plots is that one was generated with ISMOOTH on and the other with ISMOOTH off.

176

Appendix D

ESTRO 2011

Figure D.1: Poster accepted for ESTRO 2011

Appendix E

ESTRO 2012

Figure E.1: Poster accepted for ESTRO 2012

Appendix F

ESTRO 2012

Figure F.1: Poster accepted for ESTRO 2012

Appendix G

ESTRO 2013

Figure G.1: Poster accepted for ESTRO 2013

Appendix H

ESTRO 2016

Figure H.1: Poster accepted for ESTRO 2016

Appendix I

BEAMnrc input file for MLCi2 linac

```
#!GUI1.0
AIR700ICRU
0, 0, 0, 0, 0, 2, 0, IWATCH ETC.
1500000000.0, 33, 97, 2, 2, 1000, 0, 0, NCASE ETC.
10, 63.31, 3, 4, 1, 15.1, DIRECTIONAL BREM OPTIONS
-1, 19, -0.2, 0, 0, 1, 1.35, -0.24, 0.0, 0.0, IQIN, ISOURCE + OPTIONS
1, SPECTRUM
_nobackup_umdj_EGSnrc_HEN_HOUSE_spectra_Elekta_UFFF_5.spectrum
1
0, 0, 0.7, 0.1, 0, 1, 2, 0 , ECUT,PCUT,IREJCT,ESAVE
0, 0, 0, 0, 0, PHOTON FORCING
1, 10, SCORING INPUT
0,0
0, DOSE COMPONENTS
-0.35, Z TO FRONT FACE
*********** start of CM FLATFILT with identifier target ***********
3, RMAX
Target
-0.35, ZMIN
3, NUMBER OF LAYERS
1, 0.261, # CONES, ZTHICK OF LAYER 1
0.28,
0.28,
1, 0.089, # CONES, ZTHICK OF LAYER 2
0.28,
0.28,
1, 1, # CONES, ZTHICK OF LAYER 3
2.9,
2.9,
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
CU700ICRU
0, 0, 0, 0,
WRE700
0, 0, 0, 0,
CU700ICRU
0, 0, 0, 0,
CU700ICRU
0, 0, 0, 0,
CU700ICRU
*********** start of CM FLATFILT with identifier pricol ***********
4, RMAX
Primary collimator
1.5, ZMIN
1, NUMBER OF LAYERS
1, 10.1, # CONES, ZTHICK OF LAYER 1
0.645,
3.14,
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
W700ICRU
*********** start of CM FLATFILT with identifier Flatten ***********
4.65, RMAX
6MV Flattening Filter
13.3, ZMIN
7, NUMBER OF LAYERS
1, 0.1, # CONES, ZTHICK OF LAYER 1
0.0,
0.26,
1, 0.49, # CONES, ZTHICK OF LAYER 2
0.26,
0.9,
1, 0.49, # CONES, ZTHICK OF LAYER 3
0.9,
1.45,
1, 0.55, # CONES, ZTHICK OF LAYER 4
1.45,
2.25,
1, 0.56, # CONES, ZTHICK OF LAYER 5
2.25,
3.25,
1, 0.2, # CONES, ZTHICK OF LAYER 6
4.65,
4.65,
1, 0.2, # CONES, ZTHICK OF LAYER 7
4.65,
4.65,
0.7, 0.01, 0, 0,
```

```
STEEL700ICRU
0.7, 0.01, 0, 0,
AIR700ICRU
0.7, 0.01, 0, 0,
STEEL700ICRU
0.7, 0.01, 0, 0,
AIR700ICRU
0.7, 0.01, 0, 0,
STEEL700ICRU
0.7, 0.01, 0, 0,
AIR700ICRU
0.7, 0.01, 0, 0,
STEEL700ICRU
0.7, 0.01, 0, 0,
AIR700ICRU
0.7, 0.01, 0, 0,
STEEL700ICRU
0.7, 0.01, 0, 0,
AIR700ICRU
0.7, 0.01, 0, 0,
STEEL700ICRU
0.7, 0.01, 0, 0,
AIR700ICRU
0.7, 0.01, 0, 0,
AL700ICRU
0.7, 0.01, 0, 0,
AIR700ICRU
*********** start of CM FLATFILT with identifier chamber ***********
7.4, RMAX
Chamber
16.44, ZMIN
21, NUMBER OF LAYERS
1, 0.0012, # CONES, ZTHICK OF LAYER 1
7.3,
7.3,
1, 0.1, # CONES, ZTHICK OF LAYER 2
4.5,
4.5,
1, 0.0012, # CONES, ZTHICK OF LAYER 3
7.3,
7.3,
3, 0.1, # CONES, ZTHICK OF LAYER 4
4.5, 5.4, 6.625,
4.5, 5.4, 6.625,
2, 0.0012, # CONES, ZTHICK OF LAYER 5
5.4, 6.625,
5.4, 6.625,
3, 0.22, # CONES, ZTHICK OF LAYER 6
4.5, 5.4, 6.625,
4.5, 5.4, 6.625,
2, 0.0012, # CONES, ZTHICK OF LAYER 7
5.4, 6.625,
5.4, 6.625,
3, 0.1, # CONES, ZTHICK OF LAYER 8
4.5, 5.4, 6.625,
4.5, 5.4, 6.625,
2, 0.0012, # CONES, ZTHICK OF LAYER 9
5.4, 6.625,
5.4, 6.625,
3, 0.1, # CONES, ZTHICK OF LAYER 10
4.5, 5.4, 6.625,
4.5, 5.4, 6.625,
2, 0.0012, # CONES, ZTHICK OF LAYER 11
5.4, 6.625,
5.4, 6.625,
3, 0.1, # CONES, ZTHICK OF LAYER 12
4.5, 5.4, 6.625,
4.5, 5.4, 6.625,
2, 0.0012, # CONES, ZTHICK OF LAYER 13
5.4, 6.625,
5.4, 6.625,
3, 0.1, # CONES, ZTHICK OF LAYER 14
4.5, 5.4, 6.625,
4.5, 5.4, 6.625,
2, 0.0012, # CONES, ZTHICK OF LAYER 15
5.4, 6.625,
5.4, 6.625,
3, 0.1, # CONES, ZTHICK OF LAYER 16
4.5, 5.4, 6.625,
4.5, 5.4, 6.625,
2, 0.0012, # CONES, ZTHICK OF LAYER 17
5.4, 6.625,
5.4, 6.625,
1, 0.31, # CONES, ZTHICK OF LAYER 18
6.625,
6.625,
1, 0.0012, # CONES, ZTHICK OF LAYER 19
6.625,
6.625,
1, 0.0676, # CONES, ZTHICK OF LAYER 20
4.5,
4.5,
1, 0.0012, # CONES, ZTHICK OF LAYER 21
7.3,
7.3,
0, 0, 0, 0,
MYLAR700ICRU
0, 0, 0, 0,
MYLAR700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
ALUMINA700ICRU
0, 0, 0, 0,
MYLAR700ICRU
0, 0, 0, 0,
MYLAR700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
ALUMINA700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
AL700ICRU
0, 0, 0, 0,
MYLAR700ICRU
```

0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
AL700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
AL700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
AL700ICRU
0, 0, 0, 0,
MYLAR700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
AL700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
ALUMINA700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
AL700ICRU
0, 0, 0, 0,
MYLAR700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
AL700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
AL700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
AL700ICRU
0, 0, 0, 0,
MYLAR700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
AL700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
ALUMINA700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
AL700ICRU
0, 0, 0, 0,
MYLAR700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
AL700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
AL700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
AL700ICRU
0, 0, 0, 0,
MYLAR700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
AL700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
ALUMINA700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
AL700ICRU
0, 0, 0, 0,
MYLAR700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
AL700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
AL700ICRU
0, 0, 0, 0,
MYLAR700ICRU
0, 0, 0, 0,
AL700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
AL700ICRU
0, 0, 0, 0,
MYLAR700ICRU
0, 0, 0, 0,
MYLAR700ICRU
*********** start of CM SLABS with identifier bkplt ***********
6, RMAX
Back scatter plate
1, NSLABS
18.25, ZMIN
0.29, 0, 0, 0, 0, 0
AL700ICRU
*********** start of CM MIRROR with identifier mirror ***********
5.18, RMAX
mirror
18.54, 7.94, ZMIN, ZTHICK
5.18, -5.19, XFMIN, XBMIN
2, # LAYERS

0.0012, thickness of layer 1
3e-5, thickness of layer 2
0, 0, 0, 0,
POLYETH700ICRU
0, 0, 0, 0,
AL700ICRU
0, 0, 0, 0,
AIR700ICRU
0, 0, 0, 0,
AIR700ICRU
********** start of CM SYNCMLCE with identifier MLC **********
26, RMAX
MLCi2 leaf bank
0, 2, ORIENT, MODE
40, NUM_LEAF
29.3, 37.5, ZMIN, ZMAX
35, 34, ZSTEPL, ZSTEPR
0.0, TGW
0.14315, 0.18395, X3, X4
1.0, 100, SPACE, SSD
0, LBROT
0, ENDTYPE
15.02, 32.836, LEAFRADIUS, CIL
_nobackup_umdj_EGSnrc_egs_home_BEAM_MLCi2_Match_1010MLC
0.7, 0.01, 0, 0,
AIR700ICRU
0.7, 0.01, 0, 0,
WRE700
********** start of CM SYNCJAWS with identifier Jaws **********
30, RMAX
Y backup and X jaws
2, 2, # PAIRED BARS OR JAWS, field type
Y
X
_nobackup_umdj_EGSnrc_egs_home_BEAM_MLCi2_Match_1010JAW
0.7, 0.01, 0, 0,
0.7, 0.1, 0, 0,
WRE700
0.7, 0.1, 0, 0,
WRE700
********** start of CM SLABS with identifier MylWin **********
30, RMAX
Mylar window
1, NSLABS
54.8, ZMIN
0.01, 0.7, 0.1, 0, 0, 0
MYLAR700ICRU
********** start of CM SLABS with identifier Airslab **********
50, RMAX
8.5cm to DAVID
1, NSLABS
54.81, ZMIN
8.5, 0.7, 0.1, 1, 0, 0
AIR700ICRU
*******************end of all CMs****************************
#######################
:Start MC Transport Parameter:

Global ECUT= 0.7
Global PCUT= 0.1
Global SMAX= 5
ESTEPE= 0.25
XIMAX= 0.5
Boundary crossing algorithm= EXACT
Skin depth for BCA= 0
Electron-step algorithm= PRESTA-II
Spin effects= On
Brems angular sampling= Simple
Brems cross sections= BH
Bound Compton scattering= Off
Compton cross sections= default
Pair angular sampling= Simple
Pair cross sections= BH
Photoelectron angular sampling= Off
Rayleigh scattering= Off
Atomic relaxations= Off
Electron impact ionization= Off
Photon cross sections= si
Photon cross-sections output= Off

:Stop MC Transport Parameter:
#######################

Bibliography

AAMP (2008), 'Task group 106 report: Accelerator beam data commissioning equipment and procedures', *Medical Physics* **35**(9), 4186–4215. Das, Indra J. and Cheng, Chee-Wai and Watts, Ronald J. and Ahnesjö, Anders and Gibbons, John and Li, X. Allen and Lowenstein, Jessica and Mitra, Raj K. and Simon, William E. and Zhu, Timothy C.

AAPM (2001), 'Task group 50 report: Basic applications of multileaf collimators'. Boyer, Arthur et al.

AAPM (2005), 'Task group 87 report: Diode in vivo dosimetry for patients receiving external beam radiation therapy'. Yorke, Ellen and Alecu, R and Ding, L and Fontenla, D and Kalend, A and Kaurin, D and Masterson-McGary, ME and Marinello, G and Matzen, T and Saini, A and others.

AAPM (2007), 'Task group 105 report: Issues associated with clinical implementation of monte carlo-based photon and electron external beam treatment planning', *Medical Physics* **34**(12). Chetty, Indrin J. and Curran, Bruce and Cygler, Joanna E. and DeMarco, John J. and Ezzell, Gary and Faddegon, Bruce A. and Kawrakow, Iwan and Keall, Paul J. and Liu, Helen and Ma, C.-M. Charlie and Rogers, D. W. O. and Seuntjens, Jan and Sheikh-Bagheri, Daryoush and Siebers, Jeffrey V.

AAPM (2009), 'Task group 142 report: Quality assurance of medical accelerators a', *Medical physics* **36**, 4197–4212. Klein, Eric E and Hanley, Joseph and Bayouth, John and Yin, Fang-Fang and Simon, William and Dresser, Sean and Serago, Christopher and Aguirre, Francisco and Ma, Lijun and Arjomandy, Bijan.

AAPM (2011), 'Task group 114 report: Verification of monitor unit calculations for non-imrt clinical radiotherapy', *Medical Physics* **38**(1), 504–530. Stern, Robin L. and Heaton, Robert and Fraser, Martin W. and Murty Goddu, S. and Kirby, Thomas H. and Lam, Kwok Leung and Molineu, Andrea and Zhu, Timothy C.

AAPM (2014), 'Task group 71 report: Monitor unit calculations for external photon and electron beams', *Medical physics* **41**(3). Gibbons, John P and Antolak, John A and Followill, David S and Huq, M Saiful and Klein, Eric E and Lam, Kwok L and Palta, Jatinder R and Roback, Donald M and Reid, Mark and Khan, Faiz M.

AAPM (2018), 'Task group 218 report: Tolerance limits and methodologies for imrt measurement-based verification qa', *Medical physics* . Miften, Moyed and

Olch, Arthur and Mihailidis, Dimitris and Moran, Jean and Pawlicki, Todd and Molineu, Andrea and Li, Harold and Wijesooriya, Krishni and Shi, Jie and Xia, Ping and others.

Adamus-Górka, M., Mavroidis, P., Lind, B. K. & Brahme, A. (2011), 'Comparison of dose response models for predicting normal tissue complications from cancer radiotherapy: application in rat spinal cord', *Cancers* **3**(2), 2421–2443.

Agnew, A., Agnew, C., Grattan, M., Hounsell, A. & McGarry, C. (2014), 'Monitoring daily mlc positional errors using trajectory log files and epid measurements for imrt and vmat deliveries', *Physics in Medicine & Biology* **59**(9), N49.

Agostinelli, S., Allison, J., Amako, K. a., Apostolakis, J., Araujo, H., Arce, P., Asai, M., Axen, D., Banerjee, S., Barrand, G. . et al. (2003), 'Geant4—a simulation toolkit', *Nuclear instruments and methods in physics research section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **506**, 250–303.

Ahnesjö, A. (1989), 'Collapsed cone convolution of radiant energy for photon dose calculation in heterogeneous media', *Medical physics* **16**(4), 577–592.

Alber, M., Broggi, S., De Wagter, C., Eichwurzel, I., Engström, P., Fiorino, C., Georg, D., Hartmann, G., Knöös, T., Leal, A. et al. (2008), 'Guidelines for the verification of imrt', *ESTRO booklet* **7**.

Alber, M. & Nüsslin, F. (1999), 'An objective function for radiation treatment optimization based on local biological measures', *Physics in Medicine & Biology* **44**(2), 479.

Ash, D. (2007), 'Lessons from epinal'.

Ash, D. & Bates, T. (1994), 'Report on the clinical effects of inadvertent radiation underdosage in 1045 patients', *Clinical Oncology* **6**(4), 214–226.

Asnaashari, K., Chow, J. C. & Heydarian, M. (2013), 'Dosimetric comparison between two mlc systems commonly used for stereotactic radiosurgery and radiotherapy: a monte carlo and experimental study', *Physica medica* **29**(4), 350–356.

Asuni, G., Jensen, J. M. & McCurdy, B. M. C. (2011), 'A monte carlo investigation of contaminant electrons due to a novel in vivo transmission detector', *Physics in Medicine and Biology* **56**(4), 1207.

Azorín, Nieto, J. (2004), Thermoluminescence dosimetry (tld) and its application in medical physics, *in* 'AIP Conference Proceedings', Vol. 724, AIP, pp. 20–27.

Bai, S., Li, G., Wang, M., Jiang, Q., Zhang, Y. & Wei, Y. (2013), 'Effect of mlc leaf position, collimator rotation angle, and gantry rotation angle errors on intensity-modulated radiotherapy plans for nasopharyngeal carcinoma', *Medical Dosimetry* **38**(2), 143–147.

Baker, S., Dahele, M., Lagerwaard, F. J. & Senan, S. (2016), 'A critical review of recent developments in radiotherapy for non-small cell lung cancer', *Radiation oncology* **11**(1), 115.

Bielajew, A. F., Mohan, R. & Chui, C.-S. (1989), 'Improved bremsstrahlung photon angular sampling in the egs4 code system', *National Research Council of Canada Report PIRS-0203* .

Bielajew, A. F. & Rogers, D. (1986), 'Presta: the parameter reduced electron-step transport algorithm for electron monte carlo transport', *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* **18**(1-6), 165–181.

BIPM (2006), 'The international system of units (si)', *Bureau International des Poids et Mesures,* .

BIR (2007), 'Towards safer radiotherapy'. British Institute of Radiology, Institute of Physics and Engineering in Medicine, National Patient Safety Agency, Society and College of Radiographers, The Royal College of Radiologists.

Boesecke, R., Doll, J., Bauer, B., Schlegel, W., Pastyr, O. & Lorenz, W. (1988), 'Treatment planning for conformation therapy using a multi-leaf collimator.', *Strahlentherapie und Onkologie: Organ der Deutschen Rontgengesellschaft...[et al]* **164**(3), 151–154.

Bogdanich, W. (2010), 'Radiation offers new cures, and ways to do harm', *The New York Times* **23**.

Bohsung, J., Strahlentherapie, K. f., Berlin, C. & Mitte, S. (2004), 'The 2d-array seven29 a new way of dosimetric verification of imrt beams', *Application note* .

Bortfeld, T. (2006), 'Imrt: a review and preview', *Physics in Medicine & Biology* **51**(13), R363.

Brahme, A. (1984), 'Dosimetric precision requirements in radiation therapy', *Acta Radiologica: Oncology* **23**(5), 379–391.

Brahme, A. (1987), 'Design principles and clinical possibilities with a new generation of radiation therapy equipment: A review', *Acta Oncologica* **26**(6), 403–412.

Brahme, A. (1988*a*), 'Optimal setting of multileaf collimators in stationary beam radiation therapy.', *Strahlentherapie und Onkologie* **164**(6), 343–350.

Brahme, A. (1988*b*), 'Optimization of stationary and moving beam radiation therapy techniques', *Radiotherapy and Oncology* **12**(2), 129–140.

Brahme, A., Roos, J.-E. & Lax, I. (1982), 'Solution of an integral equation encountered in rotation therapy', *Physics in Medicine & Biology* **27**(10), 1221.

Briesmeister, J. F. et al. (2000), 'Mcnptm-a general monte carlo n-particle transport code', *Version 4C, LA-13709-M, Los Alamos National Laboratory* .

British Standards Institution (2000), *Transition to ISO 9001: 2000: Guidance on ISO 9001: 2000, Quality Management Systems, Requirements*, British Standards Institution. MacNee, Colin and Pyle, Jim and Sutherland, Sandy.

Buckley, L. A. & Rogers, D. (2006), 'Wall correction factors, pwall, for thimble ionization chambers', *Medical physics* **33**(2), 455–464.

Budgell, G., Brown, K., Cashmore, J., Duane, S., Frame, J., Hardy, M., Paynter, D. & Thomas, R. (2016), 'IPEM topical report 1: guidance on implementing flattening filter free (FFF) radiotherapy', *Physics in Medicine & Biology* **61**(23), 8360.

Budgell, G. J., Mott, J., Williams, P. C. & Brown, K. (2000), 'Requirements for leaf position accuracy for dynamic multileaf collimation', *Physics in Medicine & Biology* **45**(5), 1211.

Caivano, R., Califano, G., Fiorentino, A., Cozzolino, M., Oliviero, C., Pedicini, P., Clemente, S., Chiumento, C. & Fusco, V. (2014), 'Clinically relevant quality assurance for intensity modulated radiotherapy plans: gamma maps and dvh-based evaluation', *Cancer investigation* **32**(3), 85–91.

Cancer Research (2019), 'Fund gives around 5,800 more cancer patients access to advanced radiotherapy treatment every year — cancer research uk', `https://www.cancerresearchuk.org/about-us/cancer-news/press-release/2013-07-11-fund-gives-around-5800-more-cancer-patients-access/-to-advanced-radiotherapy-treatment-every-year#2`. (Accessed on 02/08/2019).

Celi, S., Costa, E., Wessels, C., Mazal, A., Fourquet, A. & Francois, P. (2016), 'Epid based in vivo dosimetry system: clinical experience and results', *Journal of applied clinical medical physics* **17**(3), 262–276.

Chalmers, A. J. (2016), 'Science in focus: combining radiotherapy with inhibitors of the dna damage response', *Clinical Oncology* **28**(5), 279–282.

Chandraraj, V., Stathakis, S., Manickam, R., Esquivel, C., Supe, S. S. & Papanikolaou, N. (2010), 'Consistency and reproducibility of the vmat plan delivery using three independent validation methods', *Journal of Applied Clinical Medical Physics* **12**(1), 129–140.

Chang, J. Y. (2015), Intensity-modulated radiotherapy, not 3 dimensional conformal, is the preferred technique for treating locally advanced lung cancer, *in* 'Seminars in radiation oncology', Vol. 25, Elsevier, pp. 110–116.

Cheng, J. Y., Ning, H., Arora, B. C., Zhuge, Y. & Miller, R. W. (2016), 'Output factor comparison of monte carlo and measurement for varian truebeam 6 mv and 10 mv flattening filter-free stereotactic radiosurgery system', *Journal of applied clinical medical physics* **17**(3), 100–110.

Chibani, O., Moftah, B. & Ma, C.-M. C. (2011), 'On monte carlo modeling of megavoltage photon beams: A revisited study on the sensitivity of beam parameters', *Medical physics* **38**(1), 188–201.

Childress, N., Chen, Q. & Rong, Y. (2015), 'Parallel/opposed: Imrt qa using treatment log files is superior to conventional measurement-based method', *Journal of applied clinical medical physics* **16**(1), 4–7.

Childress, N., Stephens, E., Eklund, D. & Zhang, M. (2012), 'Mobius3d white paper: dose calculation algorithm', *Houston, TX: Mobius Medical Systems* .

Cho, B. (2018), 'Intensity-modulated radiation therapy: a review with a physics perspective', *Radiation oncology journal* **36**(1), 1.

Convery, D. & Rosenbloom, M. (1992), 'The generation of intensity-modulated fields for conformal radiotherapy by dynamic collimation', *Physics in medicine & biology* **37**(6), 1359.

CRUK (2015), 'Cancer statistics for the uk'. (Accessed on 07/03/2018).
URL: *https://www.cancerresearchuk.org/health-professional/cancer-statistics-for-the-uk*

CRUK, N. (2014), 'Vision for radiotherapy 2014-24.'. (Accessed on 01/08/2018).
URL: *https://www.cancerresearchuk.org*

Cunningham, J., Shrivastava, P. & Wilkinson, J. (1972), 'Program irreg—calculation of dose from irregularly shaped radiation beams', *Computer Programs in Biomedicine* **2**(3), 192–199.

Dale, R. G., Jones, B. et al. (2007), *Radiobiological modelling in radiation oncology*, British Institute of Radiology.

Das, I. J., Ding, G. X. & Ahnesjö, A. (2008), 'Small fields: Nonequilibrium radiation dosimetry', *Medical Physics* **35**(1), 206–215.

De Vlamynck, K., Palmans, H., Verhaegen, F., De Wagter, C., De Neve, W. & Thierens, H. (1999), 'Dose measurements compared with monte carlo simulations of narrow 6 mv multileaf collimator shaped photon beams', *Medical physics* **26**(9), 1874–1882.

Delaby, N., Bouvier, J., Jouyaux, F., Barateau, A. & Lafond, C. (2017), '40. validation of a transit epid device for a clinical use: application to iviewdose (elekta)', *Physica Medica: European Journal of Medical Physics* **44**, 19–20.

Department of Health, Donaldson, L. (2006), 'On the state of public health', *2006 Annual Report of The Chief Medical Officer* .

Donaldson, L. (2007), 'Reducing harm from radiotherapy', *BMJ: British Medical Journal* **334**(7588), 272.

Edwards, C., Hamer, E., Mountford, P. & Moloney, A. (2007), 'An update survey of uk in vivo radiotherapy dosimetry practice', *The British journal of radiology* **80**(960), 1011–1014.

Edwards, C. & Mountford, P. (2009), 'Characteristics of in vivo radiotherapy dosimetry', *The British journal of radiology* **82**(983), 881–883.

Ezzell, G. A., Galvin, J. M., Low, D., Palta, J. R., Rosen, I., Sharpe, M. B., Xia, P., Xiao, Y., Xing, L. & Yu, C. X. (2003), 'Guidance document on delivery, treatment planning, and clinical implementation of imrt: report of the imrt subcommittee of the aapm radiation therapy committee', *Medical physics* **30**(8), 2089–2115.

Farkona, S., Diamandis, E. P. & Blasutig, I. M. (2016), 'Cancer immunotherapy: the beginning of the end of cancer?', *BMC medicine* **14**(1), 73.

Feldman, A., Edwards, F. M. & Hendee, W. R. (2001), 'The routine use of personal patient dosimeters is of little value in detecting therapeutic misadministrations', *Medical physics* **28**(3), 295–297.

Feller, W. (2008), *An introduction to probability theory and its applications*, Vol. 2, John Wiley & Sons.

Feng, Z., Wu, H., Zhang, Y., Zhang, Y., Cheng, J. & Su, X. (2015), 'Dosimetric comparison between jaw tracking and static jaw techniques in intensity-modulated radiotherapy', *Radiation Oncology* **10**(1), 28.

Feygelman, V., Opp, D., Javedan, K., Saini, A. & Zhang, G. (2010), 'Evaluation of a 3d diode array dosimeter for helical tomotherapy delivery qa', *Medical Dosimetry* **35**(4), 324–329.

Fippel, M. (1999), 'Fast monte carlo dose calculation for photon beams based on the vmc electron algorithm', *Medical Physics* **26**(8), 1466–1475.

Fippel, M., Haryanto, F., Dohm, O., Nüsslin, F. & Kriesen, S. (2003), 'A virtual photon energy fluence model for monte carlo dose calculation', *Medical Physics* **30**(3), 301–311.

Fippel, M., Laub, W., Huber, B. & Nüsslin, F. (1999), 'Experimental investigation of a fast monte carlo photon beam dose calculation algorithm', *Physics in Medicine & Biology* **44**(12), 3039.

Fisher, C. M. & Rabinovitch, R. (2014), 'Frontiers in radiotherapy for early-stage invasive breast cancer', *Journal of Clinical Oncology* **32**(26), 2894.

Fix, M. K., Keall, P. J., Dawson, K. & Siebers, J. V. (2004), 'Monte carlo source model for photon beam radiotherapy: photon source characteristics', *Medical physics* **31**(11), 3106–3121.

Fleckenstein, J., Jahnke, L., Lohr, F., Wenz, F. & Hesser, J. (2013), 'Development of a geant4 based monte carlo algorithm to evaluate the monaco vmat treatment accuracy', *Zeitschrift für Medizinische Physik* **23**(1), 33–45.

Fontenot, J. D. (2014), 'Evaluation of a novel secondary check tool for intensity-modulated radiotherapy treatment planning', *Journal of applied clinical medical physics* **15**(5), 207–215.

Fraass, B., Doppke, K., Hunt, M., Kutcher, G., Starkschall, G., Stern, R. & Van Dyke, J. (1998), 'Task group 53 report: Quality assurance for clinical radiotherapy treatment planning', *Medical Physics* **25**(10), 1773–1829.

Fragoso, M., Kawrakow, I., Faddegon, B. A., Solberg, T. D. & Chetty, I. J. (2009), 'Fast, accurate photon beam accelerator modeling using beamnrc: A systematic investigation of efficiency enhancing methods and cross-section data', *Medical physics* **36**(12), 5451–5466.

Franks, K. & McNair, H. (2012), 'Implementation of image-guided radiotherapy', *Clinical Oncology* **24**(9), 625–626.

Fredh, A., Scherman, J. B., Fog, L. S. & Munck af Rosenschöld, P. (2013), 'Patient qa systems for rotational radiation therapy: a comparative experimental study with intentional errors', *Medical physics* **40**(3), 031716.

Galvin, J. M., Smith, A. R. & Lally, B. (1993), 'Characterization of a multileaf collimator system', *International Journal of Radiation Oncology Biology Physics* **25**(2), 181–192.

Georg, D., Knöös, T. & McClean, B. (2011), 'Current status and future perspective of flattening filter free photon beams', *Medical physics* **38**(3), 1280–1293.

Gholampourkashi, S., Cygler, J. E., Belec, J., Vujicic, M. & Heath, E. (2019), 'Monte carlo and analytic modeling of an elekta infinity linac with agility mlc: Investigating the significance of accurate model parameters for small radiation fields', *Journal of Applied Clinical Medical Physics* **20**(1), 55–67.

Godart, J., Korevaar, E., Visser, R., Wauben, D. & van't Veld, A. A. (2011), 'Reconstruction of high-resolution 3d dose from matrix measurements: error detection capability of the compass correction kernel method', *Physics in Medicine & Biology* **56**(15), 5029.

Greene, D. D. & Williams, P. C. (1997), *Linear accelerators for radiation therapy*, 2nd ed edn, Bristol : Institute of Physics Pub. Previous ed.: 1985.

Greer, P. B., Cadman, P., Lee, C. & Bzdusek, K. (2009), 'An energy fluence-convolution model for amorphous silicon epid dose prediction', *Medical Physics* **36**(2), 547–555.

Griessbach, I., Lapp, M., Bohsung, J., Gademann, G. & Harder, D. (2005), 'Dosimetric characteristics of a new unshielded silicon diode and its application in clinical photon and electron beams', *Medical Physics* **32**(12), 3750–3754.

Grofsmid, D., Dirkx, M., Marijnissen, H., Woudstra, E. & Heijmen, B. (2010), 'Dosimetric validation of a commercial monte carlo based imrt planning system', *Medical physics* **37**(2), 540–549.

Harrison, R. & Morgan, A. (2007), 'In vivo dosimetry: hidden dangers?', *The British journal of radiology* **80**(957), 691–692.

Heath, E. & Seuntjens, J. (2003), 'Development and validation of a beamnrc component module for accurate monte carlo modelling of the varian dynamic millennium multileaf collimator', *Physics in Medicine & Biology* **48**(24), 4045.

Heilemann, G., Poppe, B. & Laub, W. (2013), 'On the sensitivity of common gamma-index evaluation methods to mlc misalignments in rapidarc quality assurance', *Medical physics* **40**(3).

Herwiningsih, S. & Fielding, A. (2016), Focal spot estimation of an elekta dedicated stereotactic linear accelerator monte carlo model, *in* 'Journal of Physics: Conference Series', Vol. 694, IOP Publishing, p. 012013.

Higgins, P., Alaei, P., Gerbi, B. & Dusenbery, K. (2003), 'In vivo diode dosimetry for routine quality assurance in imrt', *Medical physics* **30**(12), 3118–3123.

Hounsell, A. R. (1998), 'Monitor chamber backscatter for intensity modulated radiation therapy using multileaf collimators', *Physics in Medicine & Biology* **43**(2), 445.

Huq, M. S., Das, I. J., Steinberg, T. & Galvin, J. M. (2002), 'A dosimetric comparison of various multileaf collimators', *Physics in Medicine & Biology* **47**(12), N159.

Hussein, M., Adams, E. J., Jordan, T. J., Clark, C. H. & Nisbet, A. (2013a), 'A critical evaluation of the ptw 2d-array seven29 and octavius ii phantom for imrt and vmat verification', *Journal of applied clinical medical physics* **14**(6), 274–292.

Hussein, M., Clark, C. & Nisbet, A. (2017), 'Challenges in calculation of the gamma index in radiotherapy–towards good practice', *Physica Medica* **36**, 1–11.

Hussein, M., Rowshanfarzad, P., Ebert, M. A., Nisbet, A. & Clark, C. H. (2013b), 'A comparison of the gamma index analysis in various commercial imrt/vmat qa systems', *Radiotherapy and Oncology* **109**(3), 370–376.

Huyskens, D., Bogaerts, R., Verstraete, J., Lööf, M., Nyström, H., Fiorino, C., Broggi, S., Jornet, N., Ribas, M. & Thwaites, D. (2001), 'Practical guidelines for the implementation of in vivo dosimetry with diodes in external radiotherapy with photon beams (entrance dose)'.

IAEA (2013), *Development of procedures for in vivo dosimetry in radiotherapy.*, Series IA, International Atomic Energy Agency, Vienna.

ICRP (2001), 'Prevention of accidents to patients undergoing radiation therapy', *International Commission on Radiological Protection 86* **30**(3).

ICRP (2009), 'Preventing accidental exposures from new external beam radiation therapy technologies', *International Commission on Radiological Protection Publication 112* **39(4)**. P. Ortiz Lopez, J.-M. Cosset, P. Dunscombe, O. Holmberg, J.-C. Rosenwald, L. Pinillos Ashton, J.J. Vilaragut Llanes, S. Vatnitsky.

ICRU (1985), 'Stopping powers for electrons and positrons', *ICRU report* **37**. Brice, David K.

ICRU (1992), 'Photon, electron, proton and neutron interaction data for body tissues', *ICRU report* **46**. Bethesda, MD.

ICRU (1994), 'Prescribing, recording and reporting photon beam therapy', *International Commission on Radiation Units and Measurements report* **50**. Jones, Douglas.

ICRU (1999), 'Prescribing, recording and reporting photon beam therapy (supplement to icru report 50)', *International Commission on Radiation Units and Measurements report* **62**. T. Landberg J. Chavaudra J. Dobbs J. -P. Gerard G. Hanks J. -C. Horiot K. -A. Johansson T. Möller J. Purdy N. Suntharalingam H. Svensson.

IPEM (2006), 'Balancing costs and benefits of checking in radiotherapy', *Institute of Physics and Engineering in Medicine* . McKenzie, Alan.

IPEM (2008), 'Guidance for the clinical implementation of intensity modulated radiation therapy', *Institute of Physics and Engineering in Medicine (IPEM) Report* **96**. James, H and Beavis, A and Budgell, G and Clark, C and Convery, D and Mott, J and others.

IPEM (2018), 'Physics aspects of quality control in radiotherapy second edition', *Institute of Physics and Engineering in Medicine report* **81**. Patel, Weston S, Palmer A and others.

Islam, M. K., Norrlinger, B. D., Smale, J. R., Heaton, R. K., Galbraith, D., Fan, C. & Jaffray, D. A. (2009), 'An integral quality monitoring system for real-time verification of intensity modulated radiation therapy', *Medical physics* **36**(12), 5420–5428.

Izewska, J. & Andreo, P. (2000), 'The iaea/who tld postal programme for radiotherapy hospitals', *Radiotherapy and oncology* **54**(1), 65–72.

Johnson, D., Weston, S. J., Cosgrove, V. P. & Thwaites, D. I. (2014), 'A simple model for predicting the signal for a head-mounted transmission chamber system, allowing imrt in-vivo dosimetry without pretreatment linac time', *Journal of Applied Clinical Medical Physics* **15**(4), 270–279.

Joiner, M. C. & Van der Kogel, A. (2016), *Basic clinical radiobiology*, Vol. 2, CRC press.

Jones, L., Hoban, P. & Metcalfe, P. (2001), 'The use of the linear quadratic model in radiotherapy: a review', *Australasian Physics & Engineering Sciences in Medicine* **24**(3), 132–146.

Jordan, T. J. & Williams, P. C. (1994), 'The design and performance characteristics of a multileaf collimator', *Physics in Medicine & Biology* **39**(2), 231.

Jordan, T. & Williams, P. (1991), Commissioning and use of a multileaf collimator system, *in* 'Proceedings of the 1st biennial ESTRO Meeting on Physics in Clinical Radiotherapy', p. 3.

Juste, B., Miro, R., Campayo, J., Diez, S. & Verdu, G. (2008), Photon spectra calculation for an elekta linac beam using experimental scatter measurements and monte carlo techniques., *in* 'Engineering in Medicine and Biology Society, 2008. EMBS 2008. 30th Annual International Conference of the IEEE', IEEE, pp. 3289–3292.

Kairn, T., Crowe, S., Poole, C. & Fielding, A. L. (2009), 'Effects of collimator backscatter in an elekta linac by monte carlo simulation', *Australasian Physics & Engineering Sciences in Medicine* **32**(3), 129.

Kajaria, A., Sharma, N., Sharma, S., Pradhan, S., Mandal, A. & Aggarwal, L. (2017), 'Monte carlo study of a flattening filter-free 6 mv photon beam using the beamnrc code'.

Karagoz, G., Zorlu, F., Yeğiner, M., Yıldız, D. & Gürkaynak, M. (2014), 'Evaluation of mlc leaf position accuracy for static and dynamic imrt treatments using david in vivo dosimetric system', *International Journal of Radiation Oncology*Biology*Physics* **90**(1).

Kawrakow, I., Fippel, M. & Friedrich, K. (1996), '3d electron dose calculation using a voxel based monte carlo algorithm (vmc)', *Medical physics* **23**(4), 445–457.

Kawrakow, I., Mainegra-Hing, E., Rogers, D. W. O., Tessier, F. & Walters, B. (2000), The egsnrc code system: Monte carlo simulation of electron and photon transport, Technical report, National Research Council Canada.

Kawrakow, I., Rogers, D. & Walters, B. (2004), 'Large efficiency improvements in beamnrc using directional bremsstrahlung splitting', *Medical physics* **31**(10), 2883–2898.

Khan, F. M. (2010), *The physics of radiation therapy*, Walters Kluwer, Lippincott Williams & Wilkins.

Kim, J.-i., Park, S.-Y., Kim, H. J., Kim, J. H., Ye, S.-J. & Park, J. M. (2014), 'The sensitivity of gamma-index method to the positioning errors of high-definition mlc in patient-specific vmat qa for sbrt', *Radiation Oncology* **9**(1), 167.

Kinsella, P., Shields, L., McCavana, P., McClean, B. & Langan, B. (2016), 'Determination of mlc model parameters for monaco using commercial diode arrays', *Journal of applied clinical medical physics* **17**(4), 37–47.

Klein, E. E., Drzymala, R. E., Purdy, J. A. & Michalski, J. (2005), 'Errors in radiation oncology: a study in pathways and dosimetric impact', *Journal of Applied clinical medical physics* **6**(3), 81–94.

Kung, J., Chen, G. & Kuchnir, F. (2000), 'A monitor unit verification calculation in intensity modulated radiotherapy as a dosimetry quality assurance', *Medical physics* **27**(10), 2226–2230.

Künzler, T., Fotina, I., Stock, M. & Georg, D. (2009), 'Experimental verification of a commercial monte carlo-based dose calculation module for high-energy photon beams', *Physics in Medicine & Biology* **54**(24), 7363.

Laub, W. U., Kaulich, T. W. & Nüsslin, F. (1999), 'A diamond detector in the dosimetry of high-energy electron and photon beams', *Physics in Medicine & Biology* **44**(9), 2183.

Lawless, C., Landau, D., Faivre-Finn, C., Boyd, K., Lester, J., Fenwick, J., Maguire, J., McCartney, E., Paul, J., Parsons, E. et al. (2017), 'Adscan: a randomised phase ii study of accelerated, dose escalated, sequential chemoradiotherapy in non-small cell lung cancer (nsclc)', *Lung Cancer* **103**, S80–S81.

Lea, D. & Catcheside, D. (1942), 'The mechanism of the induction by radiation of chromosome aberrations intradescantia', *Journal of genetics* **44**(2-3), 216–245.

Leeds University (2018), 'Marc1 – advanced research computing', `http://arc.leeds.ac.uk/systems/marc1/`. (Accessed on 08/14/2018).

Leer, J.-W. H., McKenzie, A., Scalliet, P. & Thwaites, D. I. (1998), *Practical Guidelines for the Implementation of a Quality System in Radiotherapy: A Project of the ESTRO Quality Assurance Committee Sponsored by" Europe Against Cancer"*, ESTRO.

Lillicrap, S. C., Owen, B., Williams, J. R. & Williams, P. C. (1990), 'Code of practice for high-energy photon therapy dosimetry based on the npl absorbed dose calibration service', *Physics in Medicine and Biology* **35**(10), 1355–1360.

Ling, C. C., Zhang, P., Archambault, Y., Bocanek, J., Tang, G. & LoSasso, T. (2008), 'Commissioning and quality assurance of rapidarc radiotherapy delivery system', *International Journal of Radiation Oncology* Biology* Physics* **72**(2), 575–581.

Litzenberg, D. W., Moran, J. M. & Fraass, B. A. (2002), 'Verification of dynamic and segmental imrt delivery by dynamic log file analysis', *Journal of applied clinical medical physics* **3**(2), 63–72.

Liu, H. H., Mackie, T. R. & McCullough, E. C. (2000), 'Modeling photon output caused by backscattered radiation into the monitor chamber from collimator jaws using a monte carlo technique', *Medical physics* **27**(4), 737–744.

Liu, H. H., Verhaegen, F. & Dong, L. (2001), 'A method of simulating dynamic multileaf collimators using monte carlo techniques for intensity-modulated radiation therapy', *Physics in Medicine & Biology* **46**(9), 2283.

Lobo, J. & Popescu, I. A. (2010), 'Two new dosxyznrc sources for 4d monte carlo simulations of continuously variable beam configurations, with applications to rapidarc, vmat, tomotherapy and cyberknife', *Physics in Medicine and Biology* **55**(16), 4431–4443.

Looe, H. K., Harder, D., Rühmann, A., Willborn, K. C. & Poppe, B. (2010), 'Enhanced accuracy of the permanent surveillance of imrt deliveries by iterative deconvolution of david chamber signal profiles', *Physics in Medicine and Biology* **55**(14), 3981–3992.

LoSasso, T., Chui, C.-S. & Ling, C. C. (1998), 'Physical and dosimetric aspects of a multileaf collimation system used in the dynamic mode for implementing intensity modulated radiotherapy', *Medical physics* **25**(10), 1919–1927.

LoSasso, T., Chui, C.-S. & Ling, C. C. (2001), 'Comprehensive quality assurance for the delivery of intensity modulated radiotherapy with a multileaf collimator used in the dynamic mode', *Medical physics* **28**(11), 2209–2219.

Low, D. A., Harms, W. B., Mutic, S. & Purdy, J. A. (1998), 'A technique for the quantitative evaluation of dose distributions', *Medical physics* **25**(5), 656–661.

Lyman, J. T. (1985), 'Complication probability as assessed from dose-volume histograms', *Radiation Research* **104**(2s), S13–S19.

Ma, C.-m. & Nahum, A. (1993), 'Effect of size and composition of the central electrode on the response of cylindrical ionization chambers in high-energy photon and electron beams', *Physics in Medicine & Biology* **38**(2), 267.

Ma, C. M. & Rogers, D. W. O. (2018), 'Beamdp users manual'. National Research Council of Canada.

MacDougall, N. D., Graveling, M., Hansen, V. N., Brownsword, K. & Morgan, A. (2017), 'In vivo dosimetry in uk external beam radiotherapy: current and future usage', *The British journal of radiology* **90**(1072), 20160915.

Mackay, R. I. & Williams, P. C. (2009), 'The cost effectiveness of in vivo dosimetry is not proven', *The British journal of radiology* **82**(976), 265–266.

Mani, K. R., Upadhayay, S. & Das, K. M. (2017), 'Influence of jaw tracking in intensity-modulated and volumetric-modulated arc radiotherapy for head and neck cancers: a dosimetric study', *Radiation oncology journal* **35**(1), 90.

Mans, A., Wendling, M., McDermott, L., Sonke, J.-J., Tielenburg, R., Vijlbrief, R., Mijnheer, B., Van Herk, M. & Stroom, J. (2010), 'Catching errors with in vivo epid dosimetry', *Medical physics* **37**, 2638–2644.

Mayles, P., Nahum, A. & Rosenwald, J.-C. (2007), *Handbook of radiotherapy physics: theory and practice*, CRC Press.

McDonald, D. G., Jacqmin, D. J., Mart, C. J., Koch, N. C., Peng, J. L., Ashenafi, M. S., Fugal, M. A. & Vanek, K. N. (2017), 'Validation of a modern second-check dosimetry system using a novel verification phantom', *Journal of applied clinical medical physics* **18**(1), 170–177.

McGowan & Faddegon (2013), 'Statdose for 3d dose distributions'. National Research Council of Canada.

McParland, B. J. (1989), 'A derivation of the electron mass scattering power for electron dose calculations', *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **274**(3), 592–596.

MedCom (2018), 'Medcom online: Prosoma core'. (Accessed on 02/19/2018).
URL: *www.medcom-online.de/clinical-areas-products/teletherapy/ prosoma-core/*

Meurant, G. (2012), *The dosimetry of ionizing radiation*, Vol. 3, Elsevier.

Mijnheer, B. (2008), 'State of the art of in vivo dosimetry', *Radiation protection dosimetry* **131**(1), 117–122.

Mijnheer, B., Beddar, S., Izewska, J. & Reft, C. (2013), 'In vivo dosimetry in external beam radiotherapy', *Medical Physics* **40**(7).

Mishra, B., Mishra, S., Selvam, T. P., Chavan, S. & Pethe, S. (2018), 'Comparison of measured and monte carlo calculated dose distributions from indigenously developed 6 mv flattening filter free medical linear accelerator', *Journal of medical physics* **43**(3), 162.

Mobius3D (2019), 'Mobius medical systems, lp', `http://mobiusmed.com/ mobius3d/`. (Accessed on 01/15/2019).

Moeller, R. (1989), New concept for multileaf collimation, *in* 'Proceedings of the Sixth Varian European Users Meeting, San Remo, Italy, 1989', Varian, Palo Alto, CA, pp. 126–131.

Mohammed, M., Chakir, E., Boukhal, H., Saeed, M. & El Bardouni, T. (2016), 'Evaluation of variance reduction techniques in beamnrc monte carlo simulation to improve the computing efficiency', *Journal of Radiation Research and Applied Sciences* **9**(4), 424–430.

Monseux, A., Baltieri, V., Sottiaux, A., Tomsej, M. & Leclercq, C. (2016), 'Assessment of portal dosimetry accuracy as a qa tool for vmat clinical treatment plans using dolphin/compass tools', *Physica Medica: European Journal of Medical Physics* **32**, 355–356.

Morrison, K., Tree, A., Khoo, V. & Van As, N. J. a. (2018), 'The pace trial: International randomised study of laparoscopic prostatectomy vs. stereotactic body radiotherapy (sbrt) and standard radiotherapy vs. sbrt for early stage organ-confined prostate cancer.', *Journal of Clinical Oncology* **36**, TPS153–TPS153.

Mu, G., Ludlum, E. & Xia, P. (2007), 'Impact of mlc leaf position errors on simple and complex imrt plans for head and neck cancer', *Physics in Medicine & Biology* **53**(1), 77.

Munro, A. (2007), 'Hidden danger, obvious opportunity: error and risk in the management of cancer', *The British journal of radiology* **80**(960), 955–966.

Munro, T. & Gilbert, C. (1961), 'The relation between tumour lethal doses and the radiosensitivity of tumour cells', *The British journal of radiology* **34**(400), 246–251.

Nakaguchi, Y., Ono, T., Maruyama, M., Shimohigashi, Y. & Kai, Y. (2017), 'Validation of a method for in vivo 3d dose reconstruction in sbrt using a new transmission detector', *Journal of applied clinical medical physics* **18**(4), 69–75.

NATCANSAT (2019), 'National radiotherapy dataset - rtds', `http://www.natcansat.nhs.uk/rt/rtds.aspx`. (Accessed on 02/08/2019).

National Cancer Peer Review Programme (2014), 'Radiotherapy service evidence guide', *Radiotherapy Service Evidence Guide* .

Nelson, C., Garcia, M., Calderon, E. & Kirsner, S. (2016), 'Su-f-t-306: Validation of mobius 3d and fx for elekta linear accelerators', *Medical physics* **43**(6Part17), 3533–3533.

Niemierko, A. (1997), 'Reporting and analyzing dose distributions: a concept of equivalent uniform dose', *Medical physics* **24**(1), 103–110.

Nithiyanantham, K., Mani, G. K., Subramani, V., Mueller, L., Palaniappan, K. K. & Kataria, T. (2015), 'Analysis of direct clinical consequences of mlc positional errors in volumetric-modulated arc therapy using 3d dosimetry system', *Journal of applied clinical medical physics* **16**(5), 296–305.

Oborn, B. M., Williams, M., Bailey, M. & Carolan, M. (2014), Imrt treatment monitor unit verification using absolute calibrated beamnrc and geant4 monte carlo simulations, *in* 'Journal of Physics: Conference Series', Vol. 489, IOP Publishing, p. 012020.

Oderinde, O. & du Plessis, F. (2016), 'O3. accurate monte carlo modelling of an elekta synergy equipped with an agility 160-leaf mlc', *Physica Medica* **32**, 141–142.

Ohira, S., Ueda, Y., Isono, M., Masaoka, A., Hashimoto, M., Miyazaki, M., Takashina, M., Koizumi, M. & Teshima, T. (2017), 'Can clinically relevant dose errors in patient anatomy be detected by gamma passing rate or modulation complexity score in volumetric-modulated arc therapy for intracranial tumors?', *Journal of radiation research* **58**(5), 685–692.

Oliver, M., Gagne, I., Bush, K., Zavgorodni, S., Ansbacher, W. & Beckham, W. (2010), 'Clinical significance of multi-leaf collimator positional errors for volumetric modulated arc therapy', *Radiotherapy and oncology* **97**(3), 554–560.

Otto, K. (2008), 'Volumetric modulated arc therapy: Imrt in a single gantry arc', *Medical physics* **35**(1), 310–317.

Padilla-Cabal, F., Pérez-Liva, M., Lara, E., Alfonso, R. & Lopez-Pino, N. (2015), 'Monte carlo calculations of an elekta precise sl-25 photon beam model', *Journal of Radiotherapy in Practice* **14**(3), 311–322.

Palma, D. A., Verbakel, W. F., Otto, K. & Senan, S. (2010), 'New developments in arc radiation therapy: a review', *Cancer treatment reviews* **36**(5), 393–399.

Pasler, M., Michel, K., Marrazzo, L., Obenland, M., Pallotta, S., Björnsgard, M. & Lutterbach, J. (2017), 'Error detection capability of a novel transmission detector: a validation study for online vmat monitoring', *Physics in Medicine & Biology* **62**(18), 7440.

Pena, J., Sánchez-Doblado, F., Capote, R., Terrón, J. & Gómez, F. (2006), 'Monte carlo correction factors for a farmer 0.6 cm3 ion chamber dose measurement in the build-up region of the 6 mv clinical beam', *Physics in Medicine & Biology* **51**(6), 1523.

Piermattei, A., Fidanzio, A., Azario, L., Grimaldi, L., D'Onofrio, G., Cilla, S., Stimato, G., Gaudino, D., Ramella, S., D'Angelillo, R. et al. (2007), 'Application of a practical method for the isocenter point in vivo dosimetry by a transit signal', *Physics in Medicine & Biology* **52**(16), 5101.

Pitcher, G. M., Hogstrom, K. R. & Carver, R. L. (2016), 'Radiation leakage dose from elekta electron collimation system', *Journal of applied clinical medical physics* **17**(5), 157–176.

Pitcher, G. M., Hogstrom, K. R. & Carver, R. L. (2017), 'Improved electron collimation system design for elekta linear accelerators', *Journal of applied clinical medical physics* **18**(5), 259–270.

Popescu, I. A., Shaw, C., Zavgorodni, S. F. & Beckham, W. A. (2005), 'Absolute dose calculations for monte carlo simulations of radiotherapy beams', *Physics in Medicine & Biology* **50**(14), 3375.

Poppe, B., Looe, H. K., Chofor, N., Rühmann, A., Harder, D. & Willborn, K. C. (2010), 'Clinical performance of a transmission detector array for the permanent supervision of imrt deliveries', *Radiotherapy and Oncology* **95**(2), 158–165.

Poppe, B., Thieke, C., Beyer, D., Kollhoff, R., Djouguela, A., Rühmann, A., Willborn, K. C. & Harder, D. (2006), 'David—a translucent multi-wire transmission

ionization chamber for in-vivo verification of imrt and conformal irradiation techniques', *Physics in Medicine and Biology* **51**(5), 1237–1248.

PTW (2010), 'A comparison of detectors used in detector arrays'. (Accessed: August 2018.
URL: *http://www.ptw.de/3095.html*

Qi, P. & Xia, P. (2013), 'Relationship of segment area and monitor unit efficiency in aperture-based imrt optimization', *Journal of Applied Clinical Medical Physics* **14**(3), 232–243.

QUANTEC (2010), 'Quantitative analyses of normal tissue effects in the clinic (quantec): an introduction to the scientific issues', *International Journal of Radiation Oncology Biology Physics* **76**(3), S3–S9. Bentzen, Søren M and Constine, Louis S and Deasy, Joseph O and Eisbruch, Avi and Jackson, Andrew and Marks, Lawrence B and Ten Haken, Randall K and Yorke, Ellen D.

Rangaraj, D., Zhu, M., Yang, D., Palaniswaamy, G., Yaddanapudi, S., Wooten, O. H., Brame, S. & Mutic, S. (2013), 'Catching errors with patient-specific pretreatment machine log file analysis', *Practical radiation oncology* **3**(2), 80–90.

Rangel, A. & Dunscombe, P. (2009), 'Tolerances on mlc leaf position accuracy for imrt delivery with a dynamic mlc', *Medical physics* **36**(7), 3304–3309.

RCR (2008*a*), 'Implementing in vivo dosimetry'. Institute of Physics and Engineering in Medicine, National Patient Safety Agency, Society and College of Radiographers, The Royal College of Radiologists.

RCR (2008*b*), 'On target: ensuring geometric accuracy in radiotherapy'. Royal College of Radiologists, Institute of Physics and Engineering in Medicine, Society and college of Radiographers, *Authors:* Barrett, Jane and Hoskin, P and Gaze, M and Greener, T and Kirby, M and McNair, H and Powell, M and Routsis, D and Tait, D.

Redpath, A. & Thwaites, D. (1991), 'A 3-dimensional scatter correction algorithm for photon beams (dosimetry)', *Physics in Medicine & Biology* **36**(6), 779.

Ricketts, K., Navarro, C., Lane, K., Blowfield, C., Cotten, G., Tomala, D., Lord, C., Jones, J. & Adeyemi, A. (2016), 'Clinical experience and evaluation of patient treatment verification with a transit dosimeter', *International Journal of Radiation Oncology• Biology• Physics* **95**(5), 1513–1519.

Rogers, D., Bielajew, A. et al. (1990), 'Monte carlo techniques of electron and photon transport for radiation dosimetry', *The dosimetry of ionizing radiation* **3**, 427–539.

Rogers, D. W. O. (2006), 'Fifty years of monte carlo simulations for medical physics', *Physics in Medicine and Biology* **51**(13), R287.

Rogers, D. W. O., Faddegon, B. A., Ding, G. X., Ma, C.-M., We, J. & Mackie, T. R. (1995), 'Beam: A monte carlo code to simulate radiotherapy treatment units', *Medical Physics* **22**(5), 503–524.

Rogers, D., Walters, B. & Kawrakow, I. (2001), BEAMnrc users manual, Report, National Research Council Canada.

Roussakis, Y. G. (2016), Strategies for adaptive radiotherapy: towards clinically efficient workflows, PhD thesis, University of Birmingham.

Salvat, F., Fernández-Varea, J. M. & Sempau, J. (2006), Penelope-2008: A code system for monte carlo simulation of electron and photon transport, *in* 'Workshop Proceedings', Vol. 4.

Schreiner, L., Holmes, O. & Salomons, G. (2013), Analysis and evaluation of planned and delivered dose distributions: practical concerns with γ-and χ-evaluations, *in* 'Journal of Physics: Conference Series', Vol. 444, IOP Publishing, p. 012016.

SCoR (2015), 'Radiotherapy board—intensity modulated radiotherapy (imrt) in the uk: current access and predictions of future access rates.', *Society and College of Radiographers, Institute of Physics and Engineering in Medicine, Royal College of Radiologists* . Taylor, R and Hassan, S and D'Souza, D.

Scottish Executive (2006), Unintended overexposure of patient lisa norris during radiotherapy treatment at the beatson oncology centre, glasgow in january 2006, *in* 'Report of an Investigation by the Inspector Appointed by the Scottish Ministers for the Ionising Radiation (Medical Exposures) Regulations 2000', Scottish Government.

Seco, J. & Verhaegen, F. (2016), *Monte Carlo Techniques in Radiation Therapy*, Imaging in Medical Diagnosis and Therapy, CRC Press.

Semenenko, V. A., Reitz, B., Day, E., Qi, X. S., Miften, M. & Li, X. A. (2008), 'Evaluation of a commercial biologically based imrt treatment planning system', *Medical Physics* **35**(12), 5851–5860.

Seuntjens, J., Kawrakow, I., Borg, J., Hobeila, F. & Rogers, D. (2002), Calculated and measured air-kerma response of ionization chambers in low and medium energy photon beams, *in* 'Recent Developments in Accurate Radiation Dosimetry: Proc. Int. Workshop', pp. 69–84.

Sheikh-Bagheri, D. & Rogers, D. (2002*a*), 'Monte carlo calculation of nine megavoltage photon beam spectra using the beam code', *Medical physics* **29**(3), 391–402.

Sheikh-Bagheri, D. & Rogers, D. (2002*b*), 'Sensitivity of megavoltage photon beam monte carlo simulations to electron beam and other parameters', *Medical physics* **29**(3), 379–390.

Sheikh-Bagheri, D., Rogers, D., Ross, C. K. & Seuntjens, J. P. (2000), 'Comparison of measured and monte carlo calculated dose distributions from the nrc linac', *Medical physics* **27**(10), 2256–2266.

Siantar, C. H., Walling, R., Daly, T., Faddegon, B., Albright, N., Bergstrom, P., Bielajew, A., Chuang, C., Garrett, D., House, R. et al. (2001), 'Description and dosimetric verification of the peregrine monte carlo dose calculation system for photon beams incident on a water phantom', *Medical physics* **28**(7), 1322–1337.

Sikora, M. & Alber, M. (2009), 'A virtual source model of electron contamination of a therapeutic photon beam', *Physics in Medicine & Biology* **54**(24), 7329.

Sikora, M., Dohm, O. & Alber, M. (2007), 'A virtual photon source model of an elekta linear accelerator with integrated mini mlc for monte carlo based imrt dose calculation', *Physics in Medicine & Biology* **52**(15), 4449.

Smith, K., Balter, P., Duhon, J., White, G. A., Vassy, D. L., Miller, R. A., Serago, C. F. & Fairobent, L. A. (2017), 'Aapm medical physics practice guideline 8. a.: linear accelerator performance tests', *Journal of applied clinical medical physics* **18**(4), 23–39.

Song, T., Zhou, L. & Jiang, S. (2012), 'Su-e-t-496: Monte carlo simulation of a 6mv varian truebeam without flattening filter linac', *Medical physics* **39**(6Part17), 3819–3819.

Sontag, M. R. & Cunningham, J. R. (1978), 'The equivalent tissue-air ratio method for making absorbed dose calculations in a heterogeneous medium', *Radiology* **129**(3), 787–794.

Spezi, E., Angelini, A., Romani, F. & Ferri, A. (2005), 'Characterization of a 2d ion chamber array for the verification of radiotherapy treatments', *Physics in Medicine & Biology* **50**(14), 3361.

Stasi, M., Attili, A., Brusasco, C., Cirio, R., Denis, J., Donetti, M., Evans, P., Giordanengo, S., Giuliacci, A., Givehchi, N. et al. (2010), 'Architecture and development of a detector for on-line verification in imrt.', *IBA dosimetry* .

Stathakis, S., Myers, P., Esquivel, C., Mavroidis, P. & Papanikolaou, N. (2013), 'Characterization of a novel 2d array dosimeter for patient-specific quality assurance with volumetric arc therapy', *Medical physics* **40**(7).

Stieler, F., Wolff, D., Schmid, H., Welzel, G., Wenz, F. & Lohr, F. (2011), 'A comparison of several modulated radiotherapy techniques for head and neck cancer and dosimetric validation of vmat', *Radiotherapy and Oncology* **101**(3), 388–393.

Stock, M., Kroupa, B. & Georg, D. (2005), 'Interpretation and evaluation of the γ index and the γ index angle for the verification of imrt hybrid plans', *Physics in Medicine & Biology* **50**(3), 399.

Storm, L. & Israel, H. I. (1970), 'Photon cross sections from 1 kev to 100 mev for elements z= 1 to z= 100', *Atomic Data and Nuclear Data Tables* **7**(6), 565–681.

Thoelking, J., Fleckenstein, J., Sekar, Y., Boggula, R., Lohr, F., Wenz, F. & Wertz, H. (2016b), 'Patient-specific online dose verification based on transmission detector measurements', *Radiotherapy and Oncology* **119**(2), 351–356.

Thoelking, J., Sekar, Y., Fleckenstein, J., Lohr, F., Wenz, F. & Wertz, H. (2016a), 'Characterization of a new transmission detector for patient individualized online plan verification and its influence on 6mv x-ray beam characteristics', *Zeitschrift für Medizinische Physik* **26**(3), 200–208.

Thompson, C., Weston, S., Cosgrove, V. & Thwaites, D. (2014), 'A dosimetric characterization of a novel linear accelerator collimator', *Medical physics* **41**(3).

Thwaites, D. (2013), Accuracy required and achievable in radiotherapy dosimetry: have modern technology and techniques changed our views?, *in* 'Journal of Physics: Conference Series', Vol. 444, IOP Publishing, p. 012006.

Thwaites, D. I., Mijnheer, B. & Mills, J. A. (2005), 'Quality assurance of external beam radiotherapy', *Radiation oncology physics: a handbook for teachers and students* pp. 470–50.

Toft, B. (2004), 'Independent review of the circumstances surrounding a serious adverse incident that occurred in the **REDACTED**'.
URL: *http://www.who.int/patientsafety/news/Radiotherapyadverseevent Toftreport.pdf*

Tuazon, B., Narayanasamy, G., Papanikolaou, N., Kirby, N., Mavroidis, P. & Stathakis, S. (2018), 'Evaluation and comparison of second-check monitor unit calculation software with pinnacle 3 treatment planning system', *Physica Medica* **45**, 186–191.

Tyagi, N., Yang, K., Gersten, D. & Yan, D. (2012), 'A real time dose monitoring and dose reconstruction tool for patient specific vmat qa and delivery', *Medical physics* **39**(12), 7194–7204.

Valdenaire, S., Mailleux, H. & Fau, P. (2016), 'Modeling of flattening filter free photon beams with analytical and monte carlo tps', *Biomedical Physics & Engineering Express* **2**(3), 035010.

Van de Walle, J., Martens, C., Reynaert, N., Palmans, H., Coghe, M., De Neve, W., De Wagter, C. & Thierens, H. (2003), 'Monte carlo model of the elekta sliplus accelerator: validation of a new mlc component module in beam for a 6 mv beam', *Physics in Medicine & Biology* **48**(3), 371.

Van Elmpt, W., McDermott, L., Nijsten, S., Wendling, M., Lambin, P. & Mijnheer, B. (2008), 'A literature review of electronic portal imaging for radiotherapy dosimetry', *Radiotherapy and Oncology* **88**(3), 289–309.

Van Esch, A., Clermont, C., Devillers, M., Iori, M. & Huyskens, D. P. (2007), 'On-line quality assurance of rotational radiotherapy treatment delivery by means of a 2d ion chamber array and the octavius phantom', *Medical physics* **34**(10), 3825–3837.

Van Herk, M. (2004), Errors and margins in radiotherapy, *in* 'Seminars in radiation oncology', Vol. 14, Elsevier, pp. 52–64.

Vanderstraeten, B., Chin, P. W., Fix, M., Leal, A., Mora, G., Reynaert, N., Seco, J., Soukup, M., Spezi, E., De Neve, W. et al. (2007), 'Conversion of ct numbers into tissue parameters for monte carlo dose calculations: a multi-centre study', *Physics in Medicine & Biology* **52**(3), 539.

Vatnitsky, S., Lopez, P. O., Izewska, J., Meghzifene, A. & Levin, V. (2001), 'The radiation overexposure of radiotherapy patients in panama 15 june 2001', *Radiotherapy and Oncology* **60**(3), 237–238.

Vazquez-Quino, L., Huerta-Hernandez, C. & Rangaraj, D. (2017), Clinical experience with mobius fx software for 3d dose verification for prostate vmat plans and comparison with physical measurements, *in* 'Journal of Physics: Conference Series', Vol. 847, IOP Publishing, p. 012060.

Venkataraman, S., Malkoske, K. E., Jensen, M., Nakonechny, K. D., Asuni, G. & McCurdy, B. M. (2009), 'The influence of a novel transmission detector on 6 mv x-ray beam characteristics', *Physics in Medicine & Biology* **54**(10), 3173.

Venselaar, J., Welleweerd, H. & Mijnheer, B. (2001), 'Tolerances for the accuracy of photon beam dose calculations of treatment planning systems', *Radiotherapy and Oncology* **60**(2), 191 – 201.

Verhaegen, F. & Seuntjens, J. (2003), 'Monte carlo modelling of external radiotherapy photon beams', *Physics in medicine & biology* **48**(21), R107.

Verhaegen, F., Symonds-Tayler, R., Liu, H. & Nahum, A. (2000), 'Backscatter towards the monitor ion chamber in high-energy photon and electron beams: charge integration versus monte carlo simulation', *Physics in Medicine & Biology* **45**(11), 3159.

Visser, R., Wauben, D., De Groot, M., Godart, J., Langendijk, J., van't Veld, A. & Korevaar, E. (2013), 'Efficient and reliable 3d dose quality assurance for imrt by combining independent dose calculations with measurements', *Medical physics* **40**(2).

Wack, G., Lalande, F. & Seligman, M. (2007), 'Summary of asn on the epinal radiotherapy accident', *French Nuclear Safety Authority* .

Waldron, T. (2002), The theory and operation of computer-controlled medical linear accelerator, *in* 'MEDICAL PHYSICS', Vol. 29, AMER ASSOC PHYSICISTS MEDICINE AMER INST PHYSICS STE 1 NO 1, 2 HUNTINGTON QUADRANGLE, MELVILLE, NY 11747-4502 USA, pp. 1277–1277.

Walters, B. (2016), 'Discussion concerning the observed higher dose in small voxels'.

Walters, B., Kawrakow, I. & Rogers, D. (2002), 'History by history statistical estimators in the beam code system', *Medical Physics* **29**(12), 2745–2752.

Walters, B. & Rogers, D. (2003), DOSXYZnrc users manual, Report, National Research Council Canada.

Webb, S. (2003), 'The physical basis of imrt and inverse planning', *The British journal of radiology* **76**(910), 678–689.

WHO (2008), 'Radiotherapy risk profile'. World Health Organization.

Williams, M. & McKenzie, A. (2008), 'Can we afford not to implement in vivo dosimetry?', *The British journal of radiology* **81**(969), 681–684.

Wu, H., Jiang, F., Yue, H., Hu, Q., Zhang, J., Liu, Z., Gong, J., Li, S., Geng, J. & Zhang, Y. (2016), 'A comparative study of identical vmat plans with and without jaw tracking technique', *Journal of applied clinical medical physics* **17**(5), 133–141.

Wulff, J., Heverhagen, J. & Zink, K. (2008), 'Monte-carlo-based perturbation and beam quality correction factors for thimble ionization chambers in high-energy photon beams', *Physics in Medicine & Biology* **53**(11), 2823.

Zhen, H., Nelms, B. E. & Tomé, W. A. (2011), 'Moving from gamma passing rates to patient dvh-based qa metrics in pretreatment dose qa', *Medical physics* **38**(10), 5477–5489.

Zhen, H., Nelms, B. E. & Tomé, W. A. (2013), 'On the use of biomathematical models in patient-specific imrt dose qa', *Medical physics* **40**(7).