

The Generation and Application of Monte Carlo
Calculated Beamlet Dose Distributions in
Radiation Therapy

by

Karl Kenneth Bush
B.Sc., University of Victoria, 2003.

Thesis Submitted in Partial Fulfillment of the
Requirements for the Degree of

MASTER OF SCIENCE

in the Department of Physics and Astronomy.

© Karl Kenneth Bush, 2006
University of Victoria

*All rights reserved. This thesis may not be reproduced in whole or in part,
by photocopying or other means, without the permission of the author.*

ABSTRACT

The use of beamlets as a dose calculation tool in Intensity Modulated Radiation Therapy (IMRT) treatment planning is widespread and well documented. A beamlet can simply be defined as the contribution of radiation passing through a particular geometrically defined subdivision of a given linear accelerator's emerging radiation field. The most common classes of algorithms used today to calculate the dose distributions deposited by beamlets are the pencil beam convolution and collapsed cone classes of algorithms. Using BEAMnrc [1], a Monte Carlo (MC) based radiation transport simulation software package, this thesis presents a novel method of calculating MC beamlet dose distributions with a level of accuracy not achievable using the above analytic dose calculation methods.

In a first application, the MC beamlet dose distributions generated in this thesis are used to fine-tune the output of the MC or "virtual" linear accelerator from which they are produced. This is achieved through the adjustment of individual beamlet weights to align the output of the virtual accelerator to the experimentally measured output of the modeled accelerator in water.

In a second application, MC beamlets are used to derive corrections to particular Multileaf Collimator (MLC) leaf sequences of IMRT treatment plans that have been miscalculated by a convolution-based dose calculation algorithm. These calculation inaccuracies (up to as much as 15%) arise due to the well known fact that convolution-based algorithms do not accurately model dose deposition in inhomogeneous media, such as lung [2] [3] [4].

In a final application, the MC beamlet generation method described in this thesis is implemented into a direct aperture optimization (DAO) algorithm. The implementation of MC beamlet generation in DAO forms the basis for a purely MC based inverse treatment planning system.

Supervisors:

Dr. I. A. Popescu (Department of Physics and Astronomy)

Dr. J. M. Roney (Chair, Department of Physics and Astronomy)

Contents

Abstract	ii
Table of Contents	iv
List of Figures	vii
List of Tables	xv
Acknowledgements	xvi
1 Introduction	1
1.1 The Goal of Modern Radiotherapy	1
1.2 The Physics of Ionizing Radiation	2
1.2.1 Photon Interactions in Radiotherapy	3
1.2.2 Electron Interactions in Radiotherapy	10
1.2.3 Kerma and Absorbed Dose	12
1.2.4 The Modern Radiotherapy Linear Accelerator	15
1.3 The Monte Carlo Method	21
1.4 The Importance of Accuracy in Dose Calculation	26
1.5 The Scope of This Thesis	28
2 Simulation of Radiation Therapy Plans Using BEAMnrc	30

2.1	An Introduction to BEAMnrc	30
2.2	The BEAMnrc Phase Space	32
2.3	BEAMnrc Component Modules	33
2.4	The Clinac 21EX Accelerator Model	34
2.5	Modulating the Intensity of a Phase Space	38
2.6	Creating a Monte Carlo Phantom	41
2.7	DOSXYZnrc	45
2.8	Variance Reduction Techniques in BEAMnrc	49
3	The Generation of Monte Carlo Beamlets	51
3.1	Assigning Beamlet Numbers	51
3.2	Running DOSXYZnrc With Beamlets	57
3.2.1	DOSXYZnrc Modifications For Simulating Beamlets Using LATCH Bit Filtering	58
3.2.2	DOSXYZnrc Modifications For Simulating Beamlets Using Phase Space Separation	61
3.2.3	Accounting For Uncertainties In DOSXYZnrc	62
3.3	Collecting Beamlets	66
3.4	Converting Monte Carlo Dose to Absolute Dose	67
3.5	Verification of Beamlet Dose Distributions	68
4	Applications of Beamlets	69
4.1	Fine Tuning the Output of an Accelerator	69
4.2	Remodulation of Miscalculated Treatment Plans	74
4.3	Inverse Treatment Planning Using Direct Aperture Optimization and MC Beamlets	78

5	Results and Discussion	81
5.1	Generation of Beamlets	82
5.2	The Fine-Tuning of Phase Space A	87
5.3	Treatment Plan Remodulation	128
5.4	Direct Aperture Optimization	136
6	Conclusions	143
A	NRC Monte Carlo Code Changes	150

List of Figures

1.1	Variation of photon interaction types with energy for coherent scattering, photoelectric effect, Compton scattering, and pair production interactions in water.	8
1.2	A digitally reconstructed radiograph (DRR) of a patient's hip using 80 kV photons and a portal image of the same area taken using 6 MV photons.	11
1.3	The relationship between <i>Kerma</i> and <i>Absorbed Dose</i>	14
1.4	An artists rendition of the Varian Clinac 21EX accelerator and accelerator head. (With permission from Varian Medical Systems)	15
1.5	A cutout of a medical accelerator waveguide [5].	17
1.6	A simplified view of the electric field directions in a standing wave accelerator at three different times [5].	17
1.7	A simplified view of the electric field directions in a standing wave accelerator with the zero-field cavities offset [5].	18
1.8	Schematic of the 3-piece bending magnet used to redirect the electron beam 270° toward the target.	19
1.9	The angular bremsstrahlung distributions for various incident electron energies [6].	20

1.10	A cross-section of the Clinac 21EX 6 MV flattening filter and 18 MV flattening filter.	22
1.11	The Klein-Nishina cross-section as a function of scattering angle θ plotted for various energies ($\alpha = 0.01, 0.5, 1.0, 5.0, 12.0$).	25
1.12	Monte Carlo histogram results for simulating photon scattering angles for $\alpha = 0.01, 0.5, 1.0, 5.0, 12.0$	27
1.13	Graph showing the relationship of tumour control probability (TCP) and normal tissue complication probability (NTCP).	28
2.1	An intensity modulated radiotherapy plan designed to deliver a radiation field modulated to a portrait of Albert Einstein as measured with film and calculated using Eclipse and DOSXYZnrc. [7]	31
2.2	BEAMnrc preview of 6 MV and 18 MV accelerator configurations and the materials used in each component module.	35
2.3	6 MV phase space A scored at 27.3 cm from the photon source.	36
2.4	6 MV phase space B scored at 45.0 cm from the photon source.	37
2.5	Schematics of the Varian 21EX setup. [8]	38
2.6	EGS Windows accelerator picture of the Clinac 21EX accelerator showing the position of phase space A and B and various history tracks.	39
2.7	Close-up view of the 120 leaf Multi Leaf Collimator on the 21EX conforming to an arbitrary tumour volume. [Varian Medical Systems]	42
2.8	A comparison of the resulting dose distributions from a standard conformal 3-field treatment of a head and neck tumour and a 7-field Intensity Modulated Radiation Therapy (IMRT) plan for the same patient	43

2.9	The default ramp for converting CT values to material and density in CTCreat [9]. (With permission from the National Research Council of Canada)	44
2.10	An axial CT image containing traces of the CT couch and head support, the CT image after all material outside of the body contour has been set to the optical density of air and and axial phantom slice after cropping to the body contour.	45
2.11	The spherical coordinate system used in DOSXYZnrc.	46
2.12	Monte Carlo dose distribution for a clinical treatment plan overwritten into the CadPlan dose distribution for comparison with CadPlan. . .	50
3.1	The grid used to assign beamlet numbers.	54
3.2	Superimposing the collimating jaws and a grid onto phase space A. .	55
3.3	A single beamlet within the defined field and beamlets 1601 and a random beamlet from a phase space of 1601 beamlets turned on using the modified inclusive/exclusive LATCH filter.	56
3.4	Monte Carlo dose distributions scored for various beamlets within an artificial lung interface phantom.	63
5.1	Comparison of beamlet and non-beamlet parallel plane profiles of a 4×4 cm ² field incident on an artificial lung interface.	83
5.2	Comparison of beamlet and non-beamlet depth dose profiles of a 4×4 cm ² field incident on an artificial lung interface.	84
5.3	The 6 MV percent depth dose (PDD) curves for the best matched incident electron energy (6 MeV) and full width half maximum (0.75 mm).	89

5.4	The 18 MV percent depth dose (PDD) curves for the best matched incident electron energy (18 MeV) and full width half maximum (0.75 mm).	90
5.5	Initial 6 MV Monte Carlo and CadPlan profiles from a 40×40 cm ² field at depths of 1.5 cm (d_{max}), 10.0 cm, and 20.0 cm in a $50.0 \times 50.0 \times 50.0$ cm ³ water tank with a source-to-surface distance (SSD) of 100.0 cm. .	94
5.6	Corrected 6 MV, 40.0×40.0 cm ² Monte Carlo and CadPlan profiles at a depth of 1.5 cm (d_{max}), in a $50.0 \times 50.0 \times 50.0$ cm ³ water tank (100.0 cm SSD).	95
5.7	A closer look at the 1.5 cm (d_{max}) profiles of Figure 5.5.	96
5.8	A closer look at the 1.5 cm (d_{max}) profiles of Figure 5.6.	97
5.9	6 MV, 40.0×40.0 cm ² Monte Carlo and CadPlan profiles at depths of 1.5 cm (d_{max}), 10.0 cm, and 20.0 cm in a $50.0 \times 50.0 \times 50.0$ cm ³ water tank (100.0 cm SSD) after weighting the incident phase space A and re-running the simulation with different random number seeds.	98
5.10	6 MV, 40×40 cm ² initial Monte Carlo, measured, and CadPlan profiles at a depth of 1.5 cm (d_{max}) in a $50.0 \times 50.0 \times 50.0$ cm ³ water tank (100.0 cm SSD).	100
5.11	6 MV, 40×40 cm ² initial Monte Carlo and measured profiles at a depth of 1.5 cm (d_{max}) in a $50.0 \times 50.0 \times 50.0$ cm ³ water tank (100.0 cm SSD) with higher voxel resolution.	101
5.12	Corrected 6 MV, 40×40 cm ² Monte Carlo and measured profiles at a depth of 1.5 cm (d_{max}) in a $50.0 \times 50.0 \times 50.0$ cm ³ water tank (100.0 cm SSD).	102
5.13	Delta profile plot of the initial dose differences between MC and measured dose distributions for a 6 MV beam in a plane at 1.5 cm depth.	103

5.14	Delta profile plot of the dose differences between corrected MC and measured dose distributions for 6 MV in a plane at 1.5 cm depth.	104
5.15	Delta profile plot of the dose differences between re-run MC and measured dose distributions for 6 MV in a plane at 1.5 cm depth.	105
5.16	Measured and MC profiles resulting from re-running the 6 MV weighted phase space through $10 \times 10 \text{ cm}^2$ jaws.	107
5.17	Measured and MC profiles resulting from re-running the 6 MV weighted phase space through $4 \times 4 \text{ cm}^2$ jaws.	108
5.18	Measured and MC percent depth dose curves resulting from re-running the 6 MV weighted phase space through $10 \times 10 \text{ cm}^2$ jaws in a $50.0 \times 50.0 \times 22.5 \text{ cm}^3$ water tank (100.0 cm SSD).	109
5.19	18 MV, $40 \times 40 \text{ cm}^2$ initial MC, measured, and CadPlan profiles at a depth of 3.3 cm (d_{max}) in a $50.0 \times 50.0 \times 50.0 \text{ cm}^3$ water tank (100.0 cm SSD).	111
5.20	18 MV, $40 \times 40 \text{ cm}^2$ MC and measured ion chamber profiles at a depth of 3.5 cm (d_{max}), in a $50 \times 50 \times 50 \text{ cm}^3$ water tank (100.0 cm SSD).	112
5.21	Optimized 18 MV, $40 \times 40 \text{ cm}^2$ MC and measured ion chamber profiles at a depth of 3.5 cm (d_{max}) in a $50 \times 50 \times 50 \text{ cm}^3$ water tank (100.0 cm SSD).	113
5.22	18 MV, $40 \times 40 \text{ cm}^2$ MC and measured ion chamber profiles at a depth of 3.5 cm (d_{max}) in a $50 \times 50 \times 50 \text{ cm}^3$ water tank (100.0 cm SSD) after weighting the incident phase space A and re-running the simulation with different random number seeds.	114
5.23	A closer look at the 3.5 cm (d_{max}) profile in Figure 5.20.	115
5.24	A closer look at the 3.5 cm (d_{max}) profiles in Figure 5.21.	116

5.25	Delta profile plot of the initial dose differences between MC and measured dose distributions for 18 MV in a plane at 3.5 cm depth.	117
5.26	Delta profile plot of the dose differences between corrected MC and measured dose distributions for 18 MV in a plane at 3.5 cm depth.	117
5.27	Measured and MC percent depth dose curves resulting from re-running the 18 MV weighted phase space through 10 × 10 cm ² jaws.	118
5.28	Measured and MC percent depth dose curves resulting from re-running the 18 MV weighted phase space through 4 × 4 cm ² jaws.	119
5.29	Measured and MC percent depth dose curves resulting from re-running the 18 MV weighted phase space through 10 × 10 cm ² jaws in a 50.0 × 50.0 × 50.0 cm ³ water tank (100.0 cm SSD).	120
5.30	Dose subtraction between the original 18 MV Monte Carlo and ion chamber measurements in a plane at d_{max} (3.5 cm) for a 40 × 40 cm ² field prior to fine-tuning (normalized to maximum dose difference).	122
5.31	Dose subtraction between the corrected Monte Carlo and ion chamber measurements in a plane at d_{max} (3.5 cm) for an 18 MV, 40 × 40 cm ² field (normalized to dose maximum prior to fine-tuning).	122
5.32	Dose subtraction between uncorrected Monte Carlo and ion chamber measurements in a plane at d_{max} (1.5 cm) for an 6 MV, 40 × 40 cm ² field prior to fine-tuning (normalized to maximum dose difference).	123
5.33	Dose subtraction between the corrected Monte Carlo and ion chamber measurements in a plane at d_{max} (1.5 cm) for an 6 MV, 40 × 40 cm ² field (normalized to dose maximum prior to fine-tuning).	123
5.34	18 MV, 40 × 40 cm ² Monte Carlo and CadPlan profiles at depths of 3.5 cm (d_{max}), 10 cm, and 20 cm in a 50 × 50 × 50 cm ³ water tank (100.0 cm SSD).	125

5.35	Optimized 18 MV, 40×40 cm ² Monte Carlo and CadPlan profiles at a depth of 3.5 cm (d_{max}) in a $50 \times 50 \times 50$ cm ³ water tank (100.0 cm SSD).	126
5.36	18 MV, 40×40 cm ² Monte Carlo and CadPlan profiles at depths of 3.5 cm (d_{max}), 10 cm, and 20 cm in a $50 \times 50 \times 50$ cm ³ water tank (100.0 cm SSD) after weighting the incident phase space A and re-running the simulation with different random number seeds.	127
5.37	An axial CT image of the artificial lung interface phantom used for remodulating miscalculated treatment plans and an image of the reconstructed 3-D CT phantom.	130
5.38	Monte Carlo dose profiles resulting from 4×4 cm ² field incident on the lung-air interface.	131
5.39	Film, CadPlan (TPS), and MC dose profiles taken at a lung-tissue interface in an artificial phantom at a depth of 10 cm and 15 cm for 18MV.	131
5.40	Dose volume histograms prior to optimization within the lung-air interface PTV for Eclipse and Monte Carlo dose distributions.	132
5.41	Dose volume histograms after optimization within the lung-air interface PTV for Eclipse and Monte Carlo dose distributions.	133
5.42	Close-up dose subtraction between Eclipse and Monte Carlo dose distributions prior to optimization within the lung-air interface PTV.	134
5.43	Close-up dose subtraction after optimization within the lung-air interface PTV for Eclipse and Monte Carlo dose distributions.	134
5.44	Field set-up for a 7-field head and neck IMRT plan.	139
5.45	Beam's eye view of optimized apertures from Field 2 for the 7-Field direct aperture optimization [10].	140

5.46	Dose distributions from direct aperture optimization using MC beamlets and the pencil beam convolution calculated dose distribution using the MC determined apertures.	141
5.47	Optimized dose volume histograms within the structures defined in Figure 5.44 for MC calculated, PBC calculated and re-run MC calculated dose distributions [10].	142

1.1	Variation of absorption probability with photon energy, atomic number, electron density, and mass density.	7
-----	--	---

6.1	The total simulation times for LATCH bit filtered, phase space separated, unmodified BEAMnrc and DOSEXYZnrc, and cropped photon simulations for small and large size phase spaces.	85
-----	--	----

List of Tables

1.1	Variation of interaction probability with photon energy, atomic number, electron density, and mass density.	7
5.1	The total simulation times for LATCH bit filtered, phase space separated, unmodified BEAMnrc and DOSXYZnrc, and cropped phantom simulations for small and large size phase spaces.	85

I would like to thank my supervisor, Tony Popescu, for his wealth of brilliant ideas and for his support throughout this project. Thanks to Sergei Zavgorodni for the countless hours spent working out many of the details to make this project work. Thanks to Conor Shaw for just being Conor Shaw and to Gabe Rosenbaum for pretending my project was better than his. Special thanks to Wayne Beckham for bringing together an amazing group of physicists, creating a top notch work environment, and for being such an all around cool guy.

Chapter 1

Introduction

In order to familiarize readers with some of the concepts needed to understand the topics presented in this thesis, a brief introduction has been included. The introduction will define many of the key concepts used in medical physics as well as provide a brief introduction to Monte Carlo methods.

1.1 The Goal of Modern Radiotherapy

The primary goal of modern radiotherapy is to improve the lives of patients through the application of ionizing radiation. In radiotherapy, many forms of ionizing radiation are used to treat many different diseases; however, the treatment of cancerous tumors using x-rays is the most widely used form of radiotherapy. Beta, proton and other heavy ion radiations are also used in therapeutic applications.

1.2 The Physics of Ionizing Radiation

Consider a beam of ionizing photons incident on an infinitesimal slab of material of width, dx . The number of photons removed from the slab by scatter or absorption, dN , is related to the total number of incident photons by:

$$dN = -\mu N dx \quad (1.1)$$

where the constant of proportionality, μ , is called the *linear attenuation coefficient*. The physical dimension of μ is $[\text{length}]^{-1}$. If N_0 photons are incident on a slab of thickness x , then the number of photons remaining in the beam can be found by integrating Equation 1.1:

$$\int_{N_0}^N \frac{dN}{N} = -\mu \int_0^x dx$$

which gives

$$N = N_0 e^{-\mu x} \quad (1.2)$$

The linear attenuation coefficient, μ , is dependent on both the energy and density of the slab of material. It is therefore convenient to normalize μ by the material density. μ/ρ is called the *mass attenuation coefficient*. Equation 1.3 then becomes

$$N = N_0 e^{-\frac{\mu}{\rho} t} \quad (1.3)$$

where $t = x\rho$ is called the material "thickness".

1.2.1 Photon Interactions in Radiotherapy

In radiotherapy, incident photon and electron energies typically span from ~ 100 keV up to ~ 25 MeV¹. For photons within this range, several different ionizing interactions can occur, including coherent scattering, photoelectric absorption, Compton scattering, and pair production. Coherent scattering occurs when a photon is absorbed by an atom which then emits another photon in a slightly different direction; hence the incoming photon appears to have been scattered elastically with no loss of energy. As this effect occurs more probabilistically in high atomic number (Z) materials at low energies, it is negligible for mega-voltage photons (as in therapeutic x-rays) interacting in tissue (low Z).

As the energy of photons is increased, coherent scattering is less likely to occur. Instead, the dominant process in which photons interact is the photoelectric effect (discovered by Hertz in 1887, explained by Einstein in 1905). In the photoelectric effect, a photon of incident energy $h\nu$ collides with an atom and ejects a bound electron from one of the K, L, M, or N shells. The ejected electron emerges with an energy of $h\nu - E_s$ where E_s is the binding energy of the shell from which the electron is ejected. As a result of this process, the atom is left in an excited state, and when an electron from a higher energy shell fills the lower energy vacancy a “characteristic” photon is emitted (the term characteristic is used since the photon is characteristic of the energy difference between shells) or, in some cases, an electron is emitted from

¹Up to ~ 300 MeV for protons.

an outer shell (termed an Auger electron). The energy of this ejected electron is also $h\nu - E_s$ with $h\nu$ being the incident characteristic photon energy and E_s being the electron's binding energy. The ratio of the number of characteristic photons emitted to Auger electrons emitted is known as the fluorescent yield. The probability of photoelectric interaction, or photoelectric cross-section, varies with photon energy approximately as $1/(h\nu)^3$. For high atomic number (Z) materials the probability of interaction per unit length travelled varies approximately with Z^3 . For low Z materials the interaction probability varies approximately with $Z^{3.8}$ [6].

As the energy of incident photons ($h\nu$) is further increased (to ~ 30 keV in tissue), the photoelectric effect is less likely to occur and the dominant process becomes Compton Scattering. Sir Arthur H. Compton (1923) discovered experimentally and explained theoretically the following phenomenon. If x-ray photons interact with a target, for certain scattering angles the wavelength of the scattered photons is greater than the incident wavelength.

For this interaction, conservation of energy yields:

$$h\nu_0 = h\nu + E_{kin} = h\nu + \gamma mc^2 - mc^2$$

where E_{kin} is the kinetic energy of the recoiling electron (with rest mass m and speed v), γ is the relativistic factor

$$\gamma = \left(1 - \frac{v^2}{c^2}\right)^{-\frac{1}{2}}$$

and ν_0 and ν are the frequencies of the incident and scattered photon, respectively.

Conservation of momentum yields:

$$\vec{p}_0 = \vec{p} + \vec{p}_e$$

where \vec{p}_0 and \vec{p} are the momenta of the incident and scattered (at angle θ) photon, respectively, and \vec{p}_e is the momentum of recoiling electron (at angle ϕ). Using both conservation of energy and momentum, the change in the electron's wavelength before and after scattering is found to be:

$$\lambda - \lambda_0 = \frac{h}{mc}(1 - \cos \theta)$$

or

$$\Delta\lambda = \Lambda(1 - \cos \theta)$$

where $\Lambda = \frac{h}{mc}$ is called the Compton wavelength of the electron.

The kinetic energy and the angle ϕ of the recoiling electron can also be derived using conservation of momentum and energy as:

$$h\nu = h\nu_0 \frac{1}{1 + \alpha(1 - \cos \theta)}$$

$$\tan \phi = \frac{\cot \frac{\theta}{2}}{1 + \alpha}$$

where

$$\alpha = \frac{\Lambda}{\lambda_0}$$

For an electron we can write

$$\alpha = \frac{h\nu_0}{0.511\text{MeV}}$$

Analysis of the kinetic energy formula shows that when a low-energy photon undergoes Compton Scattering, most of the energy of the incident photon is retained by the scattered photon and only a small fraction of the energy is transferred to the electron. When a high-energy photon (as in a radiotherapy beam) undergoes Compton scattering, most of the energy is transferred to the recoiling electron and only a small fraction of the energy of the incident photon is retained by the scattered photon.

The differential cross-section for Compton scattering was calculated using quantum mechanics by Klein and Nishina in 1929 and later confirmed by quantum field theory to be:

$$\frac{d\sigma_e}{d\Omega} = \frac{d\sigma_{cl}}{d\Omega} F_{KN} \quad (1.4)$$

where σ_{cl} is the classical electron cross-section and

$$F_{KN} = \left(\frac{1}{1 + \alpha(1 - \cos \theta)} \right)^2 \left(1 + \frac{\alpha^2(1 - \cos \theta)^2}{(1 + \alpha(1 - \cos \theta))(1 + \cos^2 \theta)} \right) \quad (1.5)$$

is the quantum mechanical correction to σ_{cl} .

The last dominant photon interaction that occurs in radiotherapy is electron-positron pair production. This interaction becomes significant in tissue at ~ 5 MeV. When a photon passes near a nucleus and is subjected to the strong field of the nucleus, it can undergo a spontaneous transition into an electron-positron pair. Since the rest mass energy of the electron-positron pair is $2(0.511 \text{ MeV}) = 1.022 \text{ MeV}$, the energy of the incident photon must be greater than 1.022 MeV for the interaction to occur. Excess energy is shared among the pair as kinetic energy. Since the interaction

is considered a collision between the photon and nucleus, some energy is transferred to the recoiling nucleus during the collision. The amount is, however, small in comparison to the electron and positron. Once created, the positron ionizes and excites atoms as it travels through material in a similar way to that of an electron (as will be described in section 1.2.2). As it is finally brought to rest, the positron annihilates with a free electron, emitting opposing 0.511 MeV photons. The total energy given to the charged particles can be divided up in many ways. For this reason one cannot characteristically analyze the scattering angles and energies for this interaction as in Compton scattering.

The pair production cross-section increases rapidly with increase in energy above the 1.022 MeV threshold needed to create the $e^+ e^-$ pair. This is in fact opposite to the photoelectric and Compton processes which have a cross-section that decreases with increasing energy. Since pair production occurs in the field of the nucleus, which has a cross-section that varies with Z^2 , the cross-section for pair production also varies with Z^2 .

Interaction mode	$h\nu$	Z	ρ_e	ρ
Photoelectric effect	$(h\nu)^{-3}$	$Z^{3-3.8}$	—	ρ
Compton effect	$(h\nu)^{-1}$	—	ρ_e	ρ
Pair production	$h\nu$	Z^2	—	ρ

Table 1.1: Variation of interaction probability with photon energy, atomic number, electron density, and mass density [6].

The probability for each of the above interactions occurring is proportional

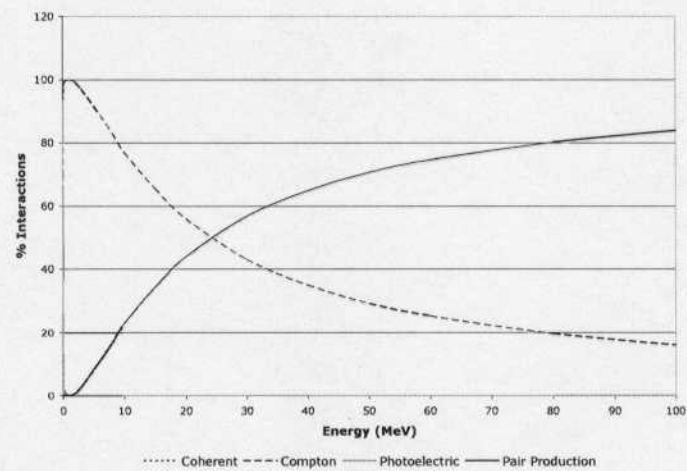



Figure 1.1: Variation of photon interaction types with energy below 200 keV (*top*) and over the full radiotherapy energy range (*bottom*) for coherent scattering, photoelectric effect, Compton scattering, and pair production interactions (Data from Johns and Cunningham [6]).

to the cross-section for that interaction. The total probability of interaction is then proportional to the sum of all the individual interaction cross-sections. Figure 1.1 shows the variation of relative frequency of the different photon interactions with energy in water (which closely approximates tissue). From this Figure one can see that up to about 30 keV, the photoelectric effect is dominant. From 30 keV to 90 keV, both the photoelectric effect and Compton scattering are dominant. From 200 keV to 2 MeV, Compton scattering is dominant. Pair production becomes important above 5 MeV. Above 50 MeV, pair production is the dominant interaction type.

From the above summary one can understand the reasons for selecting the range of energies used for various clinical applications. For radiographic imaging, choosing a photon energy in the < 90 keV range, where the photoelectric effect dominates and the cross-section varies approximately with $Z^{3.8}$, allows for better image contrast. This can be seen clearly by looking at the ratio of attenuation between bone ($Z_{eff} = 12.31$) and tissue ($Z_{eff} = 7.51$) in this energy range:

$$\frac{12.31^{3.8}}{7.51} = 6.54 \quad (1.6)$$

Photons in this energy range are 6.54 times more likely to be photoelectrically absorbed in bone than in tissue. When taking a radiographic image, areas on the film behind bone therefore are much less exposed (since there are many fewer photons that travel all the way through a patient without interacting) (see Figure 1.2). When taking a portal image² using MeV X-rays, the contrast is dramatically decreased. This is

²A film exposure, Amorphous Silicon diode array exposure, or other exposure created from x-rays

because in this energy range, Compton scattering dominates the photon interactions. Since Compton scattering varies only with ρ_e , which is not significantly different for various materials in the body, the image contrast is poor.

1.2.2 Electron Interactions in Radiotherapy

After an ionizing photon interaction occurs in a material, the electrons that have been set in motion lose energy by collisional losses or by the emission of bremsstrahlung radiation. These interactions will now be outlined.

The mechanism for collisional energy loss involves the ionization and excitation of atoms. These events give rise to local deposition of energy. The collisions occur mostly between secondary electrons (lower energy electrons that have been set in motion by high energy primary electrons) and outer shell electrons of the atoms in a medium. More frequently, the electron loses a small amount of energy in each of the successive collisions. The rate of energy loss per collision depends on electron energy, number of atomic electrons per unit volume, and on the ionization energy of the atoms. Less frequently, a large fraction of the secondary electron energy is transferred to an orbital electron in a single collision. The ejected electron, referred to as a δ -ray, is a secondary electron in its own right, causing excitations and ionizations along its track. Such collisions are known as *Møller* scattering events. When the δ -ray is produced by a positron, the process is called *Bhabha* scattering. The scattering

exiting a patient during a radiotherapy treatment.




Figure 1.2: A digitally reconstructed radiograph (DRR) of a patient's hip using 80 kV photons (*top left*) and a portal image of the same area taken using 6 MV photons (*top right*). The difference in image contrast is due to the difference in dominant photon interaction processes. Similarly, a 80 kV DRR taken of a patient's pelvis (*bottom left*) and a 6 MV portal image of the same region (*bottom right*). One can see in this image a pocket of rectal gas appears clearly (due to a smaller effective attenuation length). One can also see that this patient is slightly misaligned on the treatment couch relative to where the planned radiation is to be delivered.

cross-sections for these two processes can be calculated to a very high precision using quantum electrodynamics.

When a high energy electron passes close to the nucleus of an atom, it experiences a Coulomb force. The path of this electron is deflected and its speed is reduced as it loses energy and “falls” some amount toward the nucleus. The electron recedes from the interaction with a lower energy which is emitted as a photon of energy $h\nu$. This sudden slowing of the electron is known as bremsstrahlung (German for “braking radiation”). The cross-section for bremsstrahlung increases with energy and is proportional to Z^2 . The spectrum of bremsstrahlung photons is continuous, with the maximum photon energy less than or equal to the incident electron energy (as required by energy conservation). Bremsstrahlung radiation is the main mechanism for generating photons in the accelerator, as described in section 1.2.4.

1.2.3 Kerma and Absorbed Dose

The energy transfer from a photon to a medium can be regarded as a two step process:

- (i) The interaction of a photon (of ionizing energy) with an atom sets in motion one or more high energy *primary electrons*.
- (ii) Each primary electron transfers its energy to the medium through excitations and secondary ionizations.

The quantity characterizing step (i) is *Kerma*, which stands for Kinetic energy relaxed per unit mass,

$$K = \frac{dE_{tr}}{dm}, \quad (1.7)$$

where dE_{tr} is the energy transferred from photons to the primary electrons in a mass element dm of the medium. Kerma is measured in J/kg. The quantity that characterizes the interactions in step (ii) along the range of the primary electron is the *Absorbed Dose*. *Absorbed Dose* is also measured in J/kg or Gy (Gray) where 1 Gy=1 J/kg.

In general, at any given depth, *Kerma* does not equal *Absorbed Dose*. To help illustrate the difference between *Kerma* and *Absorbed Dose* suppose that a field of ionizing radiation is incident from air onto a slab of material. Up to a given depth within the material (which depends on the beam energy and material type), *Kerma* is higher than *Absorbed Dose*. This is because, while many electrons are being set in motion in the material near the surface (*Kerma*), there are few electrons coming from upstream (outside the surface of the material) to deposit energy (*Absorbed Dose*). This region is known as the build-up region (see Figure 1.3). With depth, the number of electrons in motion that deposit *Absorbed Dose* increases up to a maximum, d_{max} , after which *Absorbed Dose* decreases exponentially as photons are exponentially attenuated. Figure 1.3 illustrates the *Absorbed Dose-Kerma* relationship. Beyond d_{max} , *Kerma* drops below *Absorbed Dose*. This is due to the fact that electrons set in motion at any point in a material generally travel “downstream” as they deposit dose. Because of

photon attenuation, there are always more electrons set in motion just “upstream” which deposit *Absorbed Dose* “downstream” [6]. It has been shown that, ignoring bremsstrahlung radiation energy losses by travelling electrons, d_{max} occurs precisely at the point at which *Absorbed Dose* = *Kerma* [11]. When including bremsstrahlung, the point at which *Absorbed Dose* = *Kerma* is shifted to the right of d_{max} (deeper) [11].




Figure 1.3: The relationship between *Kerma* and *Absorbed Dose*.

Another quantity, *Terma*, is an acronym for Total energy relaxed per unit mass, due to primary interactions. It represents the energy of the primary photon that is (a) imparted to secondary charged particles and (b) retained by the scattered

photon. In other words, terma is the energy that leaves the primary beam per unit mass.

1.2.4 The Modern Radiotherapy Linear Accelerator

A basic overview of the acceleration of electrons and production of x-rays within a medical linac will now be covered. A typical linac (and also the one used for this thesis) is shown in Figure 1.4.




Figure 1.4: An artists rendition of the Varian Clinac 21EX accelerator (*left*) and accelerator head (*right*). [Varian Medical Systems]

Accelerating Electrons and Production of X-Rays

The acceleration of electrons in the medical linac is achieved through the application of microwaves that have been confined and structured by the use of a waveguide. The input microwaves are generated using a klystron and are typically in the 3000 MHz or S-band range (2-4 GHz range)[5]. Both traveling wave and standing wave designs are used in medical linacs. The accelerator used for this thesis is a standing wave design. The waveguide for a standing wave accelerator consists of a series of cylindrical accelerating cavities of lengths varying from 2.5 cm to 5.0 cm corresponding to the half wavelengths of the input microwaves (see Figure 1.5). To establish the standing wave both a forward traveling wave and backward traveling wave are arranged, each of which is reflected at the ends of the accelerator. Using this setup the two moving electric field maxima forward align 1/4 of the time, reverse align 1/4 of the time, and cancel each other out 1/2 the time, as shown in Figure 1.6. Because the zero-field cavities do not contribute to particle acceleration it is very common to move these cavities off axis, out of the particle path but still couple power between cavities. This is shown in Figure 1.7. By doing so the overall length of the accelerator can be significantly reduced. In order to prevent electrical arcing between disks in the waveguide the entire accelerating cavity is kept under vacuum.

By varying the aperture and length of the cavities initially traversed, the continuum of injected electron velocities delivered from the electron gun are concentrated




Figure 1.5: A cutout of a medical accelerator waveguide [5].




Figure 1.6: A simplified view of the electric field directions in a standing wave accelerator at three different times [5].




Figure 1.7: A simplified view of the electric field directions in a standing wave accelerator with the zero-field cavities offset [5].

into discrete bunches during their acceleration. Beyond the bunching region, the velocity of the electrons remains almost constant and near the speed of light. The waveguide cavities are therefore made uniform for the remainder of the acceleration period. As bunching technology improves, a greater proportion of the incident electrons is able to be captured and bunched. Current technology allows roughly one third of the incident electrons to be successfully captured, bunched, and accelerated over the length of the waveguide.

Once the electrons have reached the end of the acceleration phase they are directed toward the target by means of bending and steering magnets. In the case of the Varian Clinac series of accelerators used in this thesis the electrons are directed 270° toward the target (see Figure 1.8). The purpose of this is two-fold: to filter out and prevent low energy contaminant electrons from hitting the target, and to allow a more compact accelerator design.




Figure 1.8: Schematic of the 3-piece bending magnet used to redirect the electron beam 270° toward the target. Energy selection slits S1 and S2 can be adjusted radially inward or outward to change the accepted range of electron energies [5].

The target material used to produce x-rays is tungsten. Tungsten is chosen for its high atomic number and resistance to heat deformation. The ability to resist heat deformation is important since in a typical tungsten target only $\sim 1\%$ of the incident electron energy emerges as bremsstrahlung radiation. The rest is lost to heat in the target. This heat must be dissipated by the accelerator via a cooling system. The angular distribution of bremsstrahlung radiation is shown in Figure 1.9.




Figure 1.9: The angular bremsstrahlung distributions for various incident electron energies [6].

Shaping, Flattening and Modulating the Intensity of the Beam

Initially the bremsstrahlung lobe of x-rays produced from the target is shaped into a conical beam using a lead collimator. In the case of the Varian 21EX, the angle of this collimator (14.3°) was chosen such that a $35.0 \times 35.0 \text{ cm}^2$ square can completely

fit inside the circular projection of the primary collimator cone at a distance of 100.0 cm from the target.

The energy spectrum of the conical beam is then “flattened” using a flattening filter. The flattening filter is needed because the bremsstrahlung lobe is not radially uniform in energy and particle fluence. The 6 MV and 18 MV flattening filters used in this thesis are shown in Figure 1.10.

1.3 The Monte Carlo Method

The Monte Carlo method is a method of approximately solving mathematical and physical problems by the simulation of random variables [12]. Given a specific probability density function (PDF), the ability to generate random numbers in proportion to the distributed probabilities provides a means to simulate physical systems. This is achieved in the Monte Carlo method by transforming a distribution of uniformly distributed random numbers (which can easily be generated) into the desired distribution. The method of transforming random variables will now be outlined.

Consider the known probability density function $p(x)$, for which it is desired to generate random values X , distributed over the interval $[X_{min}, X_{max}]$ with density $p(x)$. It can be shown that these values of X are given by the equation:

$$\int_{X_{min}}^X p(x)dx = G \quad (1.8)$$

where G is a set of uniformly distributed numbers, and X_{min} is the lower bound of

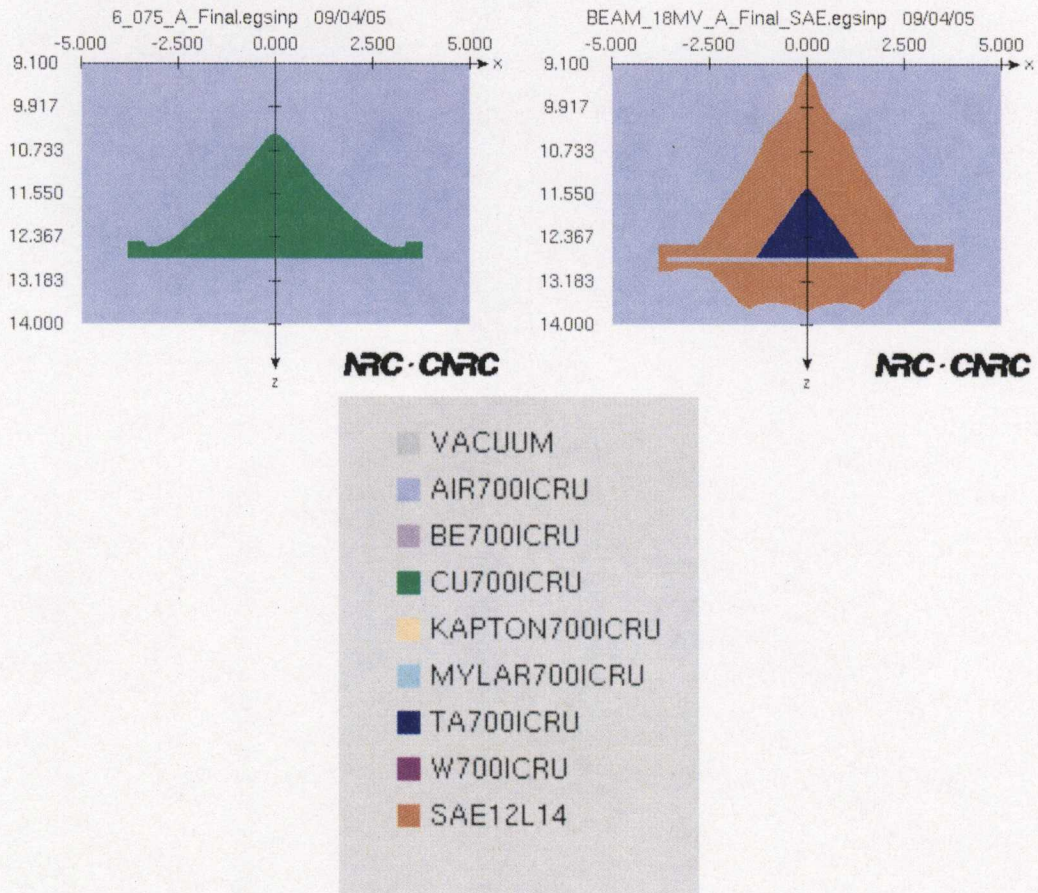


Figure 1.10: A cross-section of the Clinac 21EX 6 MV flattening filter (*left*) and 18 MV flattening filter (*right*). Both filters are cylindrically symmetric about the Z-axis. The 6 MV filter is composed completely of Copper. The 18MV filter is composed of a combination of SAE12L14 (steel alloy) and a Tantalum core.

the range of desired X values [12].

As an example suppose we take the probability density function (PDF) for determining the probability of interaction a photon penetrating a slab of material might have at depth x . Assuming that the slab is infinitely thick, the exponentially distributed PDF is:

$$p(x) = \mu e^{-\mu x}, \quad 0 \leq x \leq \infty \quad (1.9)$$

where μ is the total *linear attenuation coefficient* (in cm^{-1}) for all photon interaction types (pair production, Compton scattering, photoelectric effect, coherent scattering) for a given energy and material density.

In order to simulate the physical process using the Monte Carlo method, one could solve for X in:

$$G = \int_0^X \mu e^{-\mu x} dx \quad (1.10)$$

$$G = 1 - e^{-\mu X}$$

$$X = -\frac{1}{\mu} \ln(1 - G) \quad 0 \leq G < 1$$

By generating random numbers from the flat distribution G , one can simulate random numbers X distributed as $p(x)$.

As another example, consider the Monte Carlo modeling of Compton scattering. More specifically, suppose a way to generate random photon scattering angles θ from an incident photon energy $h\nu$ is desired. To accomplish this one could use the

Klein-Nishina formula as a PDF:

$$\frac{d\sigma}{d\theta} = \frac{1}{\sigma_c} \frac{\pi \sin(\theta) r^2 \left(1 + \cos^2(\theta) + \alpha^2 \frac{(1 - \cos(\theta))^2}{1 + \alpha(1 - \cos(\theta))} \right)}{(1 + \alpha(1 - \cos(\theta)))^2} \quad (1.11)$$

Where σ_c is a normalization factor, r is the classical electron radius, θ is the photon scattering angle, and α is the ratio of incident photon energy to electron rest mass. Note, the azimuthal dependence of the solid angle has been integrated out. The Klein-Nishina cross-section as a function of scattering angle θ has been plotted for various energies ($\alpha = 0.01, 0.5, 1.0, 5.0, 12.0$) in Figure 1.11.

The next step in the above Monte Carlo method would involve integrating the PDF and solving for X as in Equation 1.8. However, this is not an easy task since this integral is not easily inverted. Instead, one could generate the random θ 's using what is known as the *Acceptance-Rejection* method. It is easiest to understand this method geometrically by constructing a rectangular area which completely encloses the area under the PDF. Given a maximum $p(\theta)$ of P_{MAX} , the boundaries of the rectangle would be $0 \leq p(\theta) \leq P_{MAX}$ and $0 \leq \theta \leq \pi$. Suppose N uniformly distributed random points are chosen within this rectangular area by generating a pair of uniform random numbers H' and H'' in the interval $0 \leq H' \leq \pi$ and $0 \leq H'' \leq P_{MAX}$ to form the coordinates of the point. For each $\theta = H'$ generated, Equation 1.11 is evaluated and compared to H'' . In the event that H'' is less than $\frac{d\sigma(H')}{d\theta}$, the random value H' is accepted, otherwise it is rejected and another set of H' and H'' are generated. In this way a random set of values distributed as Equation 1.11 can be generated




Figure 1.11: The Klein-Nishina cross-section as a function of scattering angle θ plotted for various energies ($\alpha = 0.01, 0.5, 1.0, 5.0, 12.0$). The highest peak in the Figure corresponds to an energy of $\alpha = 12.0$, while the lowest peak corresponds to an $\alpha = 0.01$.

from two uniform distributions of random numbers. The obvious disadvantage of this method is that a portion of the random numbers generated are wasted. However, as mentioned earlier, this method eliminates the need to analytically solve Equation 1.8. The resulting randomly generated scattering angles are shown in Figure 1.12.

1.4 The Importance of Accuracy in Dose Calculation

As outlined in ICRU report 24, a goal within radiotherapy is to deliver a radiation dose to the patient within 5% of the prescribed dose. This 5% total error budget is to include all uncertainties associated within a radiotherapy treatment. In general, most of this 5% is taken up by machine calibration (up to 2%), patient set-up and movement, and imaging and diagnostic uncertainties. Given this, there is very little room left for uncertainties related to treatment planning dose calculation.

When considering the effect that radiation treatment uncertainties have on the outcome of a treatment, one can consider the quantities: *Tumour Control Probability* (TCP) and *Normal Tissue Complication Probability* (NTCP). As shown in Figure 1.13, TCP is the probability for local control³ of a tumour for a given dose delivered. A goal in radiotherapy is to achieve a TCP of 95%. However, in achieving this goal, the likelihood for healthy tissue complications increases. This is shown on the same graph as NTCP. The goal for NTCP is to aim for under 5%. As can be seen on this

³Cancer free at the site of treatment for at least 5 years.

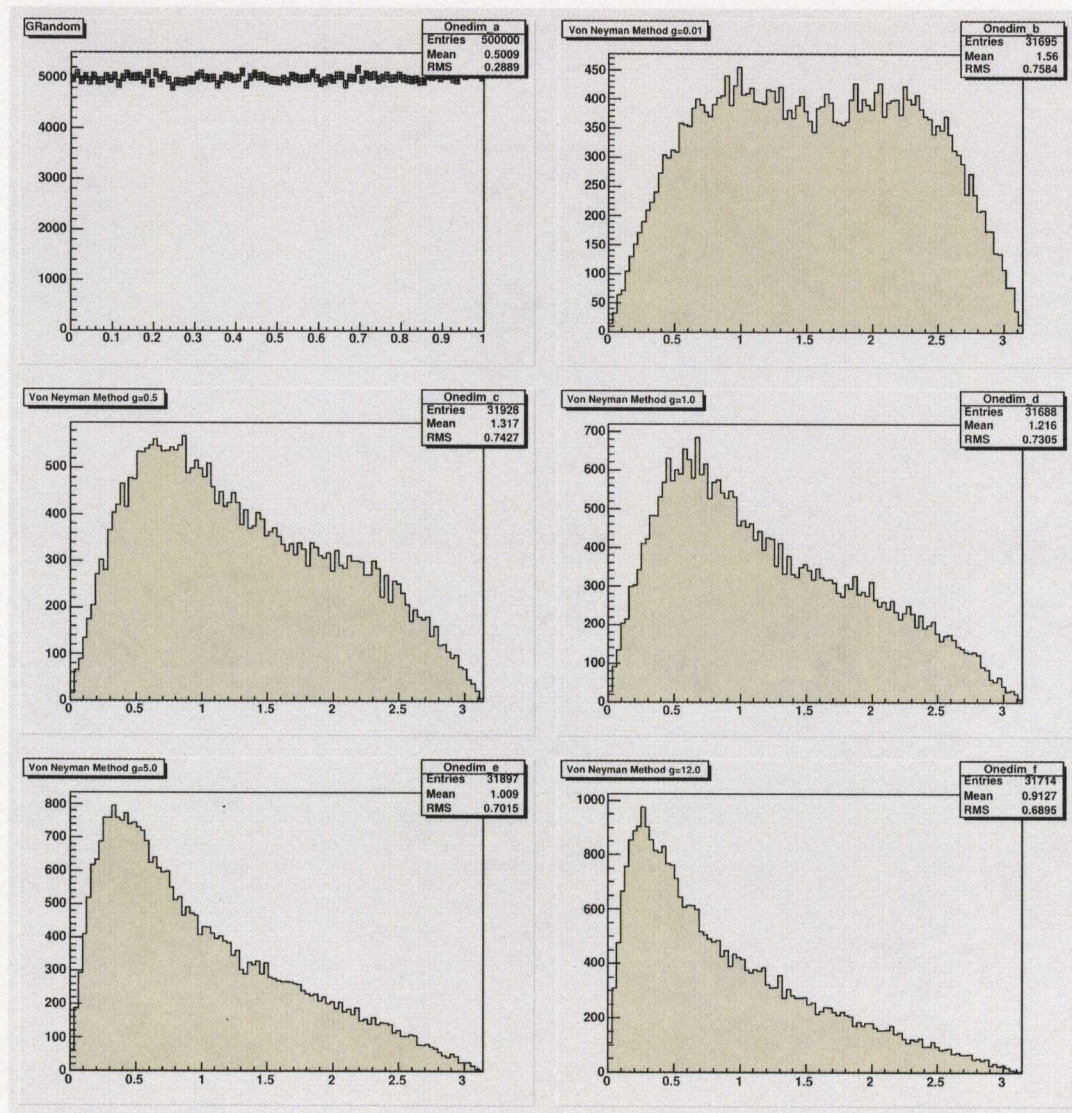


Figure 1.12: Monte Carlo histogram results for simulating photon scattering angles for $\alpha = 0.01, 0.5, 1.0, 5.0, 12.0$.

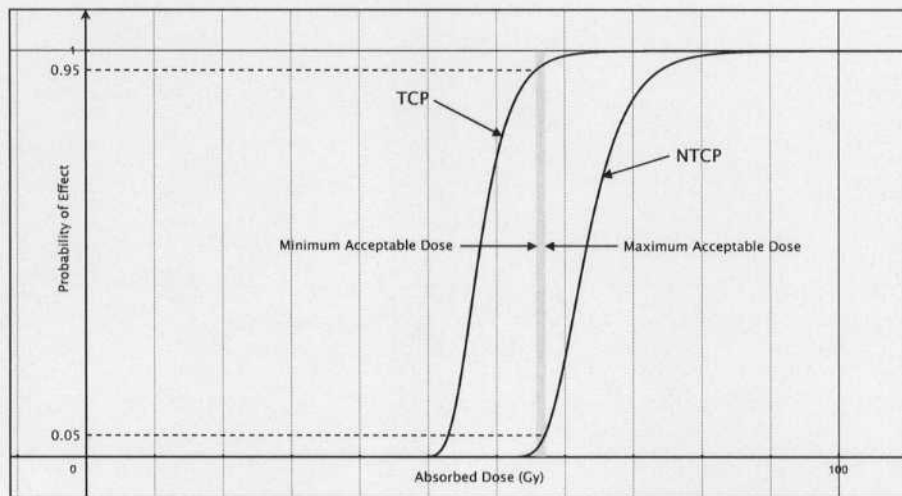


Figure 1.13: Graph showing the relationship of tumour control probability (TCP) and normal tissue complication probability (NTCP).

curve, these goals often do not leave much room for error; especially considering that small changes in dose result in large changes in TCP and NTCP when treating on the steep section of the curve. The sensitivity of these curves emphasizes the importance of minimizing treatment uncertainties.

1.5 The Scope of This Thesis

Before continuing on to further material in this thesis, a brief synopsis of the chapters to come will be presented. In the next chapter, the reader will be introduced to the basics of simulating radiation transport in radiation therapy. Here, an overview of the BEAMnrc Monte Carlo radiation simulation package will be presented and the reader will learn how the simulation of radiation therapy plans is achieved in the field

of Medical Physics. Chapter 3 will then discuss how BEAMnrc can be used to subdivide a given medical linac radiation field into beamlets, or beam elements for use in various radiotherapy applications. In Chapter 4, three distinct applications of these beamlets will be presented. In the first application, beamlets are used to fine-tune the output of the MC or “virtual” linear accelerator from which they are produced. In the second application, beamlets are used to derive corrections to the fluences of particular Intensity Modulated Radiation Therapy (IMRT) treatment plans that have been miscalculated by an inaccurate convolution-based dose calculation algorithm. Finally, in the third application, beamlets are incorporated into a Direct Aperture Optimization algorithm forming the basis for a purely MC based IMRT treatment plan optimization method. Results and a discussion of the results for these applications are then presented in Chapter 5 followed by conclusions to this body of work in Chapter 6.

Chapter 2

Simulation of Radiation Therapy Plans Using BEAMnrc

2.1 An Introduction to BEAMnrc

BEAMnrc [1] is a Monte Carlo simulation system for modeling radiotherapy sources. Built on the Stanford Linear Accelerator Center's EGS (Electron-Gamma-Shower) Code System [13], BEAMnrc was developed as part of the OMEGA project [14] with the intention of creating a 3-D treatment planning tool for use in radiation therapy. During this project the same structured set of particle interaction subroutines developed for the EGS system were used along with new subroutines to handle the geometry of a medical linear accelerator (HOWFAR and HOWNEAR) and particle scoring (AUSGAB). The result is a highly accurate means to model radiation transport from a photon source, through an accelerator head and into a CT reconstructed patient [14].

Until recently, the computing times required to simulate clinical treatment plans using Monte Carlo based methods such as BEAMnrc were not on the order



Figure 2.1: An intensity modulated radiotherapy plan designed to deliver a radiation field modulated to a portrait of Albert Einstein. The plan was scored on radiographic film (*left*), calculated in Eclipse (*middle*), and calculated using MC (*right*). The MC Einstein is, without a doubt a closer match to the film. [7]

of clinical acceptability¹. As such, faster, less accurate algorithms have predominantly been used to calculate dose distributions within a patient. Alarmingly, the 3-D pencil-beam convolution algorithm, a commonly used dose calculation algorithm, has recently been shown to miscalculate dose by as much as 15%, in the presence of tissue inhomogeneities [15] [3] [2] [4]. Even more common 2-D implementations have been shown to be much worse. Advanced analytic models incorporating tissue inhomogeneity corrections can improve calculations to a degree [16] [17] [18], but undoubtedly sacrifice a good deal of computation time. The Monte Carlo method is truly the gold standard in dose calculation with the capability of accurately simulating radiation dose deposition under almost all circumstances. In particular, Monte Carlo techniques can model backscatter from bone, surgically implanted prosthesis, scatter

¹Here the author might suggest that clinical acceptable calculation times might be on the order of <24 hrs. This is obviously quite arbitrary.

perturbations by air cavities, etc. much more accurately than any of the convolution algorithms [19] [20]. The accuracy of the Monte Carlo method does, however, come at a cost. In addition to the high computation times, the MC method relies on accurate geometrical modeling of the accelerator head and patient geometry. In order to accurately model an accelerator the exact specifications of all components must be known. Often manufacturers are reluctant to release this information freely as their designs are confidential. Even if the schematics are released, there is often still a good deal of “fine-tuning” involved in the manufacturing and commissioning of real medical linacs which cannot be accounted for in the virtual accelerator model using BEAMnrc alone.

2.2 The BEAMnrc Phase Space

One core element in BEAMnrc is its use of a phase space to record the output of the radiation beam. A phase space is simply a file containing the position, direction, energy and charge of all the particles crossing a specified plane in the accelerator model. Once a phase space has been written, it can then be used as future input for simulating the continuation of the beam further in the accelerator, further into the patient, or used to analyze the beam’s spectrum at a particular plane. Another variable stored in the phase space record, is each particles history tag and weight. The history tag, termed “LATCH” was originally intended to track the particular components of the accelerator with which a particle has interacted. However, for the

purpose of this research, the LATCH variable's use is reassigned, as will be explained in Chapter 3. The weight variable is used in the phase space to weight the dose delivered by the particular particle. The ability to weight particles is needed in variance reduction techniques incorporated in BEAMnrc as discussed later in this chapter. The weight variable is also made use of in a unique way for this research; not only in variance reduction techniques, but also to fine-tune the accelerator's output to match measured dose deposition profiles (see Chapter 3).

2.3 BEAMnrc Component Modules

Another core element of BEAM is that the accelerator model is constructed from a series of individual component modules, each of which is simulated independently from the other component modules. Many component module templates are included within the BEAMnrc distribution to assist in building an accelerator. Some examples include "SLABS" (electron target), "CONS3R" (primary collimator), "CHAMBER" (monitor chamber), "FLATFILT" (energy-flattening filter) and "JAWS" (secondary collimator). The specific dimensions, location, and material types of each module must be entered from manufacturer specifications. Each component module is designed to occupy a slab of space perpendicular to the beam's central axis. The thickness of this slab is defined by the dimensions of the module. Using this design, the accelerator becomes in essence "plug and play" in that one can insert or remove component modules to change an accelerator without need to worry about dependencies

with other modules.

2.4 The Clinac 21EX Accelerator Model

For this thesis a BEAMnrc modeled Clinac 21EX linear accelerator was used. The model was constructed using manufacturer specifications, provided in a confidentiality agreement with Varian Medical Systems (see Figure 2.2). For this model the accelerator head consists of a tungsten target, primary collimator, vacuum window, flattening filter, ion chamber, light mirror, and X and Y jaws. The Clinac 21EX is able to produce photon beams at peak energies of 6 MeV and 18 MeV (denoted MeVp), as well as five electron beam energies, with the only difference in accelerator head geometry being the thickness of the electron target and the shape of the flattening filter used. The 6 MeV target is 0.157 cm of Cu with 0.0889 cm of W brazed onto it, whereas the 18 MeV target is 1.016 cm of Cu with 0.0635 cm of W brazed onto it. Both targets are mounted on a carousel and the appropriate target is rotated into place when an energy is selected. Similarly, the flattening filters are also mounted on a carousel and rotated into place when an energy is selected.

For a given accelerator energy within the 21EX treatment head, many of the component modules remain unchanged. This fact, along with the ability to score phase spaces at any point in the accelerator head, allows the accelerator to be modeled using a time saving *two-stage* process. First, the electron target, primary collimator, flattening filter, ion chamber, and light mirror are all simulated and a phase space




Figure 2.2: BEAMnrc preview of 6 MV (left) and 18 MV (right) accelerator configurations and the materials used in each component module.

(phase space A) is scored in a plane just above the X and Y jaws. It is important to note that these components do not change from simulation to simulation. As such, these components need only be simulated once. Once created, phase space A then acts as the starting point for all further simulations of the same energy. Also recorded during this part of the simulation is $D_{\text{ch}}^{\text{forward}}$, the dose deposited in the monitor chamber. $D_{\text{ch}}^{\text{forward}}$ is needed to determine the relation of dose output to the number of monitor units scored in the monitor chamber. An example of an x-y scatter plot of phase space A is shown in Figure 2.3.




Figure 2.3: 6 MV phase space A scored at 27.3 cm from the photon source.

The second stage of the modeling process requires simulating particle transport through the X and Y jaws. The position of the jaws are, for the most part, different




Figure 2.4: 6 MV phase space B scored at 45.0 cm from the photon source.

for each treatment and thus must be re-simulated for each treatment. During this simulation, the mirror and monitor chamber are also included even though they are above the jaws. This is to allow an accurate collection of backscattered electrons in the monitor chamber and determination of $D_{\text{ch}}^{\text{back}}$, the backscattered contribution to the dose in the monitor chamber. $D_{\text{ch}}^{\text{back}}$ is added to $D_{\text{ch}}^{\text{forward}}$ to provide the total monitor chamber dose. Once completed, phase space B can be intensity-modulated (as described in the following section), if modeling an Intensity Modulated Radiation Therapy (IMRT) beam, or shaped (if modeling a static field radiation therapy beam), or just transported into a patient model (if simply modeling open fields). An example of an x-y scatter plot of phase space B through a set of $40 \times 40 \text{ cm}^2$ jaws is shown




Figure 2.5: Schematics of the Varian 21EX setup. The $BEAM_A$ simulation (left) involves the fixed component modules of the linac: target, primary collimator, flattening filter, monitor chamber and mirror. Its output is phase space A, scored above the jaws and a dose $D_{ch}^{forward}$ scored in the monitor chamber. This phase space is then used as a source for the $BEAM_B$ simulation (right), whose component modules are the jaws and, again, the mirror and the monitor chamber. The output of this run is phase space B, scored below the jaws and a dose D_{ch}^{back} scored in the monitor chamber. This phase space can then be intensity-modulated for IMRT simulations. [8]

in Figure 2.4. A schematic diagrams of the component modules used in the Clinac 21EX model is shown in Figure 2.5. An EGS generated 3-D rendered model of the Clinac 21EX is shown (with particle tracks) in Figure 2.6.

2.5 Modulating the Intensity of a Phase Space

Recent advances in medical accelerator head designs has allowed the use of a multi-leaf collimator (MLC) (see Figure 2.7) to modulate the intensity of a radiation field. This technique, known as Intensity Modulated Radiation Therapy (IMRT), can greatly re-




Figure 2.6: EGS Windows accelerator picture of the Clinac 21EX accelerator showing the position of phase space A and B and various history tracks.

duce the amount of radiation healthy tissue receives during a radiotherapy treatment. Figure 2.8 shows a sample comparison of the dose distributions received by a patient using a 3-field conformal radiotherapy treatment plan² and a 7-field IMRT plan.

Several methods for incorporating the intensity modulation of radiation therapy plans into Monte Carlo base simulations have been developed [21][22]. Less accurate methods simply generate a probability matrix from a given intensity map and eliminate particles from a given phase space stochastically based on this map. These methods unfortunately do not account for other important phenomena such as bremsstrahlung, beam hardening from transmission photons or scattering of photons from MLC leaves. For this thesis a method that does incorporate the above phenomena was used. The algorithm was developed at Virginia Commonwealth University by Siebers *et al* [23]. In this method, the amount of attenuating material a particle encounters while crossing the MLC is first determined from a detailed mapping of the MLC leaves. The schematics of the MLC's are specified in an input file and can include leaf thicknesses as well as tongue and groove dimensions and rounded leaf tip dimensions. A probability for particle transmission is determined from the thickness of the attenuating material and the relative intensity of the particular field segment to which the particle belongs. The intensity map is derived from the MLC leaf-sequence file that is produced using a treatment planning system. In order to account for photon scattering during transmission through the leaves, a single Compton scattering

²Field shapes conform to perimeter of the planned treatment volume (PTV).

event per photon is simulated using a Monte Carlo sampling of the Klein-Nishina formula. This is a time saving feature that can be applied since the Tungsten leaf thickness (5 cm) is equal to the mean free path of 6 MeV photons in Tungsten. The transmitted particles are then scored in a new intensity-modulated phase space and can be simulated into a 3-D phantom.

2.6 Creating a Monte Carlo Phantom

A Monte Carlo phantom is constructed using a 3-D matrix of voxels (volume elements) for which each voxel contains a physical density. There are two ways by which this matrix is most commonly created. The first way is to create the phantom “by hand” by first defining a set of x, y, and z boundaries and then assigning densities to these voxels. The second, more practical, way for clinical applications is to create a MC phantom from a set of CT images taken of the patient. In this method a set of CT densities in Hounsfield Units³ (HU) from each image are interpolated onto a 3-D matrix of voxels and converted to equivalent physical densities. The conversion from CT densities to physical densities is achieved by interpolating between a list of known physical-to-optical density relations (see Figure 2.9). The matrix of physical densities is then written to a file. Also found in this file, is a brief list of known materials present in the patient model, and a map of individual voxels to material types. This map is approximately determined from the provided list of materials present in the

³ $HU = \frac{\mu - \mu_w}{\mu_w} \times 1000.$




Figure 2.7: Close-up view of the 120 leaf Multi Leaf Collimator on the 21EX conforming to an arbitrary tumour volume. [Varian Medical Systems]




Figure 2.8: A standard conformal 3-field treatment of a head and neck tumour (*top left*) and the resulting dose distribution (*bottom left*). A 7-field IMRT plan for the same patient (*top right*) and resulting dose distribution (*bottom right*). The IMRT plan shows much greater delivery control which, in this case results in a much lower dose to the parotid (salivary) gland. The parotid gland is very radiosensitive and can cease to produce saliva when irradiated, a very uncomfortable side effect.

patient model and each voxel's physical density. At the time of simulation the atomic properties (*eg.* electron densities) of each voxel can then be determined.

Typically, CT images contain either 512×512 pixels if stored in DICOM⁴ image format, or 256×256 pixels if stored in CadPlan's CART image format. The pixel spacing is determined by the size of the area that was imaged. The number of CT slices is generally determined by the size of the prescribed region to be imaged and the chosen slice spacing. The BEAMnrc software package includes CTCreat, a tool to create MC phantoms, as part of the distribution.

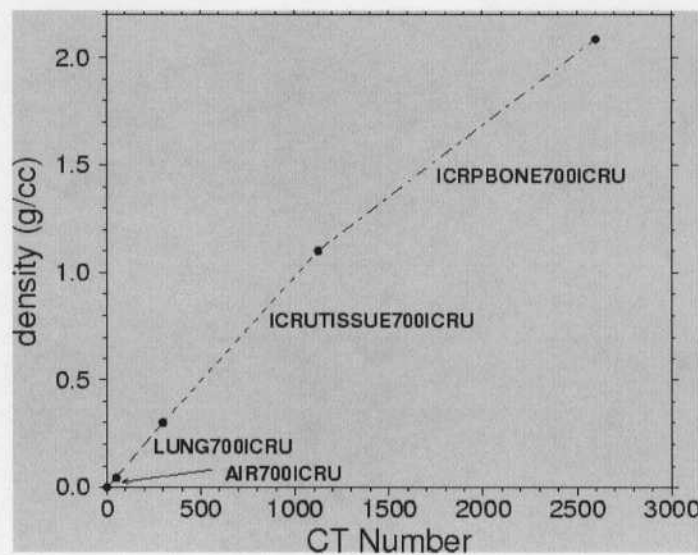


Figure 2.9: The default ramp for converting CT values to material and density in CTCreat [9]. (With permission from the National Research Council of Canada)

Often, raw CT images contain structures (other than the patient) that, if left in the image, can lead to incorrectly modeling the phantom dose distributions. The

⁴Digital Image and Communications in Medicine standard format.

most common source of these structures are the CT couch, which is not made of the same material as the treatment couch (The treatment couch is Carbon Fiber whereas the CT couch is plastic) and head support. To eliminate the structures from the CT images, a program was written to set all pixels outside of a defined contour to the Hounsfield value. A phantom can then be created with the new modified CT slices. This is shown in the left and middle images of Figure 2.10.




Figure 2.10: An axial CT image (with the Y axis flipped) containing traces of the CT couch and head support (*left*). CT image after all material outside of the body contour has been set to the optical density of air (*middle*). Axial phantom slice after cropping to the body contour (*right*).

2.7 DOSXYZnrc

Once BEAMnrc has simulated particle transport in the accelerator head and the output has been intensity modulated (if using IMRT), the resulting phase space can be transported into a MC phantom using DOSXYZnrc (included in the BEAMnrc distribution). Attempts have been made to include DOSXYZnrc as a component

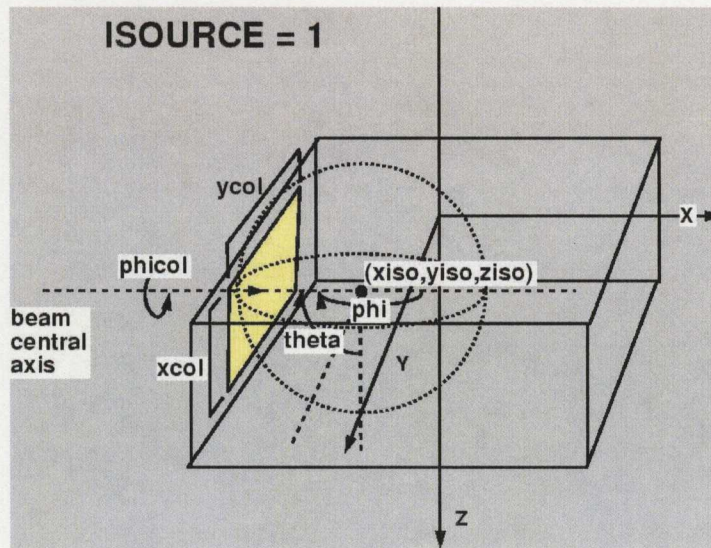


Figure 2.11: The spherical coordinate system used in DOSXYZnrc. The isocenter, $(x_{iso}, y_{iso}, z_{iso})$ defines the origin of the coordinate system. The beam collimator's position is then defined using the angles, θ and ϕ . In this figure, with the beam incident on the centre of the negative Y face of the phantom, θ is set to 90° and ϕ is set to 180° . x_{col} and y_{col} define the x and y dimensions of the secondary collimator (the jaws). ϕ_{col} defines the rotation of this collimator [9]. (With permission from the National Research Council of Canada)

module in BEAMnrc; however, currently DOSXYZnrc is a separate program [24].

In DOSXYZnrc a spherical coordinate system is defined on the phantom with an origin set to the the isocenter, as defined by $(x_{iso}, y_{iso}, z_{iso})$ the x, y, and z distances from the defined $(0, 0, 0)$ voxel. The incident beam angle is specified by θ , ϕ , and ϕ_{col} (see Figure 2.11). Clinically, the incident beam angle is defined on the Clinac 21EX (and most other linacs) by specifying a gantry angle, couch rotation angle, and collimator rotation. When using Varian treatment planning software, a coordinate transformation is needed to initially align the axes. This is due to the fact that axial

patient CT image sets are defined as x-y slices with the positive z axis towards the patients head (superior). From here the treatment planning system defines either a right (Eclipse) or left (CadPlan) handed coordinate system with the positive x axis pointing towards the patient's left side and positive y axis pointing towards the patient's back (Eclipse) or front (CadPlan).

Using Figure 2.11 and the right handed coordinates defined as in Varian's Eclipse software, a beam incident in DOSXYZnrc with $\theta = 0$, $\phi = 0$, and $\phi_{col} = 0$ is entering the patient from the -z direction which corresponds to a beam entering through the top of the patients head. In Eclipse this beam would be incident on the front (anterior) side of the patient. One must therefore transform ϕ , the couch rotation angle, by 270° and θ , the gantry angle by 90° . ϕ_{col} must then also be transformed by 180° . The transformation from right-handed Varian treatment planning system coordinates is⁵:

$$\begin{aligned}\theta &= \cos^{-1}(\sin(\gamma) \sin(\rho)) \\ \phi &= \tan^{-1}\left(\frac{-\cos(\gamma)}{\sin(\gamma) \cos(\rho)}\right) \\ \phi_{col} &= \pm \cos^{-1}[\cos(\omega) \cos(\gamma) \cos(\rho) \sin(\phi) - \sin(\omega) \sin(\rho) \sin(\theta) - \cos(\omega) \sin(\gamma) \cos(\phi)]\end{aligned}$$

Where γ is the gantry angle, ρ is the couch rotation, and ω is the collimator rotation in Varian Eclipse coordinates. The sign of ϕ_{col} is determined by taking the sign of:

⁵Thebaut, J. and Locke, C. *Coordinate transformations between Eclipse and DOSXYZnrc for running Monte Carlo simulations* (unpublished).

$$\cos(\omega) \sin(\rho) \sin(\phi) + \sin(\omega) [\cos(\gamma) \cos(\rho) \sin(\phi) - \sin(\gamma) \cos(\phi)]$$

Once the incident beam angle and isocenter position have been set in the DOSXYZnrc input file the phase space can then be transported into the phantom particle by particle until the number of histories requested is met. The resulting dose (and accompanying uncertainty) is scored in a 3-D array of voxels with size equal to the defined CT phantom voxels. The user can omit specified voxels from the scoring array. The doses, uncertainties, and voxel boundaries are all written in text format to a "3ddose" file. In order to make use of the tools available in the Eclipse commercial treatment planning system and to compare MC versus other dose calculation algorithms, the doses can be written into CadPlan or DICOM format. Figure 2.12 shows a comparison of MC and Pencil Beam (PB) calculated dose distributions on the first ever Monte Carlo IMRT verification performed in Canada [25].

Another modification to the CT phantom was made for this thesis. At the time of acquiring the CT images the size of the field of view is set. Often, clinically, to avoid cutting off critical areas of the patient, a larger field of view is used than is necessary. When creating a phantom, this extra image size is filled with voxels. When running DOSXYZnrc these extra unneeded voxels waste computation time. For this thesis a program was written to crop CT phantoms down to the precise size of the patient's body contour. As recorded in the results chapter, the total number of voxels

in a phantom can be reduced by 30% - 40% and simulation times are significantly reduced.

2.8 Variance Reduction Techniques in BEAMnrc

Several variance reduction techniques have been implemented into the BEAMnrc software package to speed up simulations and reduce uncertainties. Included techniques are Range Rejection and Bremsstrahlung Splitting.

Range-Rejection is incorporated in BEAMnrc simulations to avoid wasting computation time tracking low energy electrons. Once an electron's energy drops to a user specified energy (ECUT), the mean free path of the electron is evaluated. If this distance is less than the distance to escape the current scoring region the electron is no longer tracked and its energy is deposited in this region.

"Bremsstrahlung Splitting" is used to increase the number of bremsstrahlung photons created in the accelerator target per incident electron. During a bremsstrahlung interaction, N resulting photons are sampled individually and their weights are reduced by a factor of $1/N$ to conserve the total energy deposited in the phantom. In addition to improving the statistical uncertainty, Bremsstrahlung Splitting improves the efficiency of the simulation by reducing the amount of time spent tracking electron histories.

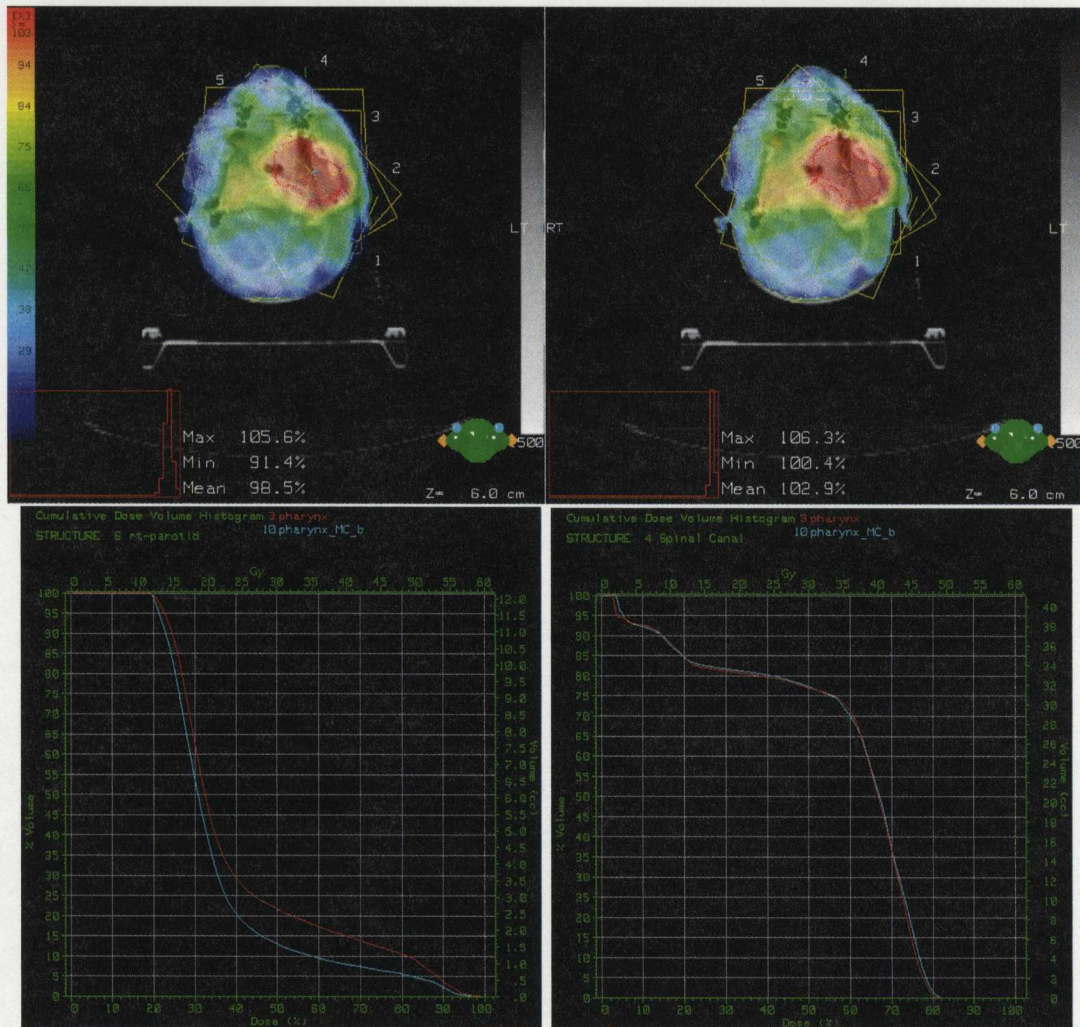


Figure 2.12: Monte Carlo dose distribution for a clinical treatment plan overwritten into the CadPlan dose distribution for comparison with CadPlan (*top left*) (as presented by Popescu *et al* at the BC Cancer Annual Cancer Conference announcing the first MC simulation of a clinical IMRT plan in Canada [25]). Original CadPlan calculated dose distribution (*top right*). DVH comparison of Monte Carlo and CadPlan dose distributions within the right parotid (*bottom left*) and spinal canal (*bottom right*).

Chapter 3

The Generation of Monte Carlo Beamlets

In this chapter the method for generating Monte Carlo beamlets will be presented in detail. The associated uncertainties and method for verifying correct beamlet dose distributions will also be covered.

3.1 Assigning Beamlet Numbers

As discussed in Chapter 2, a BEAMnrc phase space is a collection of traits scored for each particle crossing a defined Z plane. Depending on the scoring mode chosen in BEAMnrc, either 7 or 8, 32-bit variables are stored per particle. They are (in this order) LATCH, energy (MeV), X-position (cm), Y-position (cm), X-direction cosine, Y-direction cosine, particle weight, and Z-position of the last particle interaction (optional). The Z-direction cosine can be found using the relationship $\cos_x^2 + \cos_y^2 + \cos_z^2 = 1$. To avoid the need for storing another variable, the sign of the Z-direction cosine is stored as the sign of the weight variable. As mentioned earlier, the LATCH

variable is used to track the history of each particle. In the BEAMnrc input file there is an opportunity to define mappings from geometric regions of component modules to specific bits in the LATCH variable. Bits 1 to 23 of the 32-bit variable are allotted for use as a bit table for these mappings. As an example, one could map the geometric region defined by the primary collimator to bit 3 in the LATCH or map the geometric region defined by a particular layer in the flattening filter to bit 4 in the LATCH. Of note, bit 23 is permanently mapped to air regions within the accelerator. Depending on the LATCH option chosen, if a particle passes through or interacts in a geometric region with an assigned mapping the corresponding LATCH bit is set (set = 1, unset = 0). In this way one can determine from the resulting phase space which particles interacted in different regions in the accelerator.

Bits 24 to 28 in the LATCH variable are used to record the geometric region in which a secondary particle is created. For example, if a secondary particle was created in region 22, bits 24 to 28 would read 10110 for binary 22. Bits 29 and 30 store the charge of the particle and whether or not the particle is a contaminant particle. Bit 31 is used to record particles which cross the scoring plane more than once. Finally, bit 0 is used to record whether a bremsstrahlung or positron annihilation event occurs in the particle's history.

For this thesis the LATCH variable is used in a different and innovative way: to map small beam elements (beamlets) to binary beamlet numbers. The LATCH could not have been used in the usual sense because in its original design it could only

accommodate a maximum of 23 geometric regions. This project required many more regions to accommodate the large number of beamlets that each beam is divided into.

The method of generating beamlets for this thesis begins with a phase space that has already been simulated through an accelerator model. Either phase space A or B can be divided into beamlets. The choice depends on the intended use. For inverse treatment planning phase space B is used since this is the same plane as the multi-leaf collimator that will end up modulating the beam. For other applications such as fine-tuning an accelerator, phase space A is used since this is a permanent phase space input for future simulations. The chosen phase space is then divided up by placing an invisible grid of overall dimensions $\text{gridX} \times \text{gridY}$ with divisions $\text{divX} \times \text{divY}$ on it (see Figure 3.1). The grid dimensions are determined by counting out from the center of the phase space in multiples of specified divX and divY divisions until the specified outer boundary is reached. Partial beamlets on the outer border of the grid are not used. This is done to avoid beamlets with potentially low statistics. For inverse treatment planning, the outer boundary is specified by the X and Y jaws of the secondary collimator as can be seen in Figures 3.2 and 3.3. The beamlets are labeled from 1 to N as in Figure 3.1. The phase space is then read in particle by particle and the beamlet it is a member of is determined by the x and y positions in the phase space. All particles lying outside the specified grid are included in a single beamlet (labeled N+1). The bit mappings section of the LATCH, bits 1 to 21, is then assigned to correspond to the binary equivalent of the particle's beamlet number. This

differs from the original intention of the mappings section of the LATCH, but does not affect the ability to run BEAMnrc or DOSXYZnrc.




Figure 3.1: The grid used to assign beamlet numbers.

To assign the particle's beamlet number to the LATCH variable and preserve the information in the adjacent bits, a bit mask is used. Using the binary 'and' operator and a bit mask 11111111110000000000000000000001 the bit mappings for bits 1 to 21 are cleared. This erases all previous history of the particle. The binary beamlet number is then written into this region of the LATCH. It should be noted

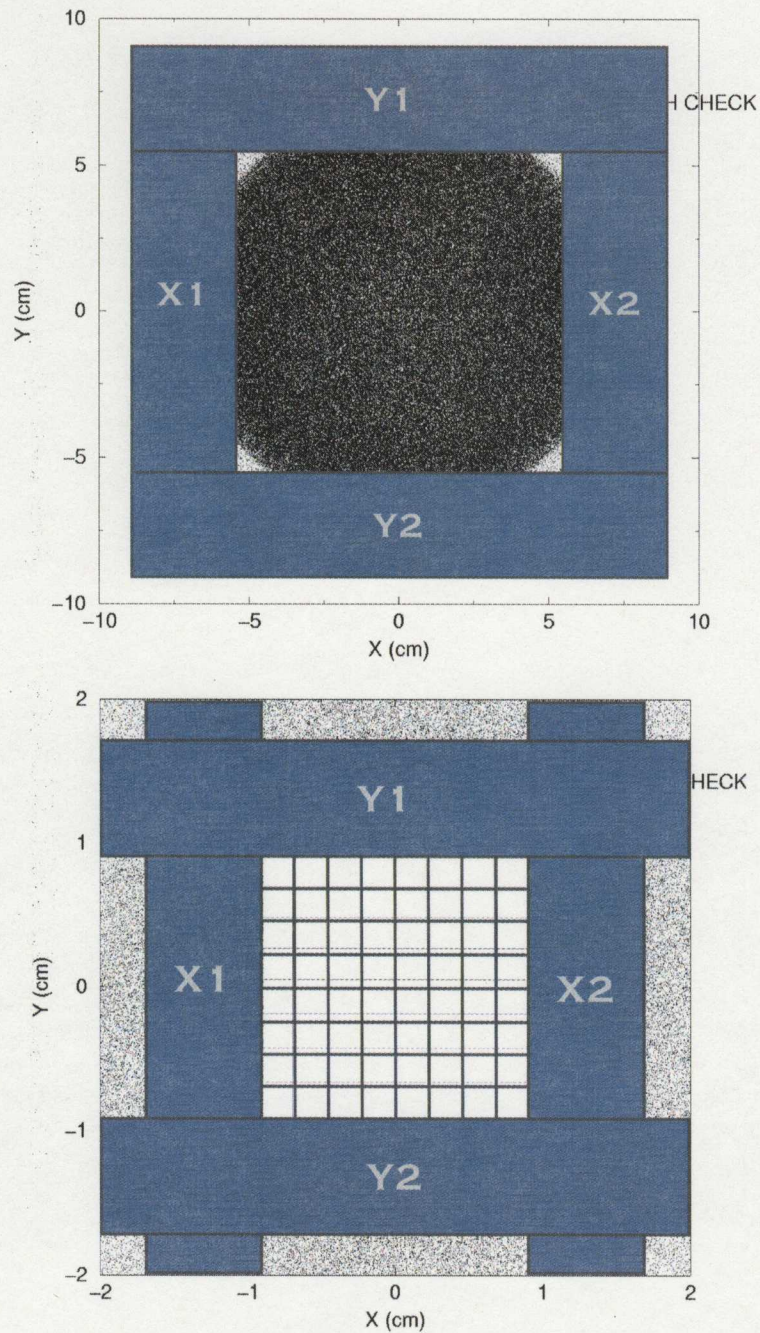


Figure 3.2: For the application of fine tuning the accelerator output X and Y jaws are superimposed on phase space A, thus determining the outer boundaries of the beamlet grid (*top*). Beamlet divisions were then specified based on the desired correction “resolution” (*bottom*).

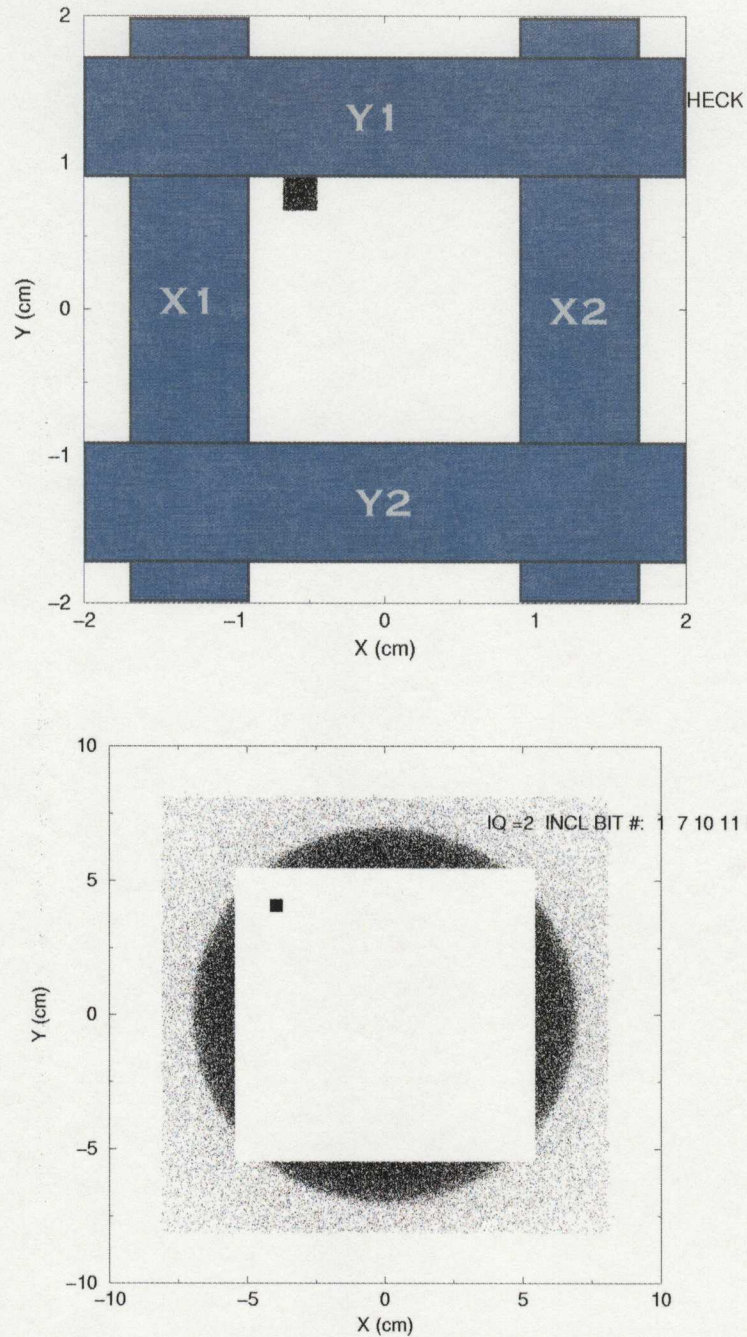


Figure 3.3: A single beamlet within the defined field (*top*). Beamlets 1601 (particles outside the grid boundary) and a random beamlet (1006) from a phase space of 1601 beamlets turned on using the modified inclusive/exclusive LATCH filter (*bottom*).

that although there are actually 23 bits reserved for bit mappings, only the first 21 are used. This was implemented after discovering that bit 23 is mapped to any air region within the accelerator model and cannot be “re-mapped”. Bit 22 is also not used in the beamlet mapping to maintain the ability to use the bit mappings in the originally intended way, for at least 1 geometric region. Extra bits can be left uncleared by the user with the sacrifice of lessening the maximum number of beamlets. This is not of great worry since $2^{21} = 2097152$ beamlets is likely more than enough for most applications. For this thesis, the beamlet divisions that were used varied between 0.25 cm and 1.0 cm with overall grid sizes of up to 40×40 cm. The total number of beamlets varied between 65 and 1601.

3.2 Running DOSXYZnrc With Beamlets

Two different approaches to running beamlets into a phantom were used for this thesis. The first method uses a modified version of the DOSXYZnrc particle LATCH bit filter to select particles of common LATCH bit setting (beamlet) from a phase space file and simulating them into a phantom. The second method uses a sorting algorithm to pre-sort particles of identical LATCH value. These particles then get written to smaller individual phase space files before running each individual phase space into the phantom.

3.2.1 DOSXYZnrc Modifications For Simulating Beamlets Using LATCH Bit Filtering

In order to run DOSXYZnrc in beamlets, minor modifications to the DOSXYZnrc source code were necessary. These modifications have been included in appendix A. The first modification made was needed to enable running beamlets into the phantom using the DOSXYZnrc inclusive/exclusive bit filter. DOSXYZnrc includes inclusive, exclusive and inclusive/exclusive bit filters to enable the ability to include and/or exclude simulating particles with certain LATCH bit settings in the simulation. These filters were specifically designed for use with the original intention of the LATCH. Small modifications were needed in order to allow filtering of binary numbers (up to 2^{21}) instead of just individual set bits (see appendix A for actual modifications).

Another modification of DOSXYZnrc was made to alter DOSXYZnrc's emphasis on delivering the exact number of histories requested by the user. In the DOSXYZnrc input file, a specific number of histories is requested (nhist). In order for exactly the requested number of histories to be delivered into the phantom, a phase space is sometimes not entirely run and sometimes it is restarted and partially rerun. When generating beamlets, because of the low number of particles in each beamlet, not running or rerunning even a small number of particles in each beamlet has a potentially dramatic affect on the acquired dose distribution. For this reason the DOSXYZnrc source code was made to force exactly *all* the particles in the phase space to be run. By making this modification, the number of particles actually simu-

lated is governed by the number of requested particle recyclings (since only multiples of the number of particles in the phase space can be run). In the end this modification does not significantly restrict particle history requests since the number of particles in each beamlet is small $\sim 1 \times 10^4$.

In DOSXYZnrc, dose output from a simulation is not given in Gy but instead in Gy per incident particle; that is, normalized by the calculated number of particles incident on the electron target required to produce the number of particles simulated into the phantom. The normalization factor is calculated by:

$$NP = NINCSRC \cdot \left[\frac{NCASE + ncaseold + nsmis + (NRCYCL + 1) \cdot (nsrjct + nsoutside)}{nshist} \right] \quad (3.1)$$

The quantity:

$$NCASE + ncaseold + nsmis + (NRCYCL + 1) \cdot (nsrjct + nsoutside)$$

is an estimate of the total number of particles read from the phase space file in the simulation. nshist is the number of particles in the phase space. The fraction of particles used out of the phase space is therefore:

$$NCASE + ncaseold + nsmis + (NRCYCL + 1) \cdot (nsrjct + nsoutside)/nshist$$

When multiplied by NINCSRC, the number of particles incident on the target to produce the phase space, the total number of particles incident on the target to produce the requested number of histories is found. In this formula NCASE is the number

of histories requested by the user, *ncaseold* is the number of histories simulated in any previous DOSXYZnrc runs (only applies to restarted runs), *nsmiss* is the number of histories that miss the phantom completely, *NRCYCL* is the number of times the phase space will be recycled¹. (*nsrjct* + *nsoutside*) is the total number of particles rejected because they had the wrong charge, LATCH bit setting, were going backwards, had crossed the phase space plane more than once, or were beyond the user-selected field size. When using inclusive/exclusive bit filters in DOSXYZnrc, NP will, in general, vary for each beamlet. This is especially due to a changing *nsrjct* for each beamlet run. This results in an inconsistent normalization of each beamlet's dose contribution which, in turn, leads to a problem when adding all the beamlet dose distributions. To solve this problem modifications were made to DOSXYZnrc's NP normalization factor formula. The modification was based on the assumption that all the particles in the phase space are forced to be simulated, as discussed above. By forcing exactly all the particles in the phase space to be run, there is no longer a need to calculate the fraction of particles run. Equation 3.1 then simplifies to:

$$NP = NINCSRC \cdot (NRCYCL + 1) \quad (3.2)$$

the number of particles required to create the incident phase space multiplied by the requested number of times each particle is recycled.

¹When recycling a phase space, the particles are reused with different random number seeds.

3.2.2 **DOSXYZnrc Modifications For Simulating Beamlets Using Phase Space Separation**

In order to simulate beamlets using the phase space separation method, no modifications were necessary to the inclusive/exclusive bit filter since it is not used in this method. All particles of a given LATCH bit setting were written into a single phase space and so the only modification needed to the source code was one to force all particles in the phase space to be run. The modification is identical to that described in the previous section. The associated modification to the normalization factor NP was also made.

When creating the smaller single beamlet phase spaces for this method a slight modification to the phase space file header was made. The phase space file header contains the file mode, the number of particles in the phase space (NPPHSP), the number of photons in the phase space (NPHOTPHSP), the maximum kinetic energy of particles in the file (EKMAXPHSP), the minimum kinetic energy of particles in the file (EKMINPHSP) and the number of particles incident from the original source (NINCSRC). When creating the individual beamlet phase spaces NPPHSP, NPHOTPHSP, EKMAXPHSP and EKMINPHSP are all written to reflect the properties of the new phase space. NINCSRC is not. Instead, NINCSRC is left to the value in the original phase space. This is done to ensure that all beamlets are normalized by the same NP value. The single beamlet phase space can then be thought of as a small scoring region of size $\text{divX} \times \text{divY}$ recording the information of particles

crossing through it from the simulation of the entire field. The simulation requires NINCSRC particles from the incident source to generate the entire field.

When combining the N beamlet dose contributions after simulation the doses to each voxel i can then simply be added:

$$\frac{D_{total}^i}{NP} = \sum_{j=1}^N \frac{D_j^i}{NP} \quad (3.3)$$

An example of several resulting beamlet dose distributions in an arbitrary patient model is shown in Figure 3.4.

3.2.3 Accounting For Uncertainties In DOSXYZnrc

When using the individual beamlet phase space method another important consideration must be taken into account to ensure correct beamlet dose uncertainty calculation. Recently, DOSXYZnrc has adopted a new, more accurate method for calculating dose uncertainties. This method is known as the “history-by-history” method [26]. The old method of determining dose uncertainties in each scoring voxel was calculated by:

$$\sigma_{\bar{D}} = \sqrt{\frac{\sum_{i=1}^N (D_i - \bar{D})^2}{N(N-1)}} \quad (3.4)$$

where N is the number of statistical batches for which the simulation is broken up into (usually 10), D_i is the dose value in batch i , and \bar{D} is the mean dose value evaluated over all batches. This method was replaced for several reasons. One reason was that by not using a large number of batches, significant fluctuations in the uncertainties arise (due to relatively small number of batches N). Another reason is




Figure 3.4: Monte Carlo dose distributions scored for various beamlets within an artificial lung interface phantom.

that by grouping histories into batches, any correlations between incident particles are ignored. To explain this further, suppose a particle incident on the linac target produces many bremsstrahlung photons. The properties recorded in the phase space (energy, position, momentum, etc.) of this set of photons are certainly correlated, especially when using variance reduction techniques such as bremsstrahlung splitting. Using the batch method of calculating uncertainties no knowledge of particle correlations is known and each particle in the phase space to be an independent history. Dose is simply scored in each voxel for each batch and the variance of each voxel is calculated over all batches using Equation 3.4.

The above method was recently replaced by a “history-by-history” method which was implemented to resolve the problems mentioned in the previous paragraph. For the history-by-history method Equation 3.4 is rewritten as:

$$\sigma_{\bar{D}} = \sqrt{\frac{1}{N-1} \left(\frac{\sum_{i=1}^N D_i^2}{N} - \left(\frac{\sum_{i=1}^N D_i}{N} \right)^2 \right)} \quad (3.5)$$

The quantities $\sum_{i=1}^N D_i^2$ and $\sum_{i=1}^N D_i$ are then calculated “on the fly”, after each history². The quantity N in Equation 3.5 now refers to the number of independent, or primary histories (i.e. incident electrons on the target) not the number of batches. In order to keep track of N , a modification was necessary to the way in which BEAMnrc phase spaces are written. Using the history-by-history method, the first particle from each primary history, that reaches the phase space scoring plane is scored in the phase

²Here the term “history” refers to all particle tracks associated with one initial particle.

space file with a negative energy³. Other particles that are created by this primary particle during the simulation and reach the scoring plane, are scored in the phase space file with a positive energy. In this way groups of correlated particles can be determined in the phase space (all particles scored from one negative energy to the next) and the actual number of primary particles that reach the phantom can be determined. Because there are generally a large number of primary histories generated (N), the error on the dose uncertainty is reduced compared to the previous batch method.

When using the history by history uncertainty calculation method when generating Monte Carlo beamlets, several considerations must be taken into account. First, all secondary particles created from a given primary history are not likely to be scored in the same geometrical beamlet area as the primary history that created them. Because of this, when creating individual beamlet phase spaces there will likely be “stray” secondary particles (with positive energy) that were created by a primary history that is recorded (with a negative energy) in a different phase space. When using a phase space file as input in DOSXYZnrc, dose deposited by particles from one negative energy to the next are grouped in the calculation of the uncertainty. Unfortunately the process of sorting a phase space file into individual beamlet phase space files “jumbles up” the grouping of independent primary histories.

Built in to the DOSXYZnrc source code is a check to determine if the phase

³This does not imply that the particle is actually of negative energy, it is simply a way of avoiding the use of another scoring variable

space uses negative energy primary history tracking. This check simply reads the energy of the first particle in the phase space and if it is negative then the program assumes that the negative energy method was used to generate the phase space. If the first particle's energy is positive, DOSXYZnrc assumes that the phase space was created without the negative energy method and proceeds to treat every particle in the phase space as an independent history. The developer of this method has shown that the effect of particle correlations on dose uncertainties is minimal [26]. As such, the Monte Carlo beamlet dose distributions generated for this thesis treat each particle as an independent history. The effect that these alterations have on the overall dose uncertainties remains to be investigated and is beyond the scope of this thesis. Future work is needed to allow better managing of correlated particle quantities using the history by history method.

3.3 Collecting Beamlets

With the appropriate modifications in place, beamlets can be run one by one into the phantom by either specifying the bits to include (the binary beamlet number) and exclude (everything else) or by queuing the appropriate beamlet phase space. Each beamlet run requires its own DOSXYZnrc input file. In this file the number of histories requested for each beamlet is determined by multiplying the number of particles in each beamlet by the number of times each particle is to be recycled. For the work in this thesis the DOSXYZnrc input files are generated by calling a C

function from a UNIX shell script.

The result of each beamlet run is a 3ddose file (an ascii file containing the doses and uncertainties to each defined voxel in the phantom). In order to reduce the number of files produced when generating a series of beamlets, the doses and uncertainties are extracted from this file and are appended into single binary dose and error files. The 3ddose file is then deleted. The process of generating appropriate input files and running DOSXYZnrc into the phantom is streamlined by using C-Shell UNIX shell scripting.

3.4 Converting Monte Carlo Dose to Absolute Dose

The doses resulting from a DOSXYZnrc simulation are scored in units of Gy per incident particle on the electron target. Multiplying by the number of incident particles on the target would give a very small dose. This is because the number of incident particles simulated during a typical MC simulation ($\sim 2.5 \times 10^8$) is, in fact, miniscule in comparison with an actual treatment ($\sim 1.0 \times 10^{17}$). Simulating the actual number of particles incident on the target would require on the order of 2.4×10^9 hours, or roughly 274 thousand years! A simulation therefore just samples a fraction of the total treatment and the resulting dose per incident particle is converted to absolute dose. A problem exists, however, in that radiation treatments are not delivered in number of incident particles on the target, they are delivered in monitor units (MU). A monitor unit is a measure of the charge collected in the monitor chamber in the linac

head. 1 MU gives an absorbed dose of 1 cGy at the depth of maximum dose (d_{MAX}) for a 10×10 cm² field with a source-to-axis distance (SAD) of 100 cm. The dose per incident particle acquired from the DOSXYXnrc simulation must therefore be converted into absolute dose corresponding to a specific number of MU. The formula, derived by Popescu *et. al.* [8], is:

$$D_{xyz,abs} = D_{xyz} \frac{(D_{ch}^{forward} + D_{ch}^{back(10 \times 10)})}{(D_{ch}^{forward} + D_{ch}^{back})} \frac{D_{xyz,abs}^{cal}}{D_{xyz}^{cal}} \cdot MU \quad (3.6)$$

where D_{xyz} is the dose in a given voxel, $D_{ch}^{forward}$ is the forward dose to the monitor chamber, D_{ch}^{back} is the dose to the monitor chamber from backscatter off the jaw configuration, $D_{ch}^{back(10 \times 10)}$ is the backscattered dose to the monitor chamber measured for a 10×10 cm² field, $D_{xyz,abs}^{cal}$ is the absolute dose at the MU calibration point, and D_{xyz}^{cal} is the dose at the MU calibration point in the MC simulation.

3.5 Verification of Beamlet Dose Distributions

Once all beamlets have been run into the phantom and the dose distributions from each individual beamlet have been added together, the resulting dose distribution can be compared with a non-beamlet simulation of the same treatment plan to verify their correctness. The resulting summed beamlet dose distributions and non-beamlet run dose distributions are shown in Chapter 5, the results chapter of this thesis.

Chapter 4

Applications of Beamlets

In this chapter, the technique of adjusting beamlet weights with the aim of converging Monte Carlo beamlet dose distributions to prescribed dose distributions is used in three distinct applications: to fine tune the output of an accelerator model, to determine a correction factor for remodulating miscalculated treatment plans, and to perform inverse treatment planning with Monte Carlo beamlets. The fine tuning of accelerator models will be described first.

4.1 Fine Tuning the Output of an Accelerator

A major barrier to a wider clinical use of MC dose calculation is the difficulty in accurately characterizing the photon source. A rigorous and time consuming commissioning process is required to ensure that the MC virtual linac produces beams of essentially the same quality as those of the real treatment unit being modeled. Typically, the adjustable parameters (the energy and spatial distribution of the electron beam incident on the linac target) are chosen by trial and error such that the

MC simulations reproduce the depth dose and parallel plane profiles measured in water. Often, after systematically varying the input parameters over a suitably chosen range, it is found that the shape of the measured profiles cannot be exactly matched. To address this problem a method was developed which allows a user with an arbitrary linear accelerator to commission a BEAMnrc virtual accelerator to accurately reproduce the measured output of the real accelerator in water.

After accurately configuring the treatment head according to the manufacturer specifications, a phase space (phase space A) is recorded just above the secondary collimator. This phase space is then divided into beamlets using the phase space separation method described in Chapter 3. The new beamlet labeled phase space is then simulated through a set of 40×40 cm² jaws. During this simulation each particle's LATCH variables is inherited and carried through to the particles in phase space B, below the jaws. The resulting beamlets are then transported one at a time into an artificially created water tank phantom using DOSXYZnrc. The resulting 3D dose distributions for each beamlet are then used in a procedure, similar to that used by Webb [27], employing a simulated annealing algorithm with a quadratic cost function. The cost function also employs an inverse dependence on the MC statistical uncertainty such that voxels with greater statistical uncertainty do not contribute to the cost function as much as those with lower uncertainty. In this way, the cost function is isomorphic with the χ^2 statistic. The cost function used for this

application is:

$$Cost = \sum_{i=1}^n \frac{1}{\sigma_i^2} \cdot [D_i^{MC} - D_i^{measured}]^2 \quad (4.1)$$

where $D_i^{measured}$ is the measured dose to voxel i , n is the number of voxels/optimization points in the phantom and σ_i is the MC uncertainty of the dose deposited in voxel i . D_i^{MC} , the Monte Carlo calculated dose to voxel i , is calculated by:

$$D_i^{MC} = \sum_{k=1}^N [W_k \cdot D_i^k] \quad (4.2)$$

where N is the number of beamlets in the optimization, W_k is the weight of beamlet k , and D_i^k is the dose to voxel i from beamlet k .

For this application the measured dose distribution, $D_i^{measured}$, is the distribution resulting from a $40 \times 40 \text{ cm}^2$ field incident on a $50 \times 50 \times 50 \text{ cm}^3$ water tank. The goal of this application is to minimize the cost function 4.1 by adjusting the weights of each beamlet dose distribution, thereby converging the Monte Carlo dose distribution to the optimal (measured) dose distribution. The resulting beamlet weights can then be assigned to the particle weight variables, producing a commissioned phase space above the jaws to act as the source of all future simulations.

During one iteration of this method, a random number is generated to determine whether a positive or negative weight change should occur. The weight change amount or "grain" is then added or subtracted to a single beamlet's weight. All beamlet dose distributions are then recombined and the cost is calculated. This cost

is then compared to the cost of the last accepted weight change:

$$dV = Cost - Cost_{previous} \quad (4.3)$$

Weight changes that reduce the cost (positive dV) are accepted immediately. Weight changes resulting in an increase in the cost (negative dV) are also accepted but only a fraction of the time. This is to help prevent the Monte Carlo dose distribution from converging to a local minimum in the cost function and is the defining characteristic of the simulated annealing algorithm. The fraction of cost increasing weight changes is decreased with the number of iterations using a "cooling schedule":

$$p = e^{-\frac{dV}{k_b T}} \quad (4.4)$$

where $k_b T$ is decreased slower than $1/\ln(i)$. This is in accordance with findings to ensure a global minimum is reached [28].

During each iteration in this process the weight change, or "grain" size, is reduced using a linear "grain reduction schedule":

$$g = -\frac{i}{\alpha} + g_o \quad (4.5)$$

where g_o is the starting grain size, i is the number of full beamlet iterations and alpha is an adjustable parameter to control the rate of grain reduction. By using a grain reduction schedule, large initial differences in dose distributions are reduced more quickly than using a small, constant grain size. For the simulated annealing algorithm to work, both negative and positive grains must be used [29].

The simulated annealing algorithm premise is analogous to the way in which metals recrystallize in the process of annealing. In this process liquid metals, initially at a high temperature and disorder, are slowly cooled such that the system is approximately in thermodynamic equilibrium at all times (if the system is cooled too quickly it can end in an amorphous metastable state of local minimum potential, rather than a true crystalline state of global minimum potential). To achieve this, at each given temperature T and energy E , perturbations to the system leading to a change in energy dE are accepted with a probability given by the Boltzmann factor $\exp(-\frac{dE}{k_b T})$.

It should be noted that there is some question as to the existence of local minima in inverse radiation therapy planning. Recent studies have shown that in the presence of non-convex feasible parameter spaces few cases with local minima arise. Those that were observed were found to produce dose distributions very close to the global minimum [30] [31]. Despite these findings the algorithm used in this thesis remains as a simulated annealing algorithm. In the event that local minima are not suspected one can “cool” the optimization faster by reducing $k_B T$. Setting $k_B T$ to zero essentially reduces the algorithm to a least squares minimization.

One unique feature of the cost function used in this thesis compared with those used by Webb [27] and Jeraj and Keall [29] is the inclusion of statistical uncertainty. Since this uncertainty is calculated for each scoring voxel in the phantom during the DOSXYZnrc simulation, it is readily available to be used in the procedure.

The algorithm is run until the cost has been reduced to a user defined value or the total number of iterations performed reaches a preset limit. Following the optimization, the last determined weights for each beamlet are recorded. Phase space A is then read particle by particle and each particle's new weight is assigned based on the beamlet of which it is a member of. This weighted phase space then acts as the source for all future simulations.

Just as a monitor chamber must be recalibrated after modifying components of an accelerator, in order to correctly convert Monte Carlo dose distributions into absolute dose after weighting the phase space the absolute dose conversion constants must be recalculated. The method for determining the absolute dose conversion constants is outlined in Chapter 3.

The results of fine tuning the output of 6 MV and 18 MV accelerators are presented in Chapter 5.

4.2 Remodulation of Miscalculated Treatment Plans

Another application of Monte Carlo beamlets in this thesis is in the determination of modulation corrections to treatment plans that have been miscalculated using pencil beam convolution. As mentioned earlier, findings have shown the algorithm often miscalculates dose in areas such as the lungs where tissue (higher density) and lung (lower density) come together at a sharp interface. The algorithm's inaccuracy arises when convoluting scatter kernels near these interfaces. At each dose calculation point

within the phantom, the density is determined and the appropriate scatter is applied and added to the dose at the surrounding points. However, if the calculation point is in one medium but points that receive the scatter are in another medium (such as at a sharp density interface) the dose received from the scatter will be calculated as though all the points were in the same medium. In general this leads to an overestimate of the dose to the air cavity. (see figure 5.39).

As stated in Chapter 2, Monte Carlo based algorithms excel in accurately modeling tissue inhomogeneities. When simulating treatment plans in regions with sharp density interfaces, Monte Carlo calculated doses closely match measured doses while pencil beam calculated doses often appear quite different. For this application a beamlet weight map is derived when Monte Carlo dose distributions (and subsequently measured dose distributions) are converged to pencil beam convolution dose distributions by optimization of beamlet weights. Corrections to the pencil beam calculated beamlet fluence map can be made by scaling the existing fluence by the MC beamlet derived weight map. The purpose for this is simply because pencil beam convolution dose distributions are currently what are approved and expected to be delivered by oncologists. In fact what is received by the patient can be different by as much as 15% [2] [3] [4]. This problem is particularly of concern when treating sites near lung, sinus, throat and esophageal regions, since these are areas with a high amount of density inhomogeneity.

The method begins by first generating beamlets for an IMRT treatment plan

that was created using a treatment planning system. The beamlets are generated at the level of phase space B after any intensity modulations have been made. The size of each beamlet is made to correspond with the width and movement of the leaves used in the multi-leaf collimator (MLC). This is to later allow 1-to-1 correspondence between beamlet weights and leaf motion. The Clinac 21EX 120 leaf MLC used in this thesis contains 60 leaf pairs. 10 pairs of width 1.0 cm¹ followed by 40 pairs of width 0.5 cm followed by another 10 pairs of width 1.0 cm. The beamlets were made 0.5 cm along the length of travel, although in reality the beamlets could be any dimension along the direction of travel. This dimension was chosen for simplicity and because it seemed to be a relatively appropriate dimension given the task at hand.

Once all beamlets were simulated for the plan, the beamlet weights were optimized. For this, the simulated annealing algorithm described in Section 4.1 was used with the following modifications. First the ability to extract optimization dose points/voxels from structures defined within the patient was developed. This was crucial to both reduce the number optimization points and to provide a means to prioritize outcomes in the optimization. Second the ability to accommodate multiple fields was added. The cost function was then modified to allow for priority weighting of structures. The cost function used was:

$$Cost = \sum_{j=1}^S \sum_{i=1}^{n(j)} \frac{1}{\sigma_i^2} * \frac{P(j)}{n(j)} \left[D_i^{MC} - D_i^{prescribed} \right]^2 \quad (4.6)$$

where $n(j)$ is the number of voxels/optimization points in structure j , σ_i is the un-

¹All dimensions are as projected to 100 cm from the electron target's upper surface.

certainty of the dose deposited in voxel i . $P(j)$ is the priority setting of structure j . This priority setting can be used to scale the contribution that voxels from each structure have in the cost function. The Monte Carlo calculated dose to voxel i , D_i^{MC} , is calculated as above in Equation 4.2. By using $\frac{P(j)}{n(j)}$ instead of just $P(j)$ for priority weighting in the cost function, the priority placed on structures is scaled by the inverse of the number of points included in the optimization. Without this in place small structures with high priority can be easily outweighed by large structures with low priority setting in the cost calculation.

Once the Monte Carlo dose distribution has acceptably converged to the pencil beam convolution dose distribution, the final beamlet weights must be converted into equivalent leaf sequences. The leaf motion calculator used to accomplish this is part of the Varian treatment planning software. It operates by converting an ideal planar fluence into a physically achievable leaf sequence. To make use of Varian's leaf motion calculator, planar fluences, as calculated with the pencil beam convolution algorithm, are scaled by the Monte Carlo beamlet weights in a fluence file. This file is then input into the leaf motion calculator and a leaf sequence is determined. During this process the ideal fluence is unfortunately smeared out to a degree. This is not only because the ideal fluence is not always achievable, but also because leaf transmission effects such as tongue and groove transmission and rounded leaf tip transmission are accounted for during the calculation. As will be discussed in the next section, using a direct aperture optimization eliminates this smearing effect since only physically obtainable

leaf sequences are modeled in the optimization. The results of remodulating treatment plans can be seen in the results section of this thesis.

4.3 Inverse Treatment Planning Using Direct Aperture Optimization and MC Beamlets

Direct aperture optimization (DAO) is an IMRT inverse planning technique whereby the shapes and relative weights of the apertures in a plan are optimized directly [32]. Traditional IMRT planning differs from DAO in that pencil beam weights are optimized to create fluence maps. These fluence maps are eventually translated into a set of deliverable aperture shapes via a leaf sequencing algorithm. DAO has been shown to produce treatment plans comparable to traditional IMRT plans, but using significantly fewer leaf arrangements. Fewer leaf arrangements leads to less beam-on time and a reduction of patient treatment times. In traditional IMRT plans, the optimally determined fluences cannot always be physically delivered. This is often the case when, in order to deliver a specific fluence, the leaves are required to move faster than physically possible. Such things as interleaf leakage and leaf tip effects also add to a smearing of the ideal fluence [32]. In DAO, there is much less smearing of the ideal fluence map because, during the optimization, only the dose distributions from deliverable leaf arrangements are used and leaf transmission is taken into account at the time of optimization. Leaf tip effects and interleaf leakage, could also be incorporated into a DAO model. For this research these effects are not included.

Incorporating MC beamlet dose distributions in DAO is a novel concept. To accomplish this the MC beamlet generation method described in the last chapter was used. However, some rearranging of the output data was needed in order for it to be used in the specific DAO software used for this thesis. The output of the beamlet simulations consists of a series of doses to a defined set of points, for a specific beamlet in a particular field. The DAO optimization, prior to using MC beamlets, was using beamlet data calculated by Varian's CadPlan treatment planning software. The format of this data was given as the contribution of each beamlet, in a specific field, to a specific point in the phantom. Besides the ordering of data, the beamlet numbering method had to be changed. In MC, the beamlets were numbered from the X1 and Y1 jaw corner as in figure 3.1 oriented as if one were looking up through the jaws at the source. In CadPlan (or Eclipse), beamlets are ordered the same way, only as if one were looking down through the jaws at the patient. To fix this beamlets were reordered to match CadPlan convention.

At the start of a DAO optimization, the leaves for each field were arranged to the cross sectional view of the PTV plus a 1 cm buffer. This position establishes the outer limit for each leaf position. During each iteration in the optimization, one of the fields is chosen randomly. From this field, one of the segments is then chosen randomly (For each field there is a set number of segments, or leaf arrangements, usually ranging from 1 to 7). From this segment, either one of the beamlets in this field is chosen randomly and the weight of this beamlet is randomly increased or

decreased by a set grain size, or the entire “segment weight” is randomly increased or decreased by a set grain size. This is then repeated for subsequent iterations.

The cost function is similar to that used by Webb [27], and similar to that used in the above section without the dependency on uncertainty. Had the author written the DAO software, the uncertainty would have been included in the cost function.

From the MC beamlet data a random number of dose points are chosen within each structure. The average density of these points is 30 dose points per cm^3 . Structure priority weighting is also used to prioritize the optimization goals for the PTV and organs at risk (OAR).

It should be noted that the scope of this thesis component was to incorporate MC beamlets in a DAO optimization, not to measure the improvement of optimizations using MC vs. Pencil Beam calculated beamlets. This is a very interesting study and is currently in progress by Bergman *et al* at the University of British Columbia[10].

Chapter 5

Results and Discussion

Before proceeding with the presentation and discussion of results, the objectives of this thesis will again be stated. The first objective was to develop a method of fine-tuning the output of a Monte Carlo virtual accelerator to obtain better agreement with the real accelerators measured data. The second objective was to address the problem of dose miscalculation in Intensity Modulated Radiation Therapy (IMRT) planning through the introduction of Monte Carlo dose calculation. Both of these objectives were met through the use of MC calculated beamlet dose distributions. The results and a discussion of the methods of calculating beamlet dose distributions will, therefore, be discussed first. Following this, the results and a discussion of the method of fine-tuning the output of a Monte Carlo virtual accelerator will be presented. Finally, the results and a discussion of the applications addressing dose miscalculation will be presented.

5.1 Generation of Beamlets

As discussed in Chapter 3, Monte Carlo beamlet dose distributions were obtained in two ways: via LATCH bit filtering and via phase space separation. The beamlet results and the advantages and disadvantages of both methods will now be presented.

To ensure that the method of running a BEAMnrc simulation using beamlets is accomplished correctly, BEAMnrc simulations run beamlet by beamlet into a phantom were verified by running the unmodified non-beamlet BEAMnrc simulation and comparing the results. Figures 5.1 and 5.2 show profile comparisons of a 4×4 cm² field scored within an artificial lung-air interface phantom. The results show that the total dose distribution from running BEAMnrc in beamlets is identical, within uncertainty, to running the original BEAMnrc code. It should be noted that all data collected for this thesis was done using the BEAMnrc02 release of the BEAM Monte Carlo System, running Linux 2.4.19 with i686 architecture. Modifications made to this release are documented in the appendix.

The advantages and disadvantages of the two methods for obtaining beamlet dose distributions will now be discussed. The advantage of using LATCH bit filtering, where a bit filter is used to run only particles of a given beamlet each DOSXYZnrc run, is simplicity. This is because all particles are contained in a single phase space from which the particles of a given beamlet are filtered out for a given DOSXYZnrc beamlet run. The disadvantage of using LATCH bit filtering is that for each DOSXYZnrc




Figure 5.1: Parallel plane profiles of a 4×4 cm² field incident on an artificial lung interface (phantom shown in Figure 5.37). The solid line profile is scored from a standard BEAMnrc simulation. The dashed line profile is scored from a reassembled BEAMnrc simulation that has been run beamlet by beamlet into the phantom.

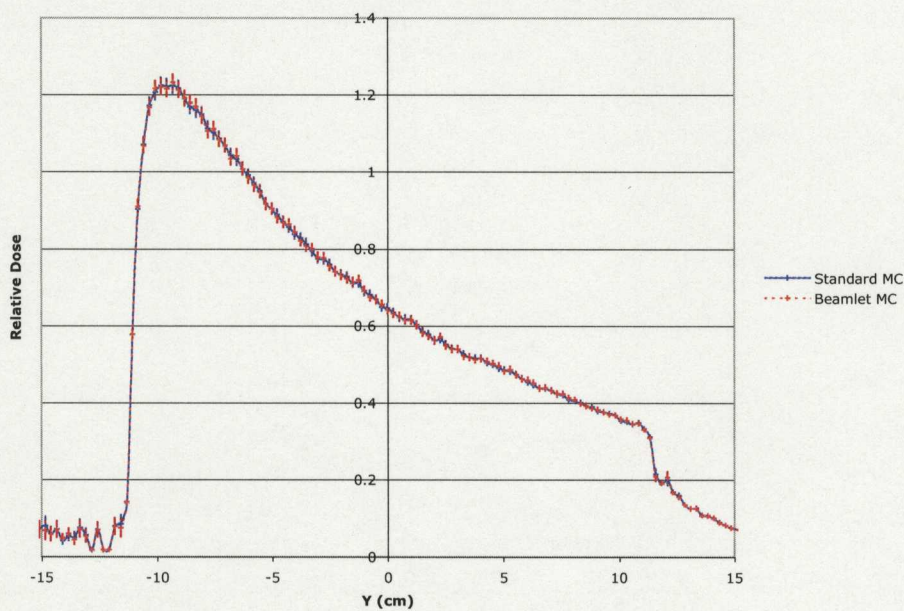


Figure 5.2: Comparison of beamlet and non-beamlet depth dose profiles of a 4×4 cm^2 field incident on an artificial lung interface (phantom shown in Figure 5.37). The solid line profile is scored from a standard BEAMnrc simulation. The dashed line profile is scored from a reassembled BEAMnrc simulation that has been run beamlet by beamlet into the phantom.

beamlet run the entire phase space must be read through to find all of the particles of a given beamlet. Depending on the size of the phase space, this can result in a significant amount of time spent simply reading the phase space file. Using this method, the phase space must be read N times, where N is the number of beamlets. When using the beamlets for clinical applications, such as inverse treatment planning where the phase spaces tend to be on the order of ~ 400 - 800 MB in size and the number of beamlets is relatively small (~ 300 - 400), this is perhaps not a big concern. However, when using beamlets to fine tune a 2 GB, 40×40 cm² phase space with 1601 beamlets, this amounts to a significant time loss (~ 100 hrs).

The phase space separation method uses a sorting algorithm on the phase space after the LATCH variables have been assigned beamlet numbers, resulting in individual phase space files ordered by beamlet number. The advantage of this method is that there is an overall time savings. The phase space must only be read $(N/100) \cdot 2$ times (based on the sorting algorithm used) instead of N times as in the previous method. The disadvantage of this method is that there are a large number of phase space files to manage.

	Unmodified	Unmodified *	LATCH Bit Filtering *	Phsp. Separation *
Sm. Phsp. Time	14.1 hrs	10.8 hrs	16.3 hrs	13.1 hrs
Lg. Phsp. Time	24.3 hrs	19.4 hrs	115.5 hrs	23.2 hrs

Table 5.1: The total simulation times for LATCH bit filtered, phase space separated, unmodified BEAMnrc and DOSXYZnrc, and cropped phantom simulations for small size (600 MB) and large size (1.8 GB) phase spaces. (* Simulated on a cropped phantom.)

To obtain an estimate of the time differences in running a simulation using bit filtered beamlets versus phase space separated beamlets, simulations were timed on a single 2.8 GHz Pentium IV processor. For a "small" 600 MB phase space with 65 beamlets, the difference in processing time was found to be 3.5 hrs (24%). For a very large 1.8 GB phase space with 1601 beamlets, the time difference was found to be 3.9 days (500%). Also determined was the time savings of cropping the CT phantom (method described in Chapter 2) prior to running a DOSXYZnrc simulation. On average for clinical plans the time savings was found to be 25% - 35%. This is based on cropping an average of 30% - 40% of the total number of voxels.

Also of interest was the time difference in running a plan using beamlets versus running a standard non-beamlet simulation. The time difference found, when using the phase space separation method on a 1.8 GB phase space file with 1601 beamlets, was 3.8 hrs (19.6%). The main difference in time is attributed to assigning beamlet numbers, sorting phase spaces, and reallocation of dose scoring and geometry arrays within DOSXYZnrc between beamlet runs. Table 5.1 summarizes the run-times for running BEAMnrc with LATCH bit filtering, phase space separation, unmodified and unmodified with phantom cropping.

5.2 The Fine-Tuning of Phase Space A

The results of fine-tuning the particle weights of phase space A¹ are presented for 6 MV and 18 MV accelerator configurations. Given these configurations, the electrons incident on the target are of energies 6 MeV and 18 MeV. It should be noted that these energies are nominal energies since the actual incident energy of the electrons on the target was selected such that the percent depth dose (PDD) profiles match ion chamber measurements in water. Often, the energy found to match PDD profiles deviates slightly from the nominal energy. However, for this accelerator, the incident energy found to best match the 6 MV PDD was in fact 6.0 MeV with a incident full width half maximum of 0.75 mm, and the incident energy found to best match the 18 MV PDD was 18.0 MeV with a incident full width half maximum (FWHM) of 0.75 mm. As mentioned earlier, it is imperative that the PDD's are matched as close as possible since the fine-tuning process does not alter the beam's PDD. The PDD's for the best matched 6 MV beam are shown in Figure 5.3 for field sizes of $40 \times 40 \text{ cm}^2$, $20 \times 20 \text{ cm}^2$, and $10 \times 10 \text{ cm}^2$. The PDD's for the best matched 18 MV beam are shown in Figure 5.4 for $40 \times 40 \text{ cm}^2$, $20 \times 20 \text{ cm}^2$, and $4 \times 4 \text{ cm}^2$. With the exception of the large 18 MV, $40 \times 40 \text{ cm}^2$ field, the PDD curves are well matched. The 18 MV, $40 \times 40 \text{ cm}^2$ field does not model dose deposition correctly in the build-up region.

Ding [33][34] found a similar result when modeling a Clinac 2100EX accelerator. Ding

¹Recall from Chapter 2 that phase space A is scored at 45 cm from the target's upstream surface, above the secondary collimator.

proposed but later disproved that the under-dose calculation in the build-up region is perhaps a result of neutron contamination created in the accelerator head for 18 MV electrons. BEAMnrc does not currently model neutron interactions. Currently the issue remains unresolved.

For this application, a slab of voxels $1.0 \times 1.0 \times 0.5 \text{ cm}^3$ was used to score the Monte Carlo dose for an open $40 \times 40 \text{ cm}^2$ field at a source-to-surface² distance (SSD) of 100 cm. For both energies, the scoring slab was placed in a $50 \times 50 \times 50 \text{ cm}^3$ DOSXYZnrc created water tank, at a depth of d_{max} , which is at a depth of 1.5 cm for the modeled 6.0 MV beam and 3.5 cm for the modeled 18.0 MV beam. The choice of using only a single slab of scoring voxels was made to keep the number of terms in the cost function to a minimum and therefore speed the fine-tuning process. More scoring slabs could certainly have been added, however, the results from just one slab were satisfactory. The choice of placing this scoring plan at d_{max} was made after examining off-axis profiles at various depths. Examining Figure 5.5 it can be seen that, due to electron scatter, the profile discrepancies “wash-out” with depth. Placing the scoring plane at d_{max} ensures that the largest discrepancies are corrected for.

The measured dose distributions for the fine-tuning goal were obtained in two ways. First, they were extracted from Varian’s CadPlan treatment planning system. Since CadPlan dose data is based on open field measurement in water, the assumption was made that CadPlan would produce accurate ($< 2\%$) open field profiles in water.

²the distance from the upstream electron target surface to the incident surface of the phantom

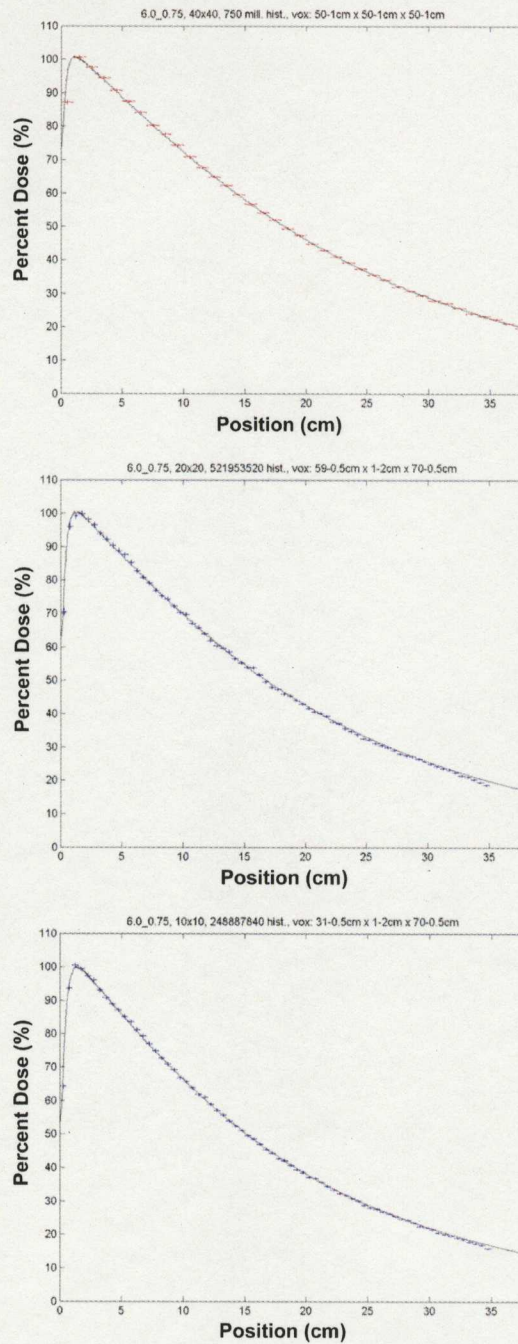


Figure 5.3: The 6 MV percent depth dose (PDD) curves for the best matched incident electron energy (6 MeV) and full width half maximum (0.75 mm) for field sizes of $40 \times 40 \text{ cm}^2$ (*top*), $20 \times 20 \text{ cm}^2$ (*middle*), $10 \times 10 \text{ cm}^2$ (*bottom*) as defined by the collimating jaws. Each curve shows the fraction of maximum dose deposited with depth (cm). The MC data is shown with error bars.

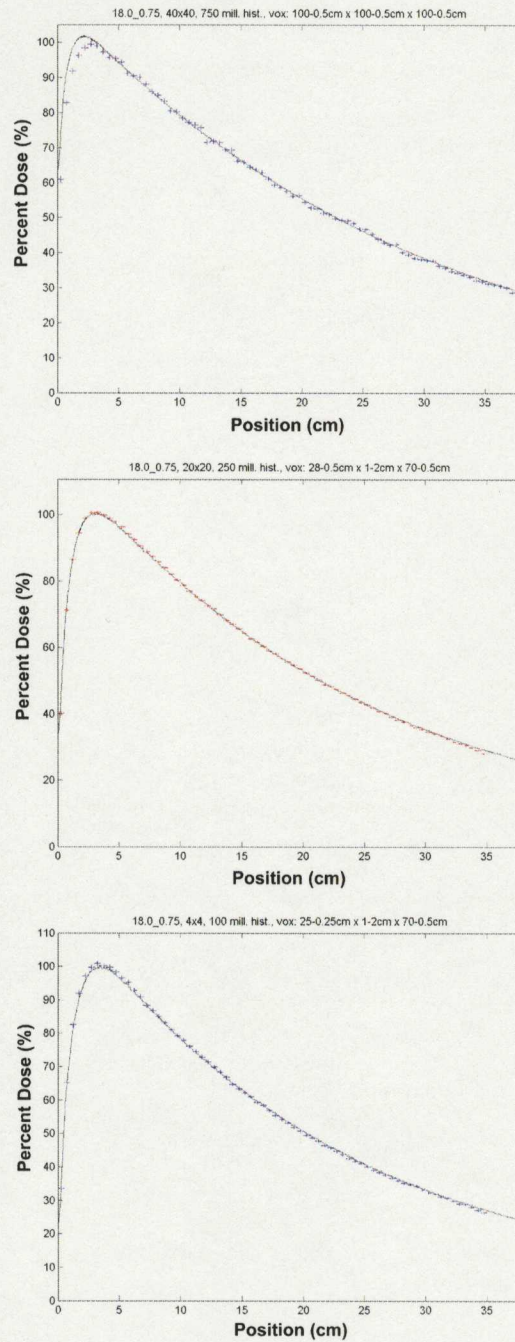


Figure 5.4: The 18 MV percent depth dose (PDD) curves for the best matched incident electron energy (18 MeV) and full width half maximum (0.75 mm) for field sizes of $40 \times 40 \text{ cm}^2$ (top), $20 \times 20 \text{ cm}^2$ (middle), $4 \times 4 \text{ cm}^2$ (bottom) as defined by the collimating jaws. Each curve shows the fraction of maximum dose deposited with depth (cm). The MC data is shown with error bars.

CadPlan dose distributions are also readily available from the treatment planning software. The CadPlan distributions were then interpolated onto a Monte Carlo voxel set using a linear interpolation program written for this thesis. When comparing the MC and CadPlan profiles in the outer penumbral region of the field, relatively large dose differences were observed. These differences are attributed to dose miscalculation by the treatment planning system. A profile comparison of the initial 6 MV 40×40 cm² MC and CadPlan profiles can be seen in Figures 5.5 and 5.7.

The second way that the measured dose distributions were obtained was by actual ion chamber measurements. Ideally a set of ion chamber measurements would have been used from the start as the fine-tuning goal. However, obtaining a full plane of ion chamber measurements is fairly time consuming when compared to the amount of time to obtain the dose distributions from a treatment planning system. Because of the time required to collect the data, measurements were only taken in the plane used in the fine-tuning process. Figure 5.11 shows the initial MC and measured profiles for the 6 MV accelerator configuration. As expected, the ion chamber measurements provided a more realistic initial profile in the beam's outer penumbral region (see Figure 5.11). Figure 5.10 shows a comparison of initial MC, measured, and CadPlan profiles to illustrate the penumbral differences. As seen in this plot, significant discrepancies exist between measured and Cadplan profiles in the beam penumbra and even within the geometrically defined field size beyond 10cm from the central axis.

Using the CadPlan dose distribution as the "measured" dose distribution, the

6 MV fine-tuning process was started with a grain size of 0.150, corresponding to initially changing each beamlet by 15%. The grain size was chosen after examining and estimating initial dose distribution profile differences between MC and CadPlan. For this application, 1601 beamlets (parameters) were used with 2500 data points. The number of degrees of freedom was then calculated to be $N-M = 2500 - 1600 = 900$, where N is the number of terms in the cost function (or data points) and M is the number of parameters to vary (or beamlet weights).

The initial χ^2 , or cost, for the 6 MV beam was calculated to be 1.512×10^6 . The fine-tuning process was run for 350 iterations of each beamlet, reducing the cost to 5.182×10^5 . The enormity of the above costs is likely due to terms in the cost function involving voxels from the penumbral regions which contribute heavily to the overall cost; a result of trying to converge realistic dose distributions to unrealistic ones. Subsequent optimizations using CadPlan data as measurement make use of a reduced "fiducial" region³ to eliminate poorly modeled terms from contributing to the cost. For this application $k_B T$, from the cooling schedule Equation 4.4, was set to zero, thus not allowing any deviation from the measured dose distribution. The results from the fine-tuning process can be seen in Figure 5.6, the profiles of which are taken in the plane of voxels included in the fine-tuning process. A closer look at the initial and final dose distributions is shown in Figures 5.7 and 5.8.

Following the fine-tuning process, the incident phase space (phase space A)

³A region where the measured data can be relied upon

was weighted according to the derived beamlet weight map. The weighted phase space was then re-run into the water tank to verify the corrected profiles matched at other depths in the phantom. Profiles for d_{max} , 10 cm and 20 cm depths are shown in Figure 5.9 for 6 MV.




Figure 5.5: Initial 6 MV Monte Carlo and CadPlan profiles from a $40 \times 40 \text{ cm}^2$ field at depths of 1.5 cm (d_{max}), 10.0 cm, and 20.0 cm in a $50.0 \times 50.0 \times 50.0 \text{ cm}^3$ water tank with a source-to-surface distance (SSD) of 100.0 cm. These profiles show the slight discrepancies in profile shapes that still exist after an initial lengthy commissioning effort.




Figure 5.6: Corrected 6 MV, $40.0 \times 40.0 \text{ cm}^2$ Monte Carlo and CadPlan profiles at a depth of 1.5 cm (d_{max}), in a $50.0 \times 50.0 \times 50.0 \text{ cm}^3$ water tank (100.0 cm SSD).




Figure 5.7: A closer look at the 1.5 cm (d_{max}) profiles of Figure 5.5.




Figure 5.8: A closer look at the 1.5 cm (d_{max}) profiles of Figure 5.6.

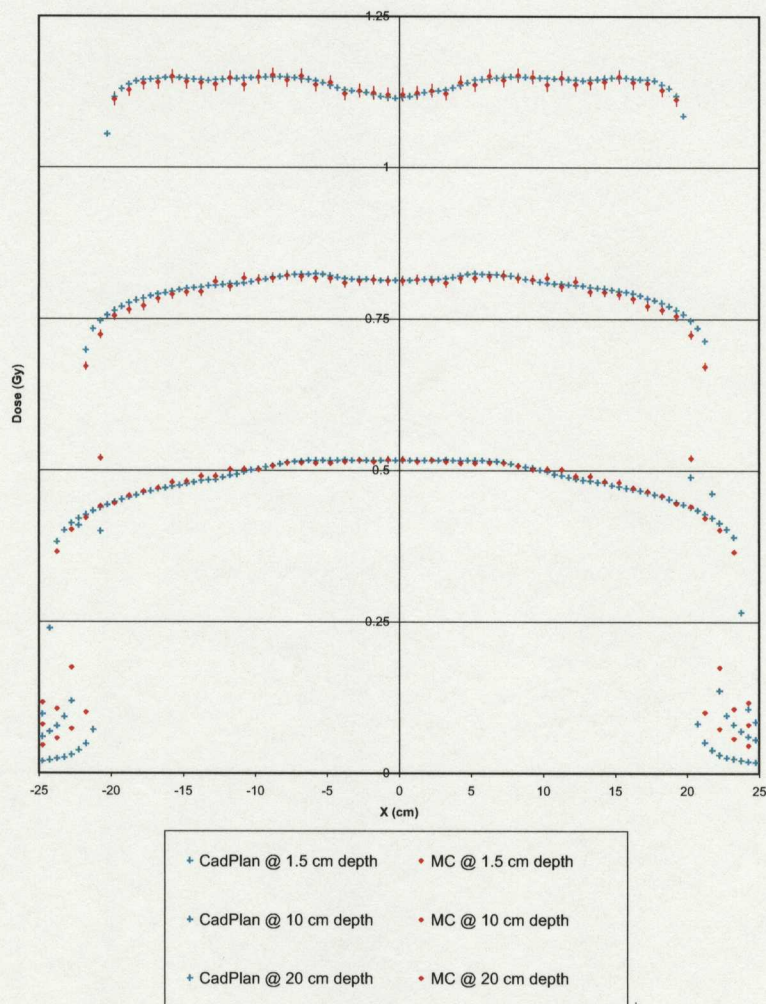


Figure 5.9: 6 MV, $40.0 \times 40.0 \text{ cm}^2$ Monte Carlo and CadPlan profiles at depths of 1.5 cm (d_{max}), 10.0 cm, and 20.0 cm in a $50.0 \times 50.0 \times 50.0 \text{ cm}^3$ water tank (100.0 cm SSD) after weighting the incident phase space A and re-running the simulation with different random number seeds.

The fine-tuning process was then repeated using ion chamber measured data in place of CadPlan data. A MC voxel resolution of $1.0 \text{ cm} \times 1.0 \text{ cm} \times 0.5 \text{ cm}$ was used for calculating the beamlet dose distributions, where 0.5 cm is the depth dimension. This voxel size was chosen to provide a reasonable statistical uncertainty on voxel doses while preserving systematic profile differences. Also, this dimension matches the depth dimension (0.5 cm) of the ion chamber used for the measurements. In order to avoid a sampling error in regions of high dose gradient (*ie.* the field edge penumbra), the MC voxel size was reduced to $0.25 \text{ cm} \times 0.25 \text{ cm} \times 0.5 \text{ cm}$. The initial dose profile differences between MC and ion chamber measured dose distributions is shown in Figure 5.11. A dose difference (delta) plot of the initial dose profile differences is shown in Figure 5.13. The initial cost was calculated (over the entire fine-tuning plane of voxels) to be 4508.0. After the fine-tuning process, the cost was calculated to be 62.1. The number of degrees of freedom for this and all subsequent fine-tuning processes was calculated to be $1600 - 1600 = 0$. The fact that the cost function does not converge precisely to 0 is likely a result of imperfections in the initial MC model. The corrected MC profile is shown compared with measurement in Figure 5.12 and a delta plot of the corrected difference is shown in Figure 5.14.

The incident phase space (phase space A) was then weighted according to the determined beamlet weights and then re-run into the water tank. Figure 5.15 shows the resulting difference plot.




Figure 5.10: 6 MV, $40 \times 40 \text{ cm}^2$ initial Monte Carlo, measured, and CadPlan profiles at a depth of 1.5 cm (d_{max}) in a $50.0 \times 50.0 \times 50.0 \text{ cm}^3$ water tank (100.0 cm SSD).




Figure 5.11: 6 MV, $40 \times 40 \text{ cm}^2$ initial Monte Carlo and measured profiles at a depth of 1.5 cm (d_{max}) in a $50.0 \times 50.0 \times 50.0 \text{ cm}^3$ water tank (100.0 cm SSD). A higher voxel resolution of 0.25 cm has been used in the MC penumbral tails to avoid sampling error. A 1.0 cm voxel resolution has been used for the rest of the field.




Figure 5.12: Corrected 6 MV, $40 \times 40 \text{ cm}^2$ Monte Carlo and measured profiles at a depth of 1.5 cm (d_{max}) in a $50.0 \times 50.0 \times 50.0 \text{ cm}^3$ water tank (100.0 cm SSD).




Figure 5.13: Delta profile plot of the initial dose differences between MC and measured dose distributions for a 6 MV beam in a plane at 1.5 cm depth. $\chi^2 = 4508.0$. Number of degrees of freedom (DOF) = 0.




Figure 5.14: Delta profile plot of the dose differences between corrected MC and measured dose distributions for 6 MV in a plane at 1.5 cm depth. $\chi^2 = 62.1$. $DOF = 0$.




Figure 5.15: Delta profile plot of the dose differences between re-run MC and measured dose distributions for 6 MV in a plane at 1.5 cm depth. $\chi^2 = 142.4$. $DOF = 0$.

To ensure that the phase space corrections held for other field sizes, in particular more clinically used field sizes, the weighted phase space A was simulated through $10 \times 10 \text{ cm}^2$ and $4 \times 4 \text{ cm}^2$ field sizes into the water tank and compared with ion chamber measurements. Figures 5.16 and 5.17 show these comparisons at depths of d_{max} , 10.0 cm and 20.0 cm. A depth dose curve for the $10 \times 10 \text{ cm}^2$ is shown in Figure 5.18⁴.

⁴For brevity, the $4 \times 4 \text{ cm}^2$ field depth dose curve was not included. It should be noted that this curve also showed excellent agreement with measurement.




Figure 5.16: Measured and MC profiles resulting from re-running the 6 MV weighted phase space through $10 \times 10 \text{ cm}^2$ jaws. Profiles shown for depths of 1.5 cm (d_{max}), 10.0 cm, and 20.0 cm in a $50.0 \times 50.0 \times 50.0 \text{ cm}^3$ water tank (100.0 cm SSD).




Figure 5.17: Measured and MC profiles resulting from re-running the 6 MV weighted phase space through $4 \times 4 \text{ cm}^2$ jaws. Profiles shown for depths of 1.5 cm (d_{max}), 10.0 cm, and 20.0 cm in a $50.0 \times 50.0 \times 50.0 \text{ cm}^3$ water tank (100.0 cm SSD).




Figure 5.18: Measured and MC percent depth dose curves resulting from re-running the 6 MV weighted phase space through $10 \times 10 \text{ cm}^2$ jaws in a $50.0 \times 50.0 \times 22.5 \text{ cm}^3$ water tank (100.0 cm SSD).

The entire fine-tuning process was then repeated for an 18 MV beam. For fine-tuning the 18 MV beam, a set of ion chamber measurements was used as a target distribution. Given the poor accuracy of the CadPlan data for the 6 MV beam, no attempt was made to fine-tune the 18 MV beam to CadPlan data. Figure 5.19 shows the initial MC, CadPlan, and measured dose profiles at d_{max} for a 40×40 cm² field. The ion chamber measurements were again only taken in one plane at d_{max} . This decision was made based on the success of fine-tuning the 6 MV beam using only a single plane. The 18 MV fine-tuning process was started with a grain size of 0.400, again based on an initial assessment of profile differences. The initial cost for the 18 MV fine-tuning process was calculated to be 2.70×10^4 . The algorithm was run for 500 iterations of each beamlet reducing the cost function to 86.9. The initial profile difference is shown in Figure 5.20. The corrected profile is shown in Figure 5.21. The re-run profile is shown in Figure 5.22.

Again, to ensure that the phase space corrections held for other field sizes, the weighted 18 MV phase space A was simulated through 10×10 cm² and 4 cm² field sizes into the water tank and compared with ion chamber measurements. Figures 5.27 and 5.28 show these comparisons at depths of d_{max} , 10.0 cm and 20.0 cm depths. A depth dose curve for the 10×10 cm² field has also been included in Figure 5.29⁵.

⁵Again, for brevity, the 4×4 cm² field depth dose curve was not included but showed excellent agreement with measurement.




Figure 5.19: 18 MV, $40 \times 40 \text{ cm}^2$ initial MC, measured, and CadPlan profiles at a depth of 3.3 cm (d_{max}) in a $50.0 \times 50.0 \times 50.0 \text{ cm}^3$ water tank (100.0 cm SSD).




Figure 5.20: 18 MV, $40 \times 40 \text{ cm}^2$ MC and measured ion chamber profiles at a depth of 3.5 cm (d_{max}), in a $50 \times 50 \times 50 \text{ cm}^3$ water tank (100.0 cm SSD). These profiles show the discrepancies in profile shapes that still exist after the initial MC benchmarking effort.




Figure 5.21: Optimized 18 MV, $40 \times 40 \text{ cm}^2$ MC and measured ion chamber profiles at a depth of 3.5 cm (d_{max}) in a $50 \times 50 \times 50 \text{ cm}^3$ water tank (100.0 cm SSD).




Figure 5.22: 18 MV, $40 \times 40 \text{ cm}^2$ MC and measured ion chamber profiles at a depth of 3.5 cm (d_{max}) in a $50 \times 50 \times 50 \text{ cm}^3$ water tank (100.0 cm SSD) after weighting the incident phase space A and re-running the simulation with different random number seeds.

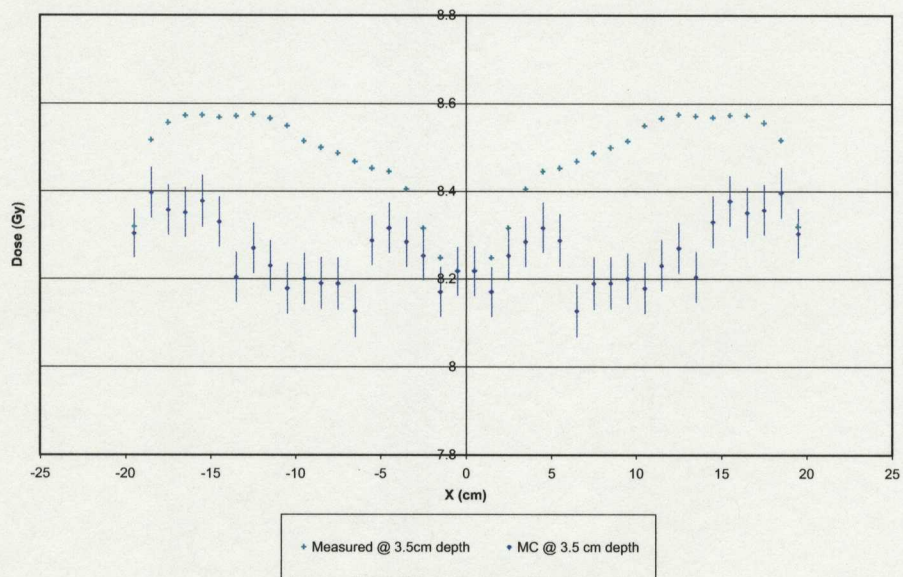


Figure 5.23: A closer look at the 3.5 cm (d_{max}) profile in Figure 5.20.




Figure 5.24: A closer look at the 3.5 cm (d_{max}) profiles in Figure 5.21.




Figure 5.25: Delta profile plot of the initial dose differences between MC and measured dose distributions for 18 MV in a plane at 3.5 cm depth. $\chi^2 = 26988.1$. $DOF = 0$.




Figure 5.26: Delta profile plot of the dose differences between corrected MC and measured dose distributions for 18 MV in a plane at 3.5 cm depth. $\chi^2 = 86.9$. $DOF = 0$.




Figure 5.27: Measured and MC percent depth dose curves resulting from re-running the 18 MV weighted phase space through $10 \times 10 \text{ cm}^2$ jaws. Profiles shown for depths of 3.5 cm (d_{max}), 10.0 cm, and 20.0 cm in a $50.0 \times 50.0 \times 50.0 \text{ cm}^3$ water tank (100.0 cm SSD).




Figure 5.28: Measured and MC percent depth dose curves resulting from re-running the 18 MV weighted phase space through $4 \times 4 \text{ cm}^2$ jaws. Profiles shown for depths of 3.6 cm (d_{max}), 10.0 cm, and 20.0 cm in a $50.0 \times 50.0 \times 50.0 \text{ cm}^3$ water tank (100.0 cm SSD).




Figure 5.29: Measured and MC percent depth dose curves resulting from re-running the 18 MV weighted phase space through $10 \times 10 \text{ cm}^2$ jaws in a $50.0 \times 50.0 \times 50.0 \text{ cm}^3$ water tank (100.0 cm SSD).

To further illustrate the dose differences between measured and MC distributions a dose subtraction was performed in a plane of voxels at d_{max} . Figures 5.32 to 5.31 show the dose differences before and after fine-tuning for both 6 MV and 18 MV beams. Prior to fine-tuning, a visible systematic dose difference is seen with an annular shape. The cause for the discrepancy is likely due to incorrect manufacturer provided data on a component of the accelerator or incorrect incident electron parameters. There is also a visible discrepancy in the corners of the $40 \times 40 \text{ cm}^2$ slab. A possible cause for the discrepancies in the corner is the absence of accelerator components in the MC model. In particular, a thin annular structure used to trim the corners of the optical light field⁶. In reality, the corners of the $40 \times 40 \text{ cm}^2$ ionizing radiation field are clipped by the primary collimator (see Chapter 2). This annular structure ensures that the corners of the optical field are also clipped. This structure is not yet modeled in the MC virtual accelerator.

⁶The optical light field is used for projecting the intended ionizing radiation field onto the patient for field set-up.




Figure 5.30: Dose subtraction between the original 18 MV Monte Carlo and ion chamber measurements in a plane at d_{max} (3.5 cm) for a $40 \times 40 \text{ cm}^2$ field prior to fine-tuning (normalized to maximum dose difference).




Figure 5.31: Dose subtraction between the corrected Monte Carlo and ion chamber measurements in a plane at d_{max} (3.5 cm) for an 18 MV, $40 \times 40 \text{ cm}^2$ field (normalized to dose maximum prior to fine-tuning).




Figure 5.32: Dose subtraction between uncorrected Monte Carlo and ion chamber measurements in a plane at d_{max} (1.5 cm) for an 6 MV, $40 \times 40 \text{ cm}^2$ field prior to fine-tuning (normalized to maximum dose difference).




Figure 5.33: Dose subtraction between the corrected Monte Carlo and ion chamber measurements in a plane at d_{max} (1.5 cm) for an 6 MV, $40 \times 40 \text{ cm}^2$ field (normalized to dose maximum prior to fine-tuning).

In an attempt to test the limits of the fine-tuning process, a correction of the output of a poorly modeled 18 MV accelerator was attempted. The extremely poor initial dose profiles are shown in Figure 5.34. After optimization, the profiles appear to match within uncertainty as seen in Figure 5.35. The weighted phase space was then rerun into the phantom to verify the correction at other depths. As seen in Figure 5.36, the profiles appear to be corrected at d_{max} , but appear to “sag” slightly in the outer field region with depth. This is likely due to the energy spectrum being incorrect to start with in the outer field region. In other words, the off-axis PDD was incorrect to start with. Because the fine-tuning process does not alter the PDD the fine tuning process cannot fully correct this phase space. In this example, initial discrepancies of up to 25% are reduced to discrepancies of up to roughly 6%. This result begins to show the limit to which the fine-tuning process can be taken. This result also demonstrates the importance of achieving initial profiles that are close to correct before applying the fine-tuning process.




Figure 5.34: 18 MV, $40 \times 40 \text{ cm}^2$ Monte Carlo and CadPlan profiles at depths of 3.5 cm (d_{max}), 10 cm, and 20 cm in a $50 \times 50 \times 50 \text{ cm}^3$ water tank (100.0 cm SSD). These profiles show the large discrepancies in profile shapes that can occur from a poorly modeled accelerator.




Figure 5.35: Optimized 18 MV, $40 \times 40 \text{ cm}^2$ Monte Carlo and CadPlan profiles at a depth of 3.5 cm (d_{max}) in a $50 \times 50 \times 50 \text{ cm}^3$ water tank (100.0 cm SSD).




Figure 5.36: 18 MV, $40 \times 40 \text{ cm}^2$ Monte Carlo and CadPlan profiles at depths of 3.5 cm (d_{max}), 10 cm, and 20 cm in a $50 \times 50 \times 50 \text{ cm}^3$ water tank (100.0 cm SSD) after weighting the incident phase space A and re-running the simulation with different random number seeds.

5.3 Treatment Plan Remodulation

The results for the remodulation of treatment plans are presented for an artificial phantom with a planned treatment volume (PTV) defined at the intersection of a vertical lung-tissue interface. This phantom was chosen as it is known to be extremely difficult for pencil beam convolution algorithms to calculate dose near density interfaces. The phantom used is shown in Figure 5.37. A PTV structure was defined within the phantom (*red*). To this phantom, a 5×5 cm² field with an SSD of 100 cm was placed. The field size was then further reduced using the MLC to create a 4×4 cm² square field. The field placement is shown in Figure 5.42. Using the MLC to create a square field is not something that would normally be done clinically (just the jaws would be used); however, the use of the MLC was needed in order to provide an initial fluence file that could later be modified once beamlet weights were determined. Given this set-up, a poorly modeled pencil beam convolution calculated dose distribution is obtained. To remind the reader of the purpose of this application, a remodulation of the poorly modeled plan will be derived using the method described in Chapter 4. This remodulation will allow the pencil beam calculated dose distribution to actually be delivered, as approved by the oncologist. Figure 5.39 shows film, CadPlan, and MC dose profiles taken at the lung-tissue interface at a depth of 10 cm (*left*) and 15 cm (*right*). The initial dose volume histograms (DVH) for the Eclipse and Monte Carlo calculated dose distributions are shown in Figure 5.40 and

a dose subtraction prior to optimization is shown in Figure 5.42. For all intents and purposes, one can consider the MC dose distribution to very closely resemble that actually delivered. The reader is encouraged to note the difference between the Eclipse calculated DVH that is assumed to be delivered, and the MC calculated DVH that is actually delivered. This application is therefore of great importance to the Medical Physicist (it is the job of the Medical Physicist to ensure that the prescribed dose distribution is actually delivered). It is for this reason that lung cancer patients are not yet treated using IMRT in the BC Cancer Agency.




Figure 5.37: An axial CT image of the artificial lung interface phantom used for remodulating miscalculated treatment plans (*top*). An image of the reconstructed 3-D CT phantom (*bottom*). A planned treatment volume (PTV) has been contoured (*red*) in a location known to be difficult for pencil beam convolution algorithms to calculate dose, at the lung air interface. Other random structures have been contoured (*magenta*) adjacent to the PTV.




Figure 5.38: Monte Carlo dose profiles resulting from $4 \times 4 \text{ cm}^2$ field incident on the lung-air interface.




Figure 5.39: Film, CadPlan (TPS), and MC dose profiles taken at a lung-tissue interface in an artificial phantom at a depth of 10 cm (*left*) and 15 cm (*right*) for 18MV. The pencil beam convolution algorithm used in CadPlan miscalculates the dose in the inhomogeneous region resulting in an unrealistically high dose to the lung.

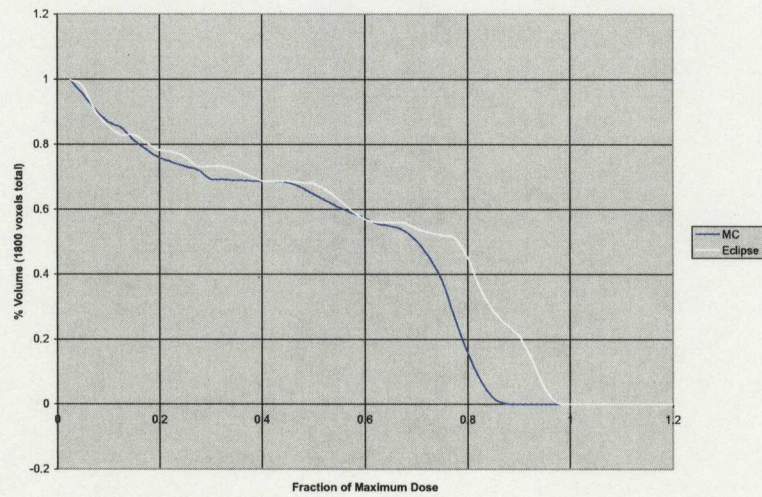


Figure 5.40: Dose volume histograms prior to optimization within the lung-air interface PTV (see Figure 5.37) for Eclipse and Monte Carlo dose distributions.

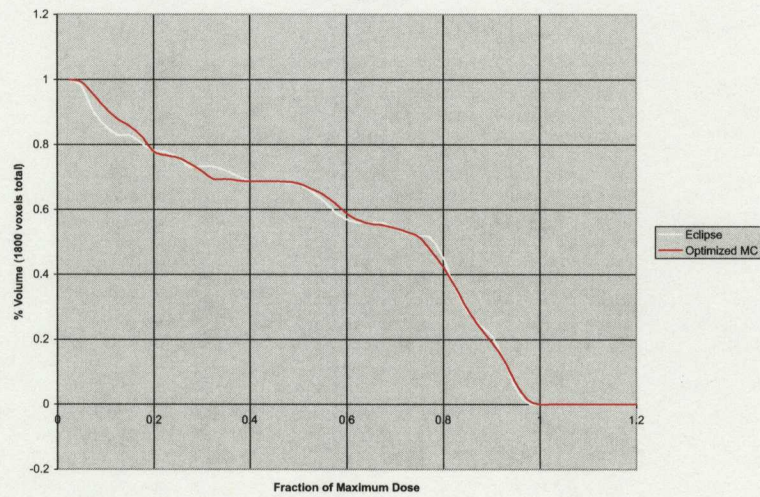


Figure 5.41: Dose volume histograms after optimization within the lung-air interface PTV (see Figure 5.37) for Eclipse and Monte Carlo dose distributions.




Figure 5.42: Close-up dose subtraction between Eclipse and Monte Carlo dose distributions prior to optimization within the lung-air interface PTV (see Figure 5.37). For display purposes the density has been removed outside of the PTV.




Figure 5.43: Close-up dose subtraction after optimization within the lung-air interface PTV (see Figure 5.37) for Eclipse and Monte Carlo dose distributions. For display purposes the density has been removed outside of the PTV.

The dimensions of the beamlets used for this optimization were chosen to be $0.5 \times 0.5 \text{ cm}^2$. These dimensions were chosen to correspond with the MLC leaf width (0.5 cm when projected 100 cm from the target) in the direction perpendicular to the MLC leaves since this is the beam delivery resolution limit. The beamlet dimension parallel to the MLC leaves could have theoretically been smaller; however, the results obtained using a resolution of 0.5 cm in this direction were satisfactory.

The optimization was run for 1000 iterations of each beamlet with an initial grain size of 0.20. The average dose to the PTV was calculated prior to optimizing as 0.392 Gy using CadPlan and 0.328 Gy using BEAMnrc or $\sim 10.5\%$ difference. After optimization, the average MC dose to the PTV was calculated to be 0.394 Gy or $\sim 0.4\%$ difference from the prescribed dose. This corresponds to a 10.1% improvement in average PTV dose accuracy for this plan. The resulting dose volume histogram after optimization are shown in Figure 5.41. A dose subtraction between CadPlan and Monte Carlo dose distributions within the PTV is shown in Figure 5.43. Within the PTV the prescribed Eclipse dose distribution has been successfully matched with the exception of the lower $4 \times 4 \text{ cm}^2$ field edge. A proposed cause for this difference is the fact that the secondary collimating jaws are not modeled correctly in the Eclipse system. In reality, the jaws consist of two pairs of Tungsten blocks that overlap with one pair above the other (As shown in Chapter 2, Figure 2.6). This overlap causes asymmetry between the X and Y field directions that is not modeled in the Eclipse system but is in the MC model. The band of dose difference is likely a result of this

difference in accelerator models.

From the determined weight map corrections, beamlet weights could then be used to scale the treatment planning system's optimal fluence map. The leaf sequence is then able to be recalculated incorporating the derived correction.

5.4 Direct Aperture Optimization

For the previous application, inverse treatment planning using Monte Carlo calculated beamlets was used to correct a miscalculated dose distribution. One could argue that MC beamlets should be used from the start. Therefore, in this application, inverse treatment planning is performed using only MC calculated beamlets. To achieve this, MC beamlet dose distributions were formatted to allow their incorporation into the Direct Aperture Optimization (DAO) algorithm described in Chapter 4⁷.

For this application, MC generated beamlet dose distributions replaced pencil beam convolution (PBC) calculated dose distributions for a pseudo-clinical 7-field head and neck IMRT plan shown in Figure 5.44. The plan was then optimized using 10 apertures per field. Three critical structures were defined for the optimization and weighted as in the actual clinical plan. The optimization was run until 100% of the PTV received 95% of the prescribed 60 Gy. A selection of the determined optimized apertures are shown in Figure 5.45. Once the optimization was completed the optimized apertures were imported into Eclipse and a pencil beam dose calculation

⁷The results presented in this section are part of a collaboration with A. Bergman, University of British Columbia, Canada.

was performed based on the apertures determined using the MC beamlets. Figure 5.46 shows a comparison of the optimized MC calculated dose distribution and pencil beam calculated dose distribution. From this figure one can see distinct dose differences, especially in the sinus cavities. Dose volume histograms (DVH) are shown for both the MC and PBC calculated dose distributions within the defined structures in Figure 5.47. Also included in this figure, are the resulting MC DVH's from re-running the MC plan after the optimized apertures have been determined. The plan was re-run to correctly model interleaf leakage and leaf-tip effects. As mentioned in Chapter 4 these effects are not included in the DAO. For this plan 57 monitor units are needed to deliver 60 Gy to the PTV. The equivalent fluence-based optimization using a dynamic MLC leaf sequence would require on the order of 97 Gy. This reduced number of monitor units is a major advantage of DAO allowing faster per fraction delivery times and potentially allowing more patients to be treated with IMRT in a given day. Reducing the number of monitor units required also results in lower leaf transmission dose as well as less total body scatter dose from the accelerator's treatment head.

As mentioned earlier, the scope of this application was to enable the use of BEAMnrc calculated Monte Carlo beamlets into the Direct Aperture Optimization algorithm, not to quantify the improvement in dose calculation accuracy by using MC beamlets. Further research by Bergman *et al* is currently underway to evaluate the improvement in dose calculation accuracy between using the MC calculated beamlets

generated using the methods described in this thesis and PBC dose calculation in Direct Aperture optimization.




Figure 5.44: Field set-up for a 7-field head and neck IMRT plan. The PTV is contoured in red. Organs at risk include: Right Temporal Lobe (*green*), Left Temporal Lobe (*pink*), and Brainstem (*white*) [10].




Figure 5.45: Beam's eye view of optimized apertures from Field 2 for the 7-Field direct aperture optimization [10].




Figure 5.46: Dose distributions from direct aperture optimization using MC beamlets (*left*) and the pencil beam convolution calculated dose distribution using the MC determined apertures (*right*). The PTV can be seen outlined in red [10].




Figure 5.47: Optimized dose volume histograms within the structures defined in Figure 5.44 for MC calculated, PBC calculated and re-run MC calculated dose distributions [10].

Chapter 6

Conclusions

A method was developed to calculate Monte Carlo (MC) beamlet dose distributions. The cumulative dose distributions from complete sets of beamlets were shown to be in good agreement with the dose distributions from standard BEAMnrc field simulations. Although the time required to obtain a complete set of MC beamlets was greater than the time required to simulate a radiotherapy field in BEAMnrc, this time was minimized by the development and application of a phase space sorting algorithm.

In a first application, MC beamlets were generated to fine-tune the output of BEAMnrc “virtual” accelerators. The weights of the MC generated beamlets were adjusted to align $40 \times 40 \text{ cm}^2$ MC dose distributions to both ion chamber measurements and CadPlan dose distributions for 6 MV and 18 MV beam configurations.

For the 6 MV configuration, the χ^2 was reduced by a factor of 98.6% when aligned to ion chamber measurements. For the 18 MV configuration, the χ^2 was reduced by a factor of 99.7% when aligned to ion chamber measurement.

For both energies, the corrected phase spaces were re-run with jaws set to 10

$\times 10 \text{ cm}^2$ and $4 \times 4 \text{ cm}^2$. The results suggest that the corrections to phase space A are valid for other field sizes.

To test the limitations of this method, an attempt was made to align the output of a very poorly modeled 18 MV beam to CadPlan data. The χ^2 was reduced, but failed to match the measured profile in the outer field regions at depth. The result helps quantify the limit to how much a phase space can be fine-tuned.

The method of fine-tuning phase spaces has the potential to dramatically reduce the time and effort needed in the initial stages of MC implementation, and could contribute to its wider use in a clinical setting. It also offers a solution to better model accelerators without sufficient knowledge of specific components and accelerators that cannot be modeled due to a lack of available component modules.

In a second application, a method was developed to determine modulation corrections to miscalculated Intensity Modulated Radiation Therapy (IMRT) treatment plans that have been derived using a pencil beam algorithm. To achieve this, MC beamlet weight optimization was implemented. To create a situation where a dose difference arises between MC and a pencil beam dose calculation, a single $4 \times 4 \text{ cm}^2$ IMRT field was placed on the lung-tissue interface of an artificial phantom. Dose distributions were obtained using Varian Eclipse treatment planning software and BEAMnrc. The BEAMnrc distribution was then converged to the Eclipse distribution by optimizing beamlet weights and a correction to the beam fluence was obtained. After optimization, the mean MC PTV dose was found to be 10% closer

to the expected mean PTV dose.

The purpose for the above optimization was not so much to develop a method of correcting miscalculated treatment plans, but instead to provide a means to initiate inverse treatment planning using MC dose calculation. By starting the MC beamlet optimization near the final optimization goal (starting from the Eclipse determined plan), the MC optimization algorithm did not need to be all that efficient. As such, a simple simulated annealing algorithm with a quadratic cost function was used.

In an attempt to develop an inverse treatment planning system using MC dose calculation, MC beamlet dose distributions were formatted to allow their incorporation into a Direct Aperture Optimization (DAO) algorithm. The MC beamlet dose distributions successfully replaced pencil beam convolution (PBC) calculated dose distributions for a 7-field head and neck IMRT plan. Further research will evaluate the improvement in dose calculation accuracy between using MC dose calculation and PBC dose calculation in DAO. By amalgamating MC beamlets and DAO, the groundwork for a MC inverse treatment planning system was established.

Bibliography

- [1] Rogers, D. W. O., *et al.* "BEAM: A Monte Carlo code to simulate radiotherapy treatment units." *Medical Physics* 22 (1995): 503-524.
- [2] Mah, E., *et al.* "Experimental evaluation of a 2D and a 3D electron pencil beam algorithm." *Physics in Medicine and Biology* 34 (1989): 1179-1194.
- [3] Krieger, T., Sauer, O. A. "Monte Carlo- versus pencil-beam-/collapsed-cone-dose calculation in a heterogeneous multi-layer phantom." *Physics in Medicine and Biology* 50(2005): 859-868.
- [4] Cranmer-Sargison, G., *et al.* "Modelling an extreme waterlung interface using a single pencil beam algorithm and the Monte Carlo method." *Physics in Medicine and Biology* 49 (2004): 1557-1567.
- [5] Karzmark, C. J. "Advances in linear accelerator design for radiotherapy." *Medical Physics* 11 (2) (1984).
- [6] Johns, H. E. and Cunningham, J. R. *The Physics of Radiology*. Charles C. Thomas, 4th Edition, 1983.
- [7] Popescu, I. A., *et al.* "Clinical implementation of a Monte Carlo based QA process for IMRT", *ESTRO 23*, Amsterdam, The Netherlands (2004).
- [8] Popescu, I. A., *et al.* "Absolute dose calculations for Monte Carlo simulations of radiotherapy beams." *Physics in Medicine and Biology* 50 (2005): 3375-3392.
- [9] Treurniet, J. A., *et al.* "BEAMnrc, DOSXYZnrc and BEAMDP GUI User's Manual." *NRC Report PIRS 0623 (rev C)* (2004).
- [10] Bergman, A. *et al.* "Implementation of Monte Carlo calculated beamlet dose distributions in a direct aperture optimization algorithm for IMRT." *ESTRO Lisbon, Portugal* (2005).
- [11] Greening J. R. *Fundamentals of Radiation Dosimetry*. Adam Hilger Ltd., 1981.

- [12] Sobol, I. M. (Messer, R. *et al*, Trans.). *The Monte Carlo Method*. Chicago: University of Chicago Press (1974).
- [13] Kawrakow, I. and Rogers, D. W. O. "The EGSnrc Code System: Monte Carlo simulation of electron and photon transport." *Technical Report PIRS-701*, National Research Council of Canada, Ottawa, Canada, (2000).
- [14] Mackie, T. R., *et al*. "The OMEGA project: Electron dose planning using Monte Carlo simulation." *Medical Physics* 17 (1990).
- [15] Cygler, J., *et al*. "Electron dose distributions in experimental phantoms: a comparison with 2D pencil beam calculations." *Physics in Medicine and Biology* 32 (1987): 1073-1086.
- [16] Jette, D. and Walker, S. "Electron dose calculation using multiple-scattering theory: Evaluation of a new model for inhomogeneities." *Medical Physics* 19 (1992): 1241-1254.
- [17] Morawska-Kaczynska, M. and Huizenga, H. "Numerical calculation of energy deposition by broad high-energy electron beams: II. Multi-layered geometry." *Physics in Medicine and Biology* 37 (1992): 2103-2116.
- [18] Shiu, A. S. and Hogstrom, K. R. "Pencil-beam redefinition algorithm for electron dose distributions." *Medical Physics* 18 (1991): 7-18.
- [19] Mackie, T. R. ed. Kase, K., *et al*. *Dosimetry of Ionizing Radiation*. Academic Press, New York (1990): 541-620.
- [20] Rogers, D. W. O. and Bielajew, A. F. "Monte Carlo techniques of electron and photon transport for radiation dosimetry." in *The Dosimetry of Ionizing Radiation*. Vol III, eds. K.R. Kase, B.E. Bjarngard, and F.H. Attix. Academic Press (1990): 427-539.
- [21] Laub, W., *et al*. "Monte Carlo dose computation for IMRT optimization." *Physics in Medicine and Biology* 45 (2000): 1741-1754.
- [22] Pawlicki, T. and Ma, C. M. "Monte Carlo simulation for MLC-based intensity-modulated radiotherapy." *Medical Dosimetry* 26 (2001): 157-168.
- [23] Siebers, J., *et al*. "A method for photon beam Monte Carlo multileaf collimator particle transport." *Physics in Medicine and Biology* 47 (2002): 3225-3249.
- [24] Reynaert, N., *et al*. "MCDE: a new Monte Carlo dose engine for IMRT" *Physics in Medicine and Biology* 49 (2004): 235-241.

- [25] Popescu, I. A., et al. "A Monte Carlo based MU calculator for IMRT", *BC Cancer Agency Annual Cancer Conference*, Vancouver, Canada (2003).
- [26] Walters, B. R. B., et al. "History by history statistical estimators in the BEAM code system." *Medical Physics* 29 (2002): 2745-2752.
- [27] Webb, S. "Optimization of conformal radiotherapy dose distributions by simulated annealing." *Physics in Medicine and Biology* 34 (10) (1989): 1349-1370.
- [28] Geman, S. and Geman, D. "Stochastic relaxation, Gibbs distribution and the Bayesian restoration of images." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6(6) (1984): 721-741.
- [29] Jeraj, R. and Keall, P. "Monte Carlo Based Inverse Treatment Planning." *Physics in Medicine and Biology* 44 (1999): 1885-1896.
- [30] Llacer, J., et al. "Absence of multiple local minima effects in intensity modulated optimization with dose-volume constraints." *Physics in Medicine and Biology* 48(2) (2003): 183-210.
- [31] Jeraj, R., et al. "Optimizer convergence and local minima errors and their clinical importance." *Physics in Medicine and Biology* 48(17) (2003): 2809-2827.
- [32] Shepard, D. M., et al. "Direct aperture optimization: A turnkey solution for IMRT." *Medical Physics* 29 1007-1018
- [33] Ding, G. X. "Energy spectra, angular spread, fluence profiles and dose distributions of 6 and 18 MV photon beams: results of Monte Carlo simulations for a Varian 2100EX accelerator." *Physics in Medicine and Biology* 47 (2002): 1025-1046.
- [34] Ding, G. X., et al. "Are neutrons responsible for the dose discrepancies between Monte Carlo calculations and measurements in the build-up region for a high-energy photon beam?" *Physics in Medicine and Biology* 47 (2002): 3251-3261.
- [35] Arnfield, M. R., et al. "A method for determining multileaf collimator transmission and scatter for dynamic intensity modulated radiotherapy." *Medical Physics* 27 (2000): 2231-2241.
- [36] Kawrakow, I., et al. "EGSnrcMP: the multi-platform environment for EGSnrc." *Technical Report PIRS-877*, National Research Council of Canada, Ottawa, Canada, (2003).
- [37] LoSasso, T., et al. "Physical and dosimetric aspects of a multileaf collimation system used in the dynamic mode for implementing intensity modulated radiotherapy." *Medical Physics* 25 (1998): 1919-27.

-
- [38] Walters, B. R. B. and Rogers, D. W. O. "QA for the BEAM System; Component Modules, Variance Reduction Options and Source Routines." *NRC Report PIRS-509k* (1995).
- [39] Wang, X., *et al.* "Dosimetric verification of intensity-modulated fields." *Medical Physics* 23 (1996): 317-326.

Appendix A

NRC Monte Carlo Code Changes

The following is a list of changes that have been made to existing Monte Carlo programs.

> displays the change made
< displays the original code

1. In order to use the existing MC nrc code for creating beamlets the following changes were made:

1.1 To /home/karlbush/egsnrc/dosxyznrc/dosxyznrc.mortran:
Note Changes to this file only affect the single user

Modification#1: When running a restart, the only thing I want carried over to the next run is the dose deposited data found in the .egsdatt file. This is because restart is not being used in the intended way. The following accomplishes this:

```
1951a1952,1965
>
>   nnread = 0;"initialize counter"
>   nsmis=0;      nmis=0; ncaseold=0;
>   nsrjct=0;     nsoutside=0;
>   planarefe=0.; planarefp=0.;
>   planarfe=0.; planarfp=0.;
>   NofREPEAT = 0; esrc=0.;      nestep=0;
>
>   IF(IPARALLEL>1 & PARNUM>0) [
>     nnphsp=INT((PARNUM-1)*nshist/IPARALLEL)+1;
>   ]
>   ELSE[
>     nnphsp=1;
>   ]
```

Modification#2: This modification prevents DOSXYZnrc from restarting (or not reaching the end) of the phase space in an attempt to deliver the number of particles, "NCASE", requested by the user. DOSXYZnrc is forced to stop after reading through the phasespace exactly once + one restarted particle (the extra particle is a trade off to not modifying srcxyznrc.mortran, the part shared by all users). This modification is significant in that it places more importance on the number of recyclings not the number of histories requested. When requesting histories, the user should request an NCASE of greater than the number of "non-rejectable" particles in the incident phasespace.

```
2098c2112
<
---
>     IF(nnread .le. nshist)[
```

```
2183c2197
<
---
>]
```

Modification#3: Changing the number of ph-sp particles simulated (N_simulated) to reflect the fact that modification#2 reads exactly the (number of particles in the phasespace + 1) and recycles every particle NRCYCL times. Note: modification#1 resets ncaseold to zero upon restart.

```
2233c2248
<     NRCYCL,NCASE+ncaseold,OUTCNT;
---
>     NRCYCL,NCASE+(NRCYCL+1)+ncaseold,OUTCNT;
```

Modification#4: Because ainflu includes all rejected particles and this number changes for each beamlet, this is not a good value to normalize to. Instead the dose is left unnormalized until the dose is converted to absolute dose (using absolute_ainflu.c). This program normalizes all beamlets to the full non-beamlet run's ainflu.

```
2374c2394
<             endep(irl) = endep(irl)*1.602e-10/(amass*ainflu);
---
>             endep(irl) = endep(irl)*1.602e-10/(amass);
```

Modification#5: In order to allow larger ainflu values to be seen the following change was made

```
2428c2448
<     (t5,'No. of particles incident from original source =',1f13.1//
```

> (t5,'No. of particles incident from original source =',1f16.1//

1.2 To home/beam/omega_home/dosxyznrc/srcxyznrc.mortran:

Note changes to this file affect all beam users

Modification#1: In order to use all combinations of the latch the following change was made. Previously one bit corresponded to one area of the accelerator. Now the latch bit is set to the binary equivalent of the beamlet number.

1981,1985c1981,1985

```
<          IF($BTEST(latchi,LATBIT(IBIT))) [EXIT;]"go to excl. check"
<          IF(IBIT=NBIT1)[
<              nsrjct=nsrjct+1;
<              goto :READ-PH-SP-DATA;;
```

```
<          ]"discard the particle"
```

```
>          IF($BTEST(latchi,LATBIT(IBIT))) []
>          ELSE[
>              nsrjct=nsrjct+1;
>              goto :READ-PH-SP-DATA;;
>          ]
```

1999c1999,2002

```
<          goto :READ-PH-SP-DATA;;
```

```
>          "goto :READ-PH-SP-DATA;;
>          ]
>          IF(nsrjct=NBIT1)[
>              goto :READ-PH-SP-DATA;; "Only reject if particle has certain
>              combination of bits"
```

1.3 To home/beam/omega_home/progs/beamdp/beamdp.mortran:

Note changes to this file affect all beam users

Modification#1: This modification allows the use of beamdp to view the beamlets based on latch.

2330c2330

```
<          IF($BTEST(LATCHI,III))EXIT;
```

```
>          IF($BTEST(LATCHI,III)) []
2332c2332,2334
<          IF(I=NBIT1)GOTO :start-of-file-read0;;
---
>          ELSE[
>      GOTO :start-of-file-read0;;
>      ]
```