

2-1-2012

Charged hadron beam therapy : fast computational physics methods

Roy William Keyes

Follow this and additional works at: https://digitalrepository.unm.edu/phyc_etds

Recommended Citation

Keyes, Roy William. "Charged hadron beam therapy : fast computational physics methods." (2012).
https://digitalrepository.unm.edu/phyc_etds/33

This Dissertation is brought to you for free and open access by the Electronic Theses and Dissertations at UNM Digital Repository. It has been accepted for inclusion in Physics & Astronomy ETDs by an authorized administrator of UNM Digital Repository. For more information, please contact disc@unm.edu.

Roy William Keyes

Candidate

Physics and Astronomy

Department

This dissertation is approved, and it is acceptable in quality and form for publication:

Approved by the Dissertation Committee:

Shuang Luan, Chairperson

Michael Holzscheiter

Philip Heintz

Bernd Bassalleck

David Dunlap

Charged Hadron Beam Therapy: Fast Computational Physics Methods

by

Roy William Keyes

B.S., Physics, Rice University, 2001
M.S., Physics, University of New Mexico, 2007
M.S., Nuclear Engineering, University of New Mexico, 2009
concentration Medical Physics

DISSERTATION

Submitted in Partial Fulfillment of the
Requirements for the Degree of

Doctor of Philosophy
Physics

The University of New Mexico

Albuquerque, New Mexico, USA

November 2011

©2011, Roy William Keyes

Dedication

*To my wife with love.
May she never have to read this.*

Acknowledgments

My dissertation committee: Michael Holzscheiter, Shuang Luan, Philip Heintz, David Dunlap, Bernd Bassalleck.

For incredible amounts of encouragement and advice, both technical and otherwise: Michael Holzscheiter, Shuang Luan, Philip Heintz, Niels Bassler.

The UNM Particle Therapy Group: Shuang Luan, Michael Holzscheiter, Daniel Riofrío, Dominic Maes, Dorian Arnold, Adam Hecht

For stimulating lectures and discussions: David Dunlap, Murray Gell-Mann, Carlton Caves, Ivan Deutsch

My fellow travelers in this process: the many physics grad students and medical physics grad students at UNM and elsewhere that have shared this experience with me.

My undergraduate advisor, Arjendu Pattanayak, for lots of advice over many years and opening my eyes to many interesting ideas.

My friends from Rice University who took the long road to the radiological sciences and inspired me along the way: Vivek Bansal, MD, PhD, Daniel Baseman, MD, Stewart Becker, PhD.

Funders (in many, many forms): The National Science Foundation (i.e. the American tax payers), the NSF GK-12 Fellowship, the University of New Mexico, Linac Systems, LLC, the State of New Mexico, the Advanced Cancer Therapy Foundation and its donors, Los Alamos National Laboratory, the Air Force Office of Scientific Research, Amazon.com, LLC.

The very helpful staff in the Physics and Astronomy, Medical Physics, and Computers Science departments, who smoothed over many administrative hurdles inherent in the university process.

And finally, my parents for decades of support and my wife for endless amounts of tolerance.

Eener alleene is nich scheene, aba eener mit eene und denne alleene - det is scheene.

A children's limerick in the Berlin dialect of German. Possibly addressing the importance of distributed computing in medical physics.

The important thing in science is not so much to obtain new facts as to discover new ways of thinking about them.

- William Lawrence Bragg

There are no physicists in the hottest parts of hell, because the existence of a 'hottest part' implies a temperature difference, and any marginally competent physicist would immediately use this to run a heat engine and make some other part of hell comfortably cool. This is obviously impossible.

- Richard Davisson

Charged Hadron Beam Therapy: Fast Computational Physics Methods

by

Roy William Keyes

B.S., Physics, Rice University, 2001

M.S., Physics, University of New Mexico, 2007

M.S., Nuclear Engineering, University of New Mexico, 2009

concentration Medical Physics

Ph.D., Physics, University of New Mexico, 2011

Abstract

Charged hadron beams have been investigated for use in radiation therapy of cancer since the 1940s due to their unique potential to place tightly conformal radiation doses deep inside tissue. This is achieved by exploiting the phenomenon of the so-called Bragg peak. In both research and clinical settings, fast and accurate radiation calculations play a crucial role in charged hadron therapy physics. Unfortunately, physicists are often faced with the fundamental trade off of speed versus accuracy in their calculations. This dissertation addresses this trade off by presenting three computational physics methods for specific and general charged hadron beam therapy calculations.

In this dissertation the pseudo-Monte Carlo method of track repeating is adapted for fast calculations of linear energy transfer (LET) and for fast estimation of dose in

the peripheral regions of the target volume (i.e. secondary dose estimation). Additionally, the first proof-of-concept framework for carrying out massively distributed parallel Monte Carlo calculations for radiation therapy using cloud computing is presented. Performance and accuracy assessments of each calculation method are also presented.

Contents

List of Figures	xiv
List of Tables	xviii
1 Introduction	1
2 History and background of radiation therapy	4
2.1 Introduction	4
2.2 History of radiation therapy	5
2.3 History of charged hadron therapy	7
2.4 The physics of radiation therapy	10
2.4.1 Fluence and Fluence rate	11
2.4.2 Energy fluence and Energy fluence rate	11
2.4.3 Absorbed Dose	12
2.4.4 Exposure	12
2.4.5 Equivalent dose	13

2.4.6	Effective dose	14
2.5	Physics of particle therapy	15
2.5.1	Stopping power and the Bethe formula	15
2.5.2	Range	17
2.5.3	Spread-out Bragg peak	18
2.5.4	Bragg peak width	19
2.5.5	Lateral scattering	20
2.5.6	Linear energy transfer	21
2.6	Radiation biology and charged hadron therapy	22
2.6.1	DNA damage	22
2.6.2	Relative biological effectiveness	23
3	Computational particle therapy background	25
3.1	Introduction	25
3.2	The Monte Carlo method	26
3.2.1	Monte Carlo fundamentals	26
3.2.2	Radiation transport via Monte Carlo	28
3.3	Monte Carlo performance enhancements and derivative techniques	29
3.3.1	Metric of simulation efficiency	30
3.3.2	Transport cutoffs	30
3.3.3	The condensed history method	31

<i>Contents</i>	xi
3.3.4 Variance reduction techniques	31
3.3.5 Particle splitting	31
3.3.6 Russian roulette	32
3.3.7 Interaction forcing	33
3.4 Deterministic methods	34
4 Fast LET mapping via track repeating	35
4.1 Introduction	35
4.2 LET optimization	35
4.2.1 Oxygen enhancement effect and LET	36
4.2.2 RBE and LET	36
4.3 Track repeating	36
4.3.1 Track repeating fundamentals	37
4.4 The LET track repeating algorithm	39
4.4.1 Track generation with Geant4	40
4.5 Benchmarking and performance	44
4.5.1 Benchmarks	44
4.5.2 Performance	47
4.6 Conclusion and outlook	48
5 Fast secondary dose estimation via track repeating	50
5.1 Introduction	50

5.2	Secondary dose and second cancer in particle therapy	51
5.2.1	Sources of secondary dose	51
5.3	Estimating secondary dose	53
5.4	SDE via track repeating	54
5.4.1	Databases	56
5.5	Performance improvement techniques	60
5.5.1	Database filtering	60
5.5.2	Secondary splitting	61
5.5.3	Transport cuts	61
5.6	Benchmarking and performance	62
5.7	Conclusions and outlook	65
6	Distributed radiation therapy simulations via cloud computing	67
6.1	Introduction	67
6.1.1	CPU intensive calculations in radiation therapy	68
6.1.2	Cloud computing	68
6.2	The McCloud distributed processing framework	70
6.2.1	Software and infrastructure	70
6.2.2	Simulations	72
6.3	Results	73
6.4	Discussion	75

<i>Contents</i>	xiii
6.5 Conclusion and outlook	77
7 Conclusions	79
References	81

List of Figures

2.1	An ideal dose distribution, with no dose outside of the tumor region.	5
2.2	Realistic physical dose distributions, with dose spilling outside of the tumor region. The top figure shows dose due to opposed proton beams, while the bottom shows dose due to opposed carbon ion beams. Figure used with permission of Niels Bassler. Similar to figures published in [1]	6
2.3	A comparison of typical single-field photon and proton depth-dose curves in liquid water. Proton data generated with Geant4.9.3 and photon data generated with FLUKA.	8
2.4	A timeline of events in charged hadron therapy. Sources [2, 3, 4].	9
2.5	ICRP 103 neutron weighting values as a function of neutron energy (see Eq. 2.9) [5].	14
2.6	Depth-dose curve of 110 MeV protons in liquid water (simulated with FLUKA 2008).	16
2.7	Depth-dose curves of 70 - 99 MeV protons in liquid water (simulated with FLUKA 2008). This plot shows the energy dependence of the beam stopping depth.	17

2.8	CSDA ranges for various ions. Data from the NIST <i>pstar</i> and <i>astar</i> tools [6].	18
2.9	A proton spread-out Bragg peak with the weighted “pure” Bragg peak components. Such a dose distribution allows for uniform coverage of an extended target. Data generated with Geant4.9.3.	19
2.10	Bragg peak longitudinal width, 100 MeV proton beam in liquid water. Data generated with Geant4.9.3.	20
2.11	The basic structure and components of DNA.	23
3.1	Monte Carlo estimation of π by randomly placing points in a quadrant of the square and determining the proportion which fall within the inscribed unit circle.	27
3.2	Asymptotic approach to the “true” value of π as the number of trials increases.	27
3.3	The particle splitting technique.	32
3.4	The Russian roulette technique.	33
4.1	The basic track repeating concept. a) Generate a database of particle tracks using a full-fledged Monte Carlo package. b) Overlay tracks from database onto scoring grid. c) Accumulate quantity of interest.	38
4.2	Kinetic energy versus step number for 200 proton tracks starting at 50 MeV in liquid water. Some tracks have abrupt drops to ≈ 0 MeV due to inelastic nuclear interactions.	42
4.3	Energy loss versus step number for 200 proton tracks starting at 50 MeV in liquid water.	42

4.4	Step length versus step number for 200 proton tracks starting at 50 MeV in liquid water.	43
4.5	Mean LET versus step number for proton tracks starting at 50 MeV in liquid water. Only tracks with an equal numbers of steps were used to create the mean LET value. This value is very similar to LET versus depth, which requires track propagation.	43
4.6	Unrestricted LET for 50 MeV protons in liquid water generated with track repeating and generated from data from NIST's pstar tool [6].	45
4.7	Absorbed dose of a 50 MeV proton beam in liquid water calculated with track repeating and Geant4.	46
4.8	Track-averaged LET of a 50 MeV proton beam in liquid water calculated with track repeating and Geant4.	47
4.9	Track-averaged LET versus depth of a 50 MeV proton beam in liquid water calculated with track repeating and Geant4.	48
5.1	Absorbed dose in the peripheral region from 75 MeV protons in liquid water (log scale). Data generated with FLUKA.	52
5.2	Run times for 20k primary protons using different Δ values in the binary search algorithm. Δ values below 0.08 caused recursion errors, since the algorithm had to continue to search for steps with energies that met such a strict criterion.	55
5.3	Kinetic energy versus step number for 215 MeV neutrons in liquid water phantom of 400 cm sides.	58
5.4	Frequency histograms of secondary neutrons produced by 250 MeV primary protons. The plot on the right has been re-binned, such that each bin has approximately equal numbers of neutrons.	59

5.5	Absorbed dose from 20,000 unfiltered 120 MeV proton tracks repeated 1,000 times. Second plot is log scale.	62
5.6	Absorbed dose from 20,000 filtered 120 MeV proton tracks repeated 1,000 times, using Filter 2. Second plot is log scale.	63
5.7	Depth-dose curves from 20,000 filtered and unfiltered 120 MeV proton tracks repeated 1,000 times (log scale).	63
5.8	Depth-dose curves in the peripheral region from 120 MeV protons in liquid water. Curves generated by FLUKA, Geant4, and track repeating. Colored bands represent $\pm\sigma$ uncertainty.	64
5.9	Depth-dose curves in the peripheral region from 120 MeV protons in liquid water. Curves generated by FLUKA, Geant4, and track repeating (log scale).	65
6.1	General scheme for calculating dose using the cloud paradigm. . . .	71
6.2	A simple broad-beam proton therapy plan calculated on 130 virtual nodes on AWS's EC2 cloud service using a voxel based phantom. . .	73
6.3	Calculation time versus number of virtual nodes for proton depth-dose curves on AWS's EC2 cloud service.	74

List of Tables

2.1	ICRP 103 radiation weighting factors [5].	13
2.2	IRCP 103 tissue/organ weighting factors used in Eq. 2.10 [5].	15
4.1	Performance of LET calculations for 50 MeV proton beams in liquid water.	48
5.1	Energies and storage sizes of neutron track databases (10k particles each).	59
5.2	Energies and storage sizes of photon track databases (10k particles each).	60
5.3	High-pass transport cut values for track repeating. Range sources [6, 7].	62
6.1	Amazon EC2 pricing for instances running Linux in US eastern region (as of April 2010). *An EC2 Compute Unit is equivalent to a 1.0-1.2 GHz 2007 Opteron or 2007 Xeon processor.	70

Chapter 1

Introduction

The subject of this dissertation sits at the intersection of many different disciplines. On the one hand this dissertation is primarily about computational physics, but it is computational physics as applied to “applied particle physics”. On the other hand this dissertation is about radiation physics and computationally determining quantities of radiobiological interest. Finally, this dissertation investigates aspects of high performance computing, internet scale computing, and the economics of medical physics calculations in a new era of utility-style resource allocation. Ultimately this dissertation seeks to shed some light on the question, “how can we produce the most accurate radiation dose calculations for cutting edge radiation therapy on clinical time scales?”

Cancer is a devastating group of diseases that affects nearly everyone, whether directly or indirectly. Worldwide, cancers accounted for approximately 13% of all deaths in 2007 [8]. According to the National Cancer Institute, as of 2010, approximately half of all cancer patients in the United States receive radiation therapy as part of their treatment [9]. Radiation therapy has proven to be not only one of the most effective therapies for cancer, but also one of the most significant applications of physics in impacting people’s lives.

Over its more than 100 year history, radiation therapy has been one of the most technologically-driven fields of medicine. Therapeutic radiation technologies have made rapid improvements in recent decades, largely due to advances in imaging (in particular CT, MRI, and PET), but fundamentally due to advances in computing. This has allowed for the characterization, calculation, and delivery of highly conformal distributions of radiation dose to target tissues.

While most radiation therapy is performed with externally applied x-rays in the MeV range and internally applied gamma rays in the keV-MeV range, charged hadron beam therapy, or “particle therapy”, is an advanced form of radiation therapy and is currently the object of intense interest in the medical and physics communities. This interest stems from the fact that charged particles heavier than electrons deposit most of their energy at the end of their finite range, forming the so-called “Bragg peak”. This property allows for the creation of extremely target-conformal dose distributions. Additionally, some charged particles have higher relative biological effect (RBE) as compared with photons. Chapter 2 covers the background information on particle therapy.

One of the crucial components in the complex chain that makes up particle therapy delivery is dose calculation. While the radiation transport physics necessary for dose calculations is well understood, in implementing dose calculations for clinical scenarios one nearly always faces the fundamental trade-off between speed and accuracy. Stochastic methods, such as Monte Carlo, can yield extremely accurate results, but tend to be extremely slow (or conversely resource intensive). Deterministic methods, such as dose kernel convolution, can be performed very quickly, but often have trouble with scenarios such as inhomogeneous media. Naturally, the desire is to have both speed and accuracy in dose calculations, but often one is chosen over the other. Generally speed is opted for in clinical settings, while accuracy is chosen in less time-constrained research settings. Chapter 3 covers the background information on computational methods employed in particle therapy.

This dissertation centers on fast computational physics methods that try to bridge the gap between speed and accuracy for certain types of calculations in particle therapy. Specifically, the pseudo-Monte Carlo method of “track repeating” will be applied to fast calculation of linear energy transfer (LET) to enable so-called LET painting techniques (Chapter 4) and secondary dose estimation (Chapter 5), i.e. dose outside the target area where fluences are very low. Additionally this dissertation examines the use of the emerging cloud computing paradigm to enable fast Monte Carlo dose calculations on large scale virtual clusters (Chapter 6).

Chapter 2

History and background of radiation therapy

2.1 Introduction

Radiation therapy is the use of ionizing radiation to treat pathological conditions such as cancer. The goal of radiation therapy is to deliver a lethal dose of radiation to the diseased tissue, while sparing surrounding healthy tissue as much as possible from side effects. This goal is realized by understanding the underlying physics of radiation transport, a requisite for dose calculation and treatment planning, and the interactions of radiation with tissues, radiobiology. Charged hadron beam therapy, or “particle therapy”, is an advanced form of radiation therapy which exploits the physics of heavier-than-electron charged particles to produce highly conformal radiation dose distributions. This section will provide background information on radiation therapy and charged hadron therapy in particular.

Figure 2.1 shows the ideal radiation dose distribution - the target tissue is covered with 100% of the “prescription dose”, while the surrounding tissue receives no dose.

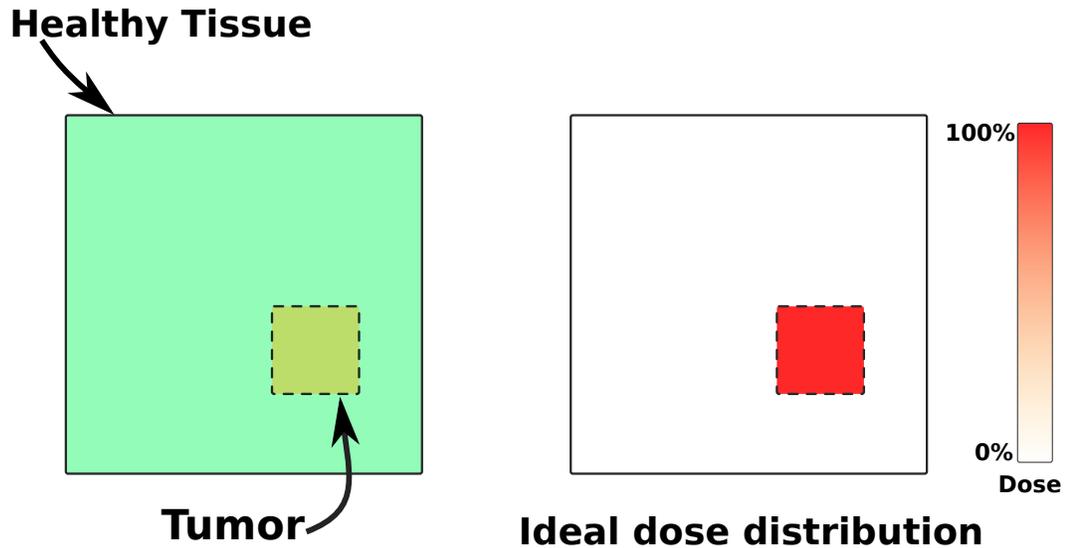


Figure 2.1: An ideal dose distribution, with no dose outside of the tumor region.

The dose distribution in Figure 2.1 is fundamentally impossible due to the physics of energy deposition in media. Figure 2.2 depicts realistic dose distributions, with uniform dose covering the target region, but also (unavoidable) dose spilling into the surrounding, non-target region. Thus, radiation therapy is fundamentally an issue of balancing dose to the target against dose to healthy surrounding tissue.

2.2 History of radiation therapy

Radiation therapy was first attempted shortly after the discovery of x-rays by Roentgen in 1895 [10]. The first reported radiation therapy treatments were performed in 1896 by Despeignes and Freund, among others, and involved exposing patients to low energy x-rays [11, 12]. Shortly thereafter important discoveries and advances

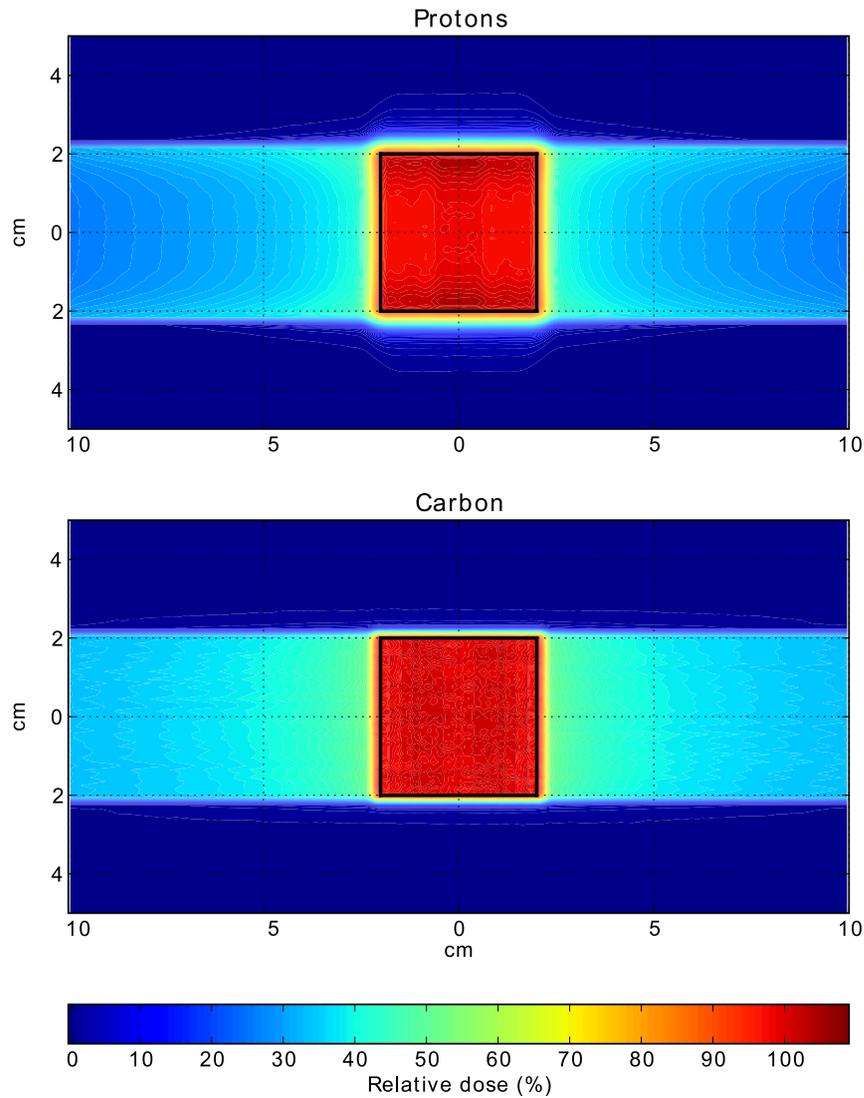


Figure 2.2: Realistic physical dose distributions, with dose spilling outside of the tumor region. The top figure shows dose due to opposed proton beams, while the bottom shows dose due to opposed carbon ion beams. Figure used with permission of Niels Bassler. Similar to figures published in [1]

were made, such as the observation of radio-carcinogenesis (1902), radium implants (1903), and fractionated treatment (1928), which laid the basis for what was to come [13]. Over the next century, treatment techniques were greatly refined as the physics of radiation and radioactivity, as well as radiobiology, were more thoroughly under-

stood. Currently the most common forms of radiation therapy involve the use of MeV range x-rays and electrons from accelerators (external beam radiotherapy) and keV-MeV range gamma rays from implantable sources (brachytherapy).

2.3 History of charged hadron therapy

The use of charged hadron beams, dubbed “particle therapy”, has been of interest to the medical and physics communities since the 1940’s. Since the early 2000’s a growing number of clinical particle therapy facilities has been commissioned and interest has increased accordingly. As of 2011, more than 84,000 patients have received particle therapy worldwide [14].

The idea of charged hadron beam therapy for cancer treatment was first proposed by Robert R. Wilson in 1946 [15]. Wilson’s idea was based on the well established property of ionization by heavier-than-electron charged particles called a “Bragg curve”, wherein the greatest amount of ionization takes place at the end of the path, the so-called “Bragg peak”. This phenomenon was discovered by William H. Bragg in 1903 [16]. In contrast, the maximum of the ionization, and thus deposited dose, from an external photon beam is near the beginning of its path (see Fig. 2.3). Wilson recognized that the location and shape of the Bragg peak would be potentially valuable in radiation therapy. Compared with photon beams, charged hadrons offered the potential for more conformal dose distributions, sparing healthy tissue.

Starting in 1948 the first efforts were made to investigate particle therapy on animals, with the first humans treated in 1954, both at Lawrence Radiation Laboratory in Berkeley, California [17]. Subsequent experiments at Berkeley and Harvard in the US, Uppsala in Sweden, Dubna and Moscow in the USSR, and other locations investigated the use of a range of beams. These included protons, deuterons, alpha

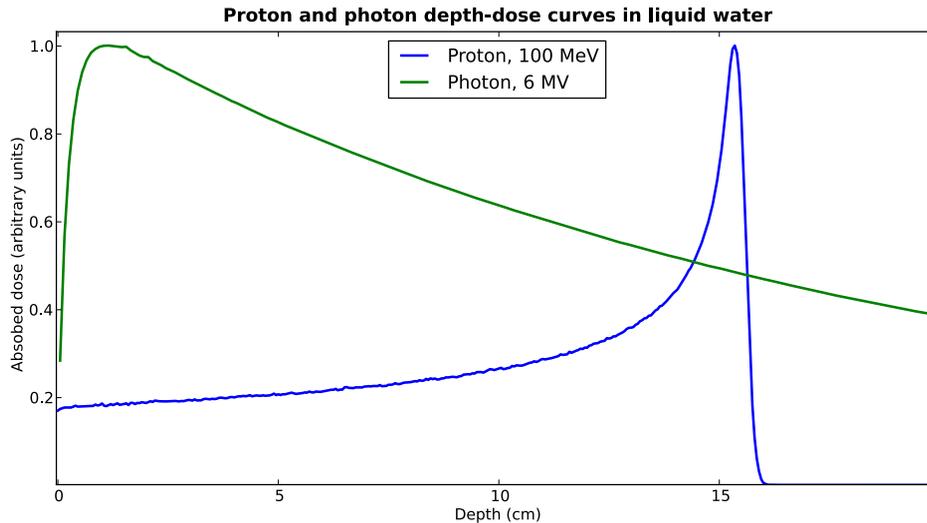


Figure 2.3: A comparison of typical single-field photon and proton depth-dose curves in liquid water. Proton data generated with Geant4.9.3 and photon data generated with FLUKA².

particles, and ions including Li, C, N, O, Ne, and Ar [2].

During much of the early research efforts into charged hadron beam therapy, the direction of research was largely dictated by the capabilities and limitations of available accelerators. This meant that some types of experiments had only short windows of time, as the necessary accelerators were modified for other experiments (e.g. beam energies increased) or decommissioned [2].

Starting in the early 1970's, a large scale research effort began to investigate negative pions for radiation therapy. Near the end of their range, negative pions are captured by nuclei, producing nuclear fragments in the so-called “star formation” process. The star formation process was hypothesized to be advantageous for radiation therapy and a number of institutions built or planned pion therapy programs [2]. Ultimately these programs were decommissioned in the early 1990's after some 1100

²Simulation of 6 MV beam from a Siemens linac by Niels Bassler, who provided this data.

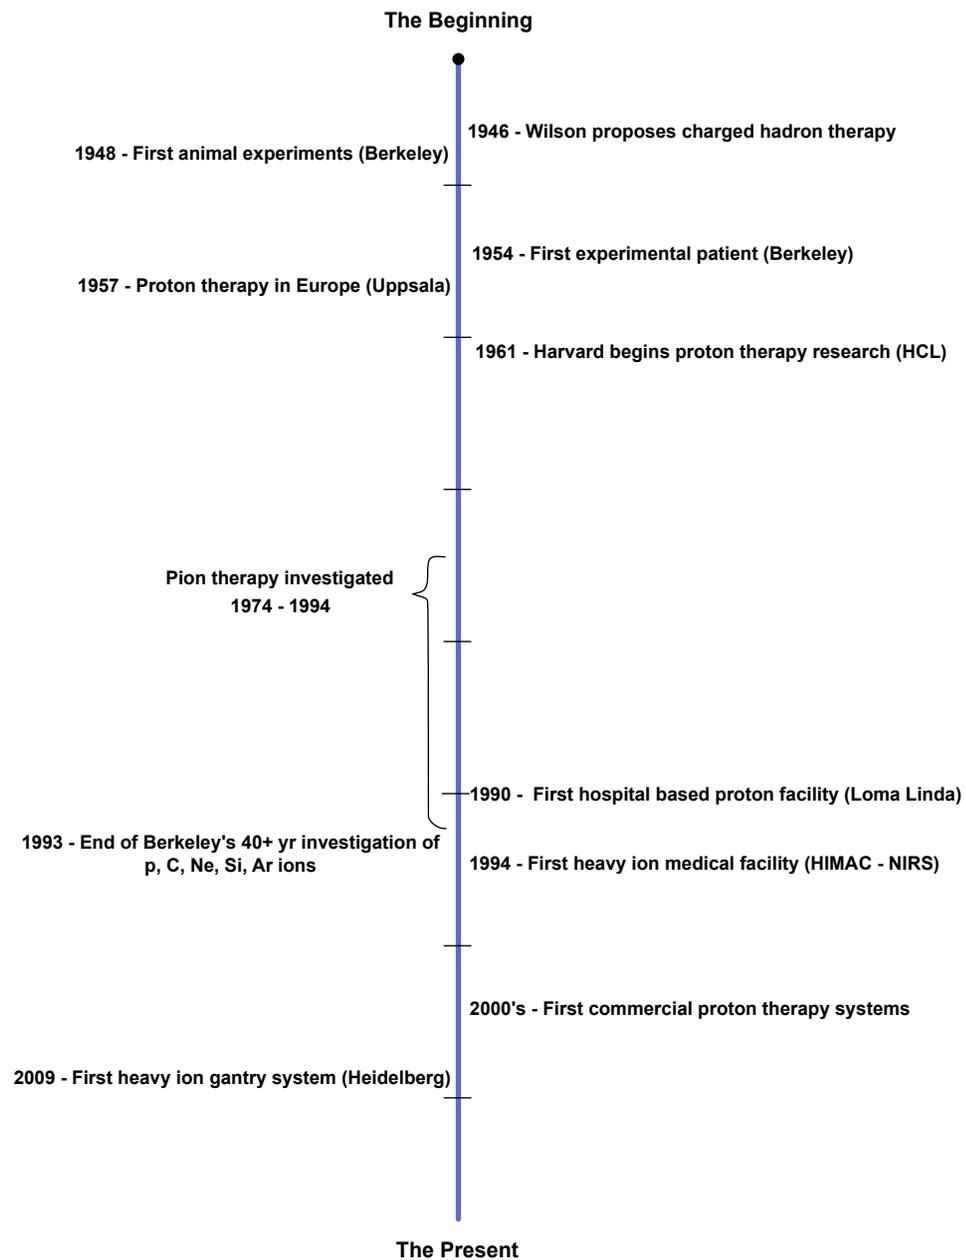


Figure 2.4: A timeline of events in charged hadron therapy. Sources [2, 3, 4].

patients were treated, but knowledge gained carried over into later clinical particle therapy efforts [14].

Another project investigating the use of exotic particles for use in radiation therapy has been the Antiproton Cell Experiment (a.k.a. AD-4/ACE) based at CERN on the border of Switzerland and France [18]. Antiprotons, being the antimatter counterparts of protons, behave in a very similar fashion as protons in media, except when they interact with nuclei and annihilate. The premise of the AD-4/ACE project is that some of this annihilation energy will deposit locally near the annihilation site, boosting the tumor dose. Additionally, heavy ion fragments will contribute to greater radiobiological effects than protons, allowing for lower total dose. Since starting in 2003, this project, lead by University of New Mexico researchers, has established the dosimetric properties of antiprotons in water and provided early estimates of the relative biological effectiveness (see Sec. 2.6.2) [19, 20, 21, 22, 23, 24, 25, 26, 1].

The modern era of particle therapy began in 1990 with the commissioning of the first hospital based proton therapy facility at Loma Linda University, using both a purpose built accelerator and a gantry allowing the rotation of the beam about the patient [27]. In 1994 NIRS in Japan commissioned its Heavy Ion Medical Accelerator facility (HIMAC), the first purpose built accelerator for carbon ions designed for medical use [28]. As of 2011, there are some thirty facilities in operation worldwide delivering proton therapy with another fifteen proposed for construction [14]. Also as of 2011, six facilities are delivering carbon ion therapy with another six planned [14].

2.4 The physics of radiation therapy

In radiological physics, two types of ionizing radiation are considered: directly ionizing particles, such as electrons and protons, and indirectly ionizing particles, such as photons and neutrons. Directly ionizing particles interact primarily through Coulomb type interactions, resulting in direct energy transfer [29]. Indirectly ionizing particles interact through a large number of mechanisms, largely dependent on

the type and energy of the particle, as well as the composition of the media [29]. Photons, for example, will interact with increasing probability as a function of energy in the following order of interactions: Rayleigh scattering, photo-electric effect, Compton scattering, and electron-positron pair production [29]. They also partake in other interactions, including photo-fission. The interactions of neutrons are more complex as result of the large range of nuclear reactions. Neutrons have highly varying interaction cross-sections depending on the medium and incident particle energy. The interactions include neutron capture, neutron induced fission, elastic scattering, and inelastic scattering.

A number of quantities are defined by convention in radiological physics. Some of the most important quantities are defined below.

2.4.1 Fluence and Fluence rate

Fluence is the differential of the expectation value of the quantity of particles passing through a great circle area of da over an arbitrary time interval.

$$\Phi = \frac{dN}{da} \quad (2.1)$$

Units such as m^{-2} are used [29]. The fluence rate is then,

$$\phi = \frac{d\Phi}{dt} = \frac{d}{dt} \left(\frac{dN}{da} \right) \quad (2.2)$$

2.4.2 Energy fluence and Energy fluence rate

Energy fluence describes the differential total energy passing through a great circle area in an arbitrary time interval. If \mathbb{E} is the total energy and $N(E)$ is the expectation value of the number of particles with energy value E , then

$$\Psi = \frac{d\mathbb{E}}{da} = \frac{\int_0^{E_{max}} E dN(E) dE}{da} = \int_0^{E_{max}} E \Phi(E) dE \quad (2.3)$$

with units of joules or MeV [29]. Similarly, energy fluence rate is then,

$$\psi = \frac{d\Psi}{dt} = \frac{d}{dt} \left(\frac{\mathbb{E}}{da} \right) \quad (2.4)$$

2.4.3 Absorbed Dose

The most important quantity considered in radiological physics is the “radiation dose” deposited in a medium by the incident radiation. The standard quantity is called “absorbed dose” and defined as the energy deposited per unit mass with the SI unit of gray (Gy), which is equal to joules per kilogram [29].

$$D = \frac{dE}{dm} \quad (2.5)$$

An absorbed dose as small as 4 Gy over the entire body can be lethal to humans [13].

2.4.4 Exposure

While absorbed dose is a clearly defined quantity, it is not always readily measurable. As such, other quantities, such as “exposure” are measured. Exposure is defined as

$$X = \frac{dQ}{dm} \quad (2.6)$$

where dQ is the charge liberated in the air of mass dm [29]. Under conditions of charged particle equilibrium (i.e. an equal number of charged particles are scattered into as out of the cross-sectional detection area), exposure in air can be directly related to absorbed dose by the equation

$$D = X \times \frac{\bar{W}}{e} \quad (2.7)$$

where \bar{W} is the mean energy expended in a gas per ion formed and e is the charge of the electron [29].

2.4.5 Equivalent dose

Equivalent dose is a quantity defined by the International Commission on Radiological Protection (ICRP), among others. It is designed to take into account the differing biological effects of different radiation types. Equivalent dose is defined as

$$H = \sum_R D_R \times w_R(E) \quad (2.8)$$

where D_R is the absorbed dose due to a given radiation type (e.g. gamma rays) and $w_R(E)$ are unitless biological weighting factors for each type of radiation, which may in turn depend on the energy of the radiation incident on the volume of interest [5]. The standard weighting factors, as defined by ICRP are found in Table 2.1. Equivalent dose is measured in the SI unit sievert (Sv), which is equivalent to joules per kilogram [29]. As such, Gy and Sv both measure deposited dose, but Sv designates that radiobiological weighting factors have been taken into account.

Radiation type	Weighting factor
Photons	1
Electrons and muons	1
Protons and pions	2
Alpha particles, fission fragments, heavy ions	20
Neutrons	Energy dependent (w_r in Eq. 2.9)

Table 2.1: ICRP 103 radiation weighting factors [5].

$$w_R(E_n) = \begin{cases} 2.5 + 18.2e^{-[\ln(E_n)]^2/6} & E_n < 1 \text{ MeV} \\ 5.0 + 17.0e^{-[\ln(2E_n)]^2/6} & 1 \text{ MeV} \leq E_n \leq 50 \text{ MeV} \\ 2.5 + 3.25e^{-[\ln(0.04E_n)]^2/6} & E_n > 50 \text{ MeV} \end{cases} \quad (2.9)$$

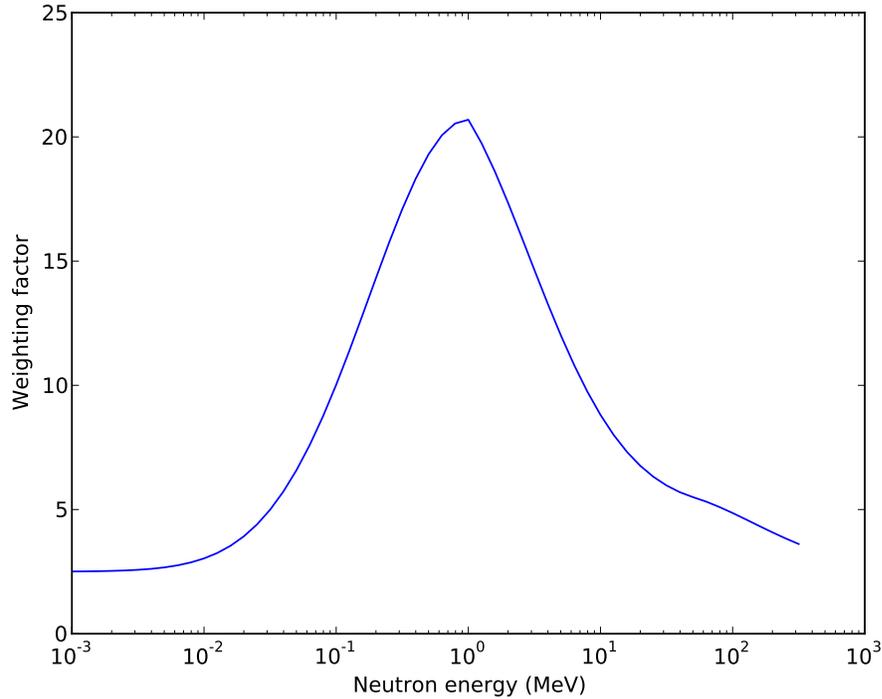


Figure 2.5: ICRP 103 neutron weighting values as a function of neutron energy (see Eq. 2.9) [5].

2.4.6 Effective dose

Effective dose is a quantity that attempts to “normalize” differing irradiation geometries that may affect a human, such that the result can be used for risk assessment, radiation protection, and regulatory purposes. In other words, effective dose attempts to offer a single quantity that allows comparison of exposure to, for example, a single hand with exposure of the entire body. This is accomplished by adding organ/tissue weighting factors to the equivalent dose definition. Effective dose is defined as follows,

$$\mathcal{E} = \sum_{T,R} D_{T,R} \times w_T \times w_R(E) \quad (2.10)$$

where $D_{T,R}$ is the dose in a given tissue/organ volume due to a given radiation type,

R , and w_T and $w_R(E)$ are the tissue/organ weighting factors and radiation weighting factors respectively. Effective dose was most recently defined in ICRP 103 [5].

Tissue/organ	w_T	$\sum w_T$
Red bone marrow, Colon, Lung, Stomach, Breasts, Remainder tissues	0.12	0.72
Gonads	0.08	0.08
Bladder, Liver, Esophagus, Thyroid	0.04	0.16
Skin, Bone surface, Salivary glands, Brain	0.01	0.04
Total		1.00

Table 2.2: ICRP 103 tissue/organ weighting factors used in Eq. 2.10 [5].

2.5 Physics of particle therapy

The advantages of charged hadron beams for radiation therapy are due to the underlying physics of charged hadrons passing through media and depositing energy.

2.5.1 Stopping power and the Bethe formula

The key feature of ion beams recognized as useful by Wilson was the characteristic depth-dose curve, or Bragg curve. The Bragg curve was theoretically described by Bethe [30] in what came to be known as the Bethe (or sometimes Bloch-Bethe) formula, Eq. 2.11.

$$\frac{dE}{dx} = \frac{4\pi}{m_e c^2} \cdot \frac{nz^2}{\beta^2} \cdot \left(\frac{e^2}{4\pi\epsilon_0}\right)^2 \cdot \left[\ln\left(\frac{2m_e c^2 \beta^2}{I \cdot (1 - \beta^2)}\right) - \beta^2 \right] \quad (2.11)$$

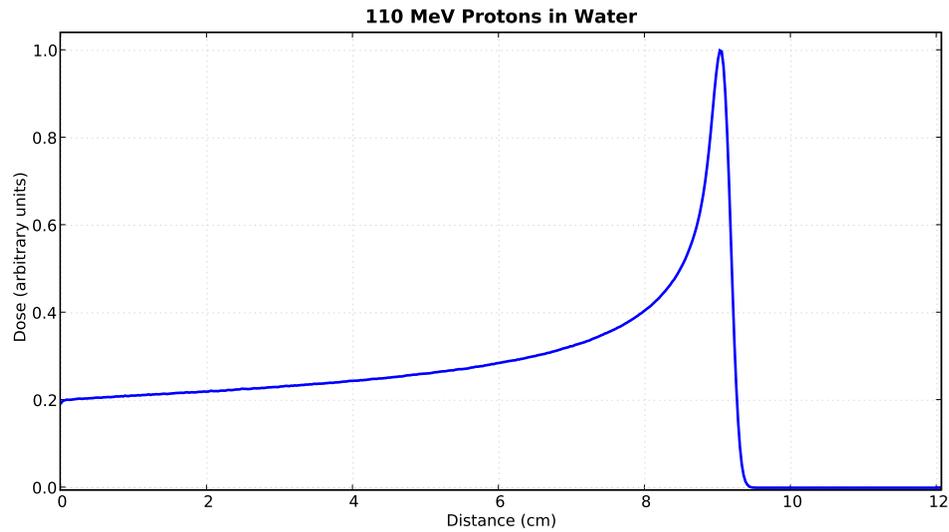


Figure 2.6: Depth-dose curve of 110 MeV protons in liquid water (simulated with FLUKA 2008).

E - energy lost to media	m_e - rest mass of the electron
x - distance traveled by the particle	n - electron density of media
z - particle charge	I - mean excitation potential media
e - charge of the electron	$\beta = v/c$

where β gives the velocity dependence. The Bethe formula is written in terms of $\frac{dE}{dx}$, a quantity known as “stopping power”. Integrating this results in the depth-dose curve, as seen in Figure 2.6. The most important feature of the Bethe formula for radiation therapy applications is the dependence of the energy deposition on β , which in turn determines the location of the Bragg peak (the maximum dose deposition point). By controlling the kinetic energy of the beam, the depth of the Bragg peak can be controlled (see Figure 2.7).

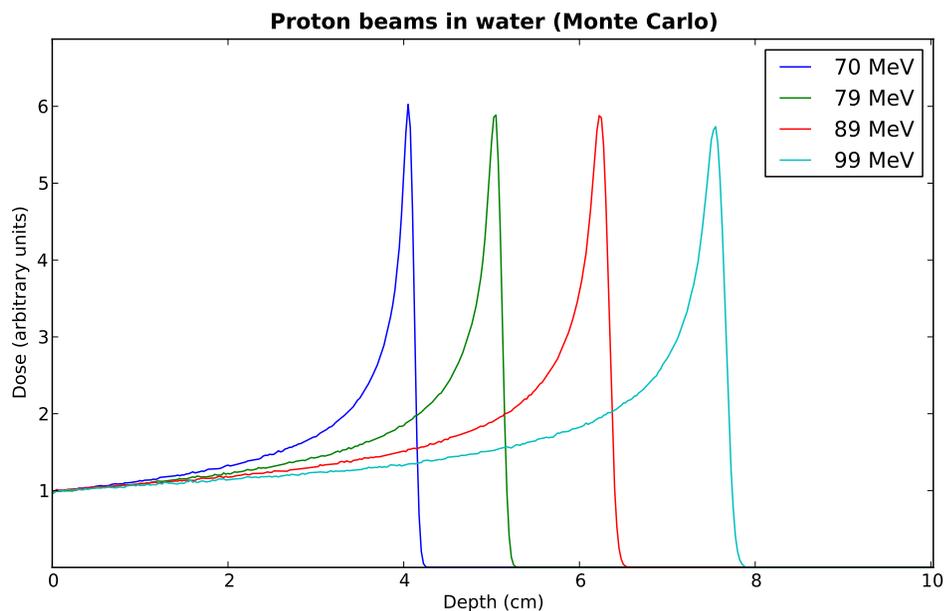


Figure 2.7: Depth-dose curves of 70 - 99 MeV protons in liquid water (simulated with FLUKA 2008). This plot shows the energy dependence of the beam stopping depth.

2.5.2 Range

The mean depth at which a charged particle comes to rest in a medium, or range, can be estimated by integrating the stopping power, Eq. 2.11, as follows,

$$R_{CSDA} = \int_0^{E_0} \left(\frac{dE}{dx} \right)^{-1} dE \quad (2.12)$$

This is known as the continuous slowing down approximation [29]. Figure 2.8 shows the CSDA ranges for various ions in liquid water over an energy range relevant to particle therapy.

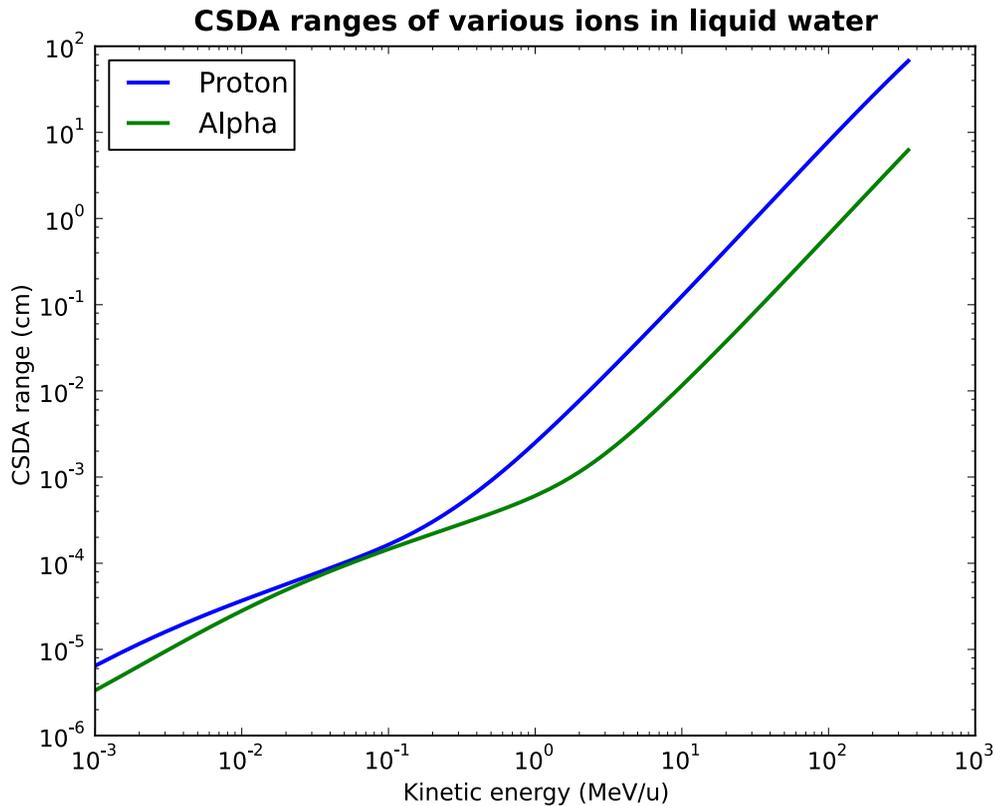


Figure 2.8: CSDA ranges for various ions. Data from the NIST *pstar* and *astar* tools [6].

2.5.3 Spread-out Bragg peak

In his 1946 paper, Wilson posited that by superposing many beams with different kinetic energies, an extended, uniform dose, or spread-out Bragg peak (SOBP) could be created to provide tumor coverage (see Figure 2.9). The SOBP concept remains the cornerstone in particle therapy today.

To construct the SOBP, the mono-energetic Bragg curves are treated as members of a basis set and added in a weighted sum to achieve the desired curve to within some tolerances. These weights are usually determined by solving a least squares or similar problem, such as described by Eq. 2.13.

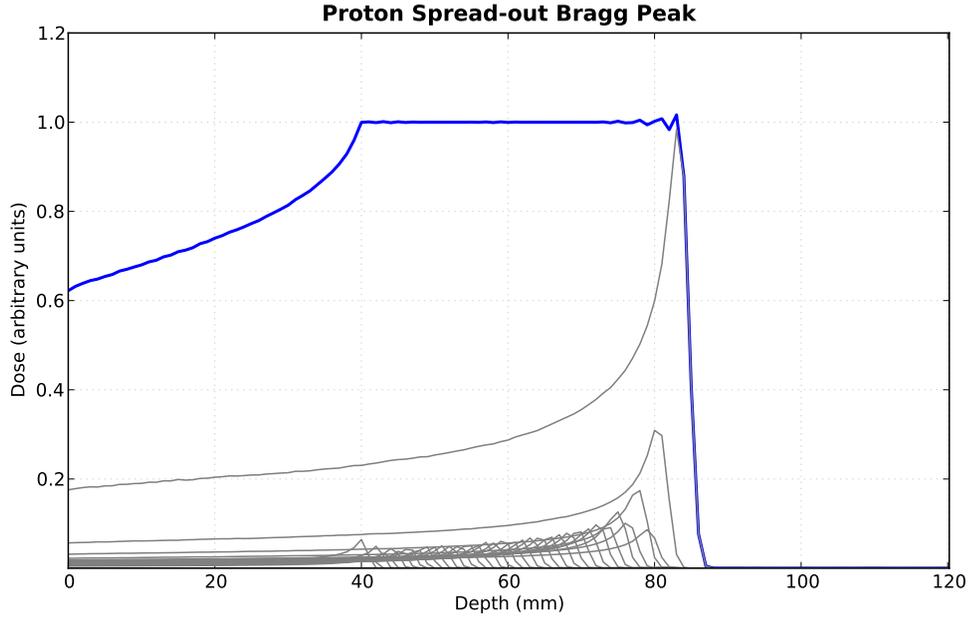


Figure 2.9: A proton spread-out Bragg peak with the weighted “pure” Bragg peak components. Such a dose distribution allows for uniform coverage of an extended target. Data generated with Geant4.9.3.

$$\min\{|\mathbf{X} \cdot \vec{\omega} - \vec{y}|^2\} \text{ with } \begin{cases} \mathbf{X} = \text{the matrix of beam doses} \\ \vec{\omega} = \text{a vector of weights for the beams} \\ \vec{y} = \text{the desired dose distribution} \end{cases} \quad (2.13)$$

2.5.4 Bragg peak width

While the Bragg curve has a very localized peak, its width (see Fig. 2.10) increases at greater depth, corresponding to greater initial energies. This phenomenon is similar to the SOBP, but is due to inherent multiple scattering and energy straggling effects, which take on a normal distribution. The Bragg peak width (FWHM) is described by the following formula [31],

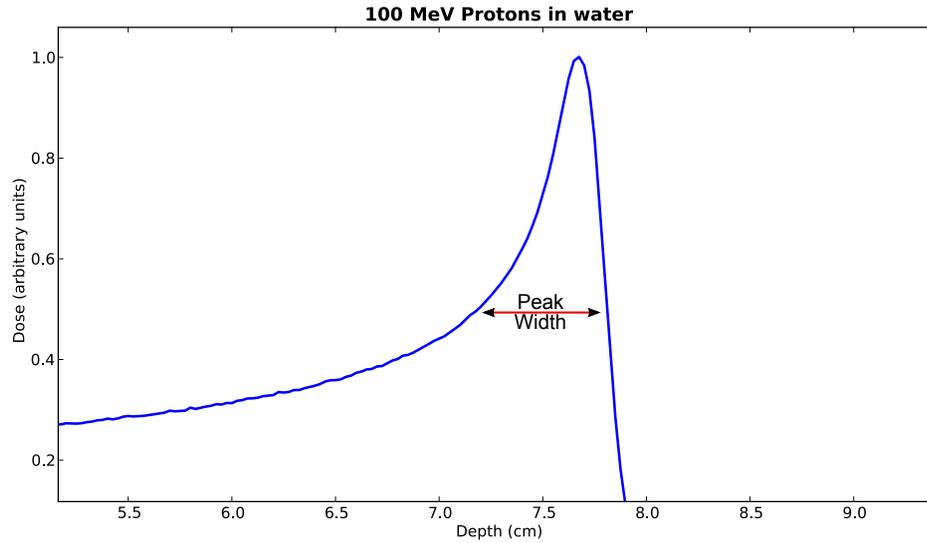


Figure 2.10: Bragg peak longitudinal width, 100 MeV proton beam in liquid water. Data generated with Geant4.9.3.

$$s(x) = \frac{1}{\sqrt{2\pi}\sigma_x} \exp - \frac{(x - R(E))^2}{2\sigma_x^2} \quad (2.14)$$

$$\sigma_x = 0.012 \times R(E)^{0.951} A^{-0.5} \quad (2.15)$$

$$FWHM = 2\sqrt{2 \ln 2} \sigma_x \quad (2.16)$$

where $s(x)$ is the distribution of ranges, $R(E)$ is the mean range (see Sec. 2.5.2), and A is the atomic number of incident particle.

2.5.5 Lateral scattering

The lateral shape of the energy deposition due to a charged hadron beam is described by Molière scattering theory [32]. Multiple scattering events lead to a Gaussian

distribution of dose around the beam center, described by Eq.'s 2.17 and 2.18.

$$f(\theta) = \frac{1}{\sqrt{2\pi}\sigma_\theta} \exp\left(-\frac{\theta^2}{2\sigma_\theta}\right) \quad (2.17)$$

$$\sigma_\theta = \frac{14.1\text{MeV}}{\beta pc} Z \sqrt{\frac{d}{L_{rad}} \left(1 - \frac{1}{9} \log_{10} \frac{d}{L_{rad}}\right)} \quad (2.18)$$

p - particle momentum L_{rad} - radiation length
 d - target thickness $\beta = v/c$
 θ - angle of scatter

These equations give the dependence of the lateral dose distribution width on particle momentum and target thickness, which in the case of therapeutic beams can be interpreted as depth into the target.

2.5.6 Linear energy transfer

Linear energy transfer, or LET, is a specific form of stopping power with particular relevance to radiobiology. Like standard stopping power (see Sec. 2.5.1), LET is a measure of energy lost per unit length, but in this case it is the energy transferred to the immediate vicinity. For photon-electron interactions, *restricted* LET, L_Δ , which does not count energy carried away by secondaries with kinetic energy over a certain threshold, Δ , is used [29].

$$L_\Delta = \frac{dE_\Delta}{dx} \quad (2.19)$$

For heavy particles, *unrestricted* LET, L_∞ , which is equal to the collisional stopping power, is most appropriate [33].

$$L_{\infty} = \frac{dE_{\infty}}{dx} \quad (2.20)$$

LET is generally weighted either by fluence of the primary particles, “track-averaged LET” (Eq. 2.21) or weighted by dose, “dose-averaged LET” (Eq. 2.22) [33].

$$L_{\text{track}}(\mathbf{x}) = \frac{\int_0^{\infty} \Phi_r(\mathbf{x}) L_{\infty}(r) dr}{\int_0^{\infty} \Phi_r(\mathbf{x}) dr} \quad (2.21)$$

$$L_{\text{dose}}(\mathbf{x}) = \frac{\int_0^{\infty} \Phi_r(\mathbf{x}) L_{\infty}^2(r) dr}{\int_0^{\infty} \Phi_r(\mathbf{x}) L_{\infty}(r) dr} \quad (2.22)$$

where Φ_r is the particle fluence with residual range r . LET generally has units of keV/ μm .

2.6 Radiation biology and charged hadron therapy

While the physics of particle beams plays a paramount role in designing a suitable treatment, radiation biology plays an equally important role. The reader is referred to Hall [13] as a primary reference for radiobiological background.

2.6.1 DNA damage

The primary mechanism of radiation leading to cell death is thought to be sufficient damage caused to the deoxyribonucleic acid (DNA) contained within the cell’s nucleus. Because of this, much of radiation biology takes place on the length scale of DNA, which has a width of approximately 2 nm.

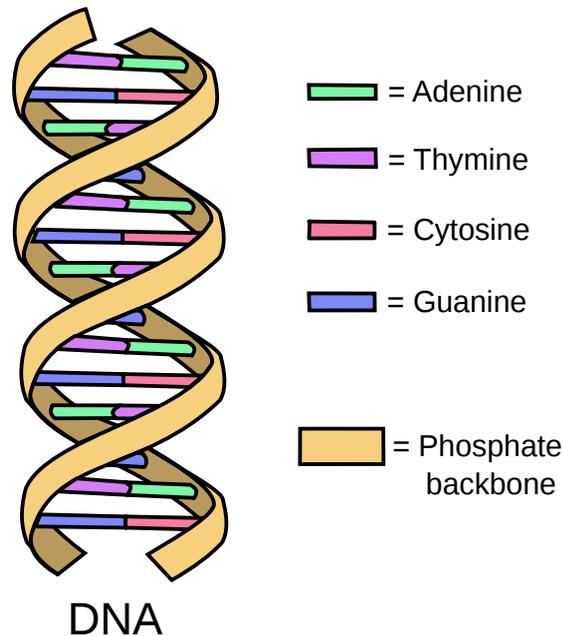


Figure 2.11: The basic structure and components of DNA.

Charged particles are able to cause direct damage to DNA bonds, while photons indirectly damage DNA, primarily through the creation of free radicals, such as hydroxyl ($\text{OH}\cdot$), which then attack DNA molecular bonds. The famous “double helix” ladder-like structure of DNA (Fig. 2.11) is susceptible to single- and double-strand breaks of the phosphate backbone on which the nucleotide bases are held. Double-strand breaks are more difficult for the cell to repair than single strand breaks and are in turn more lethal.

2.6.2 Relative biological effectiveness

On the laboratory bench scale, radiation effects are quantified using the quantity relative biological effectiveness (RBE). RBE is the relative effect that the same absorbed dose of one type of radiation has versus the same absorbed dose of a reference

type of radiation (often Co-60 gamma rays or 250 KV x-rays) [Hall]. Reference radiation is assigned an RBE of 1.0 by definition. The biological effect (or “end point”) under consideration is arbitrary, but is often cell proliferation. RBE is defined as,

$$\text{RBE} \equiv \frac{D_{\text{ref}}}{D_{\text{test}}}\bigg|_{\text{isoeffect}} \quad (2.23)$$

where D_{ref} is the reference radiation dose and D_{test} is the dose of the test radiation necessary for the same biological effect [13]. Thus a more lethal radiation, requiring a lower dose, would have an RBE greater than 1.0.

For the purposes of planning radiation therapy, absorbed dose distributions are weighted by the corresponding RBE values of the distribution. Proton beams have RBE values that are relatively uniform over the length of the Bragg curve. The average RBE value of 1.1 is used in treatment planning, as suggested by the ICRU [34]. Carbon ion beams have RBE that varies over the Bragg curve as a function of depth, but the specific values depend on several factors and are not well determined [34].

Chapter 3

Computational particle therapy background

3.1 Introduction

The goal of radiation therapy is to deliver the “best” radiation dose distribution possible to the patient. In order to determine and deliver this “best” distribution, a number of physical, geometric, and biological parameters must be known and understood and used in conjunction with physical and radiobiological models. An indispensable tool in this sequence is that of computational methods.

Computational methods allow for the *in silico* estimation and optimization of dose distributions with equal or, at times, even superior accuracy than what can be practically measured in clinical settings. Because of this, computing is a cornerstone in the sequence of radiation therapy delivery, and nowhere more so than in charged particle therapy.

This section will provide background information on some of the fundamental computational techniques used in particle therapy to calculate quantities of interest.

AAPM Task Group Report No. 105 [35] is a very well referenced guide to many techniques, in particular those related to Monte Carlo.

3.2 The Monte Carlo method

The Monte Carlo method was first developed in 1946 by Stanisław Ulam and John von Neumann for use in neutron shielding calculations at the then Los Alamos Scientific Laboratory [36]. Neutron calculations were not solvable analytically, despite relatively well characterized parameters, due to the stochastic nature of neutron interactions with media. The Monte Carlo method was designed to capture this stochastic nature and yield reliable results even in physically complex scenarios.

3.2.1 Monte Carlo fundamentals

The Monte Carlo method is a statistical computing technique that uses large numbers of independent, stochastic trials to converge on the solution of a problem [37]. This technique is particularly useful for complex, many dimensional problems, which are not adequately described in closed form and are thus intractable by analytical means [38].

An elementary example of the Monte Carlo method often presented in introductory courses is that of estimating π [39]. By inscribing a unit circle inside a square, randomly tossing “darts” into the square, and counting the number that land within the circle, the value of π can be estimated using the formulas for the area of the circle, $A_{\circ} = \pi r^2$, and the square, $A_{\square} = l^2 = (2r)^2$ (see Fig. 3.1).

$$\frac{\pi}{4} = I = \int_0^1 \sqrt{1-x^2} dx \quad (3.1)$$

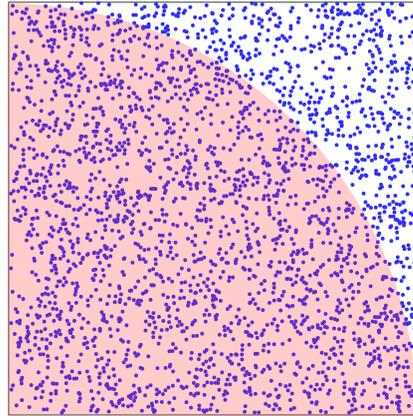


Figure 3.1: Monte Carlo estimation of π by randomly placing points in a quadrant of the square and determining the proportion which fall within the inscribed unit circle.

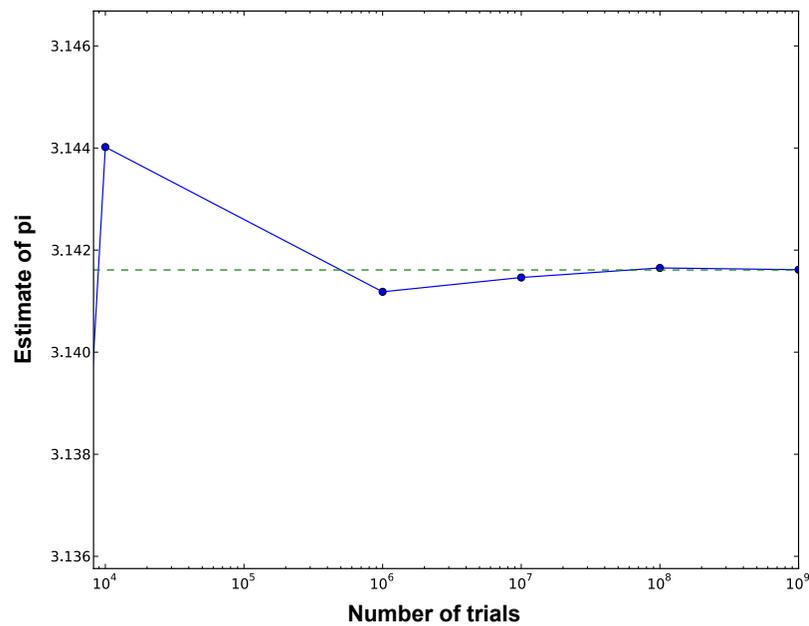


Figure 3.2: Asymptotic approach to the “true” value of π as the number of trials increases.

In effect this process numerically solves the integral in Eq. 3.1. The accuracy of the solution is a function of the number of trials performed (see Fig. 3.2).

For a general Monte Carlo calculation, steps taken in the event simulation are governed by probability distributions. Each trial randomly samples from these governing distributions to create an event “trajectory”. At the conclusion of all trials, the total results are aggregated to produce a final answer for the quantity of interest [38].

Key features of the Monte Carlo method are:

- Final outcome is product of many individual, stochastic trials.
- Inherent parallelizability due to the independence of each trial (i.e. no cross-trial interaction).
- Need for random numbers for probability distribution sampling and thus reliance on pseudo-random number generators.
- Potential for asymptotic accuracy, given adequate computing resources.

3.2.2 Radiation transport via Monte Carlo

Radiation transport is simulated using the Monte Carlo method by characterizing individual radiation interactions as probability distribution functions. Each particle trajectory is the result of many interactions, in which the outcome is determined by randomly sampling from the appropriate probability distribution function. The probability distribution functions are built from fundamental physics models and empirically obtained interaction cross sections [38].

The Monte Carlo method has been widely applied in radiation therapy physics and is considered the gold standard for dose calculations [40, 41]. Due to Monte

Carlo's stochastic nature and its treatment of particle tracks individually, Monte Carlo is especially adept at dealing with situations of highly heterogeneous media, which are often encountered in radiation therapy [35]. For radiation therapy related calculations, several Monte Carlo packages are in use, including: EGSnrc [42], FLUKA [43, 44, 45], Geant4 [46, 47, 48], MCNP [49], MCNPX [50], Shield-HIT [51], and Penelope [52].

Solving the Boltzmann equation

Radiation transport via Monte Carlo is a technique to numerically solve the Boltzmann radiation transport equation (Eq. 3.2) [38].

$$\phi(\mathbf{r}, \mathbf{v}) = \int \left[\int \phi(\mathbf{r}', \mathbf{v}') C(\mathbf{v}' \rightarrow \mathbf{v}, \mathbf{r}') d\mathbf{v}' + Q(\mathbf{r}', \mathbf{v}) \right] T(\mathbf{r}' \rightarrow \mathbf{r}, \mathbf{v}) d\mathbf{r}' \quad (3.2)$$

$\phi(\mathbf{r}, \mathbf{v})$ - particle flux density $C(\mathbf{v}' \rightarrow \mathbf{v}, \mathbf{r}')$ - collision kernel

$Q(\mathbf{r}', \mathbf{v})$ - source term $T(\mathbf{r}' \rightarrow \mathbf{r}, \mathbf{v})$ - transport kernel

\mathbf{r} - particle location \mathbf{v} - particle velocity

More intuitively, radiation Monte Carlo can be thought of as simply simulating the physics processes of many independent particles and arriving at a mean value for quantities of interest.

3.3 Monte Carlo performance enhancements and derivative techniques

While Monte Carlo has the potential to be extremely accurate, it is relatively resource intensive and often prohibitively so. Because of this, many methods have

been developed to improve the performance of Monte Carlo to make it more suitable for applications such as clinical dose calculations [35]. These techniques fall into two broad categories: approximation efficiency improvement techniques and variance reduction techniques [53].

3.3.1 Metric of simulation efficiency

A common metric for efficiency of Monte Carlo calculations is given by Eq. 3.3

$$\epsilon = \frac{1}{\sigma^2 T} \quad (3.3)$$

where ϵ is efficiency, σ^2 is the estimated variance of the quantity of interest, and T is the time to simulate the quantity with variance σ^2 [35, 54].

3.3.2 Transport cutoffs

Transport cutoffs are approximation techniques designed to discard particles which meet certain criteria in order to decrease calculation time per trial [35, 53].

Range rejection

Range rejection is a technique in which a particle is discarded if its residual range is not sufficient to travel beyond the current region into the region of interest or if the particle has started beyond the region of interest [53].

Energy cuts

Energy transport thresholds, or *energy cuts*, are generally high-pass criteria, below which particles are discarded [35]. The residual energy may or may not be deposited

locally.

3.3.3 The condensed history method

Condensed history methods are approximation techniques used for the transport of charged particles. Condensed history methods are Monte Carlo techniques which group together multiple charged particle interactions using multiple scattering theory and stopping power to enhance performance over individual track simulation [55]. The two major variants, class I and class II condensed history methods, have been widely used in radiation therapy calculations, in particular for electron-photon transport [35].

3.3.4 Variance reduction techniques

Variance reduction techniques are methods employed which lower the variance, σ^2 , of a quantity of interest while improving the efficiency of the calculation and not biasing the result [35]. Variance reduction techniques differ from approximation techniques described above, in that they do not alter the fundamental physics of the simulation. Several techniques have been used in conjunction with radiation therapy Monte Carlo calculations with varying frequency. In addition to those described below, others include simultaneous transport of particle sets, the Sobol random number method, Woodcock tracing, Macro Monte Carlo [35, 56].

3.3.5 Particle splitting

Particle splitting, most often used in the context of *bremsstrahlung splitting*, is the technique of increasing simulation statistics by cloning particles, particularly in low fluence regions or low yield secondary particles [57]. In order to maintain correct

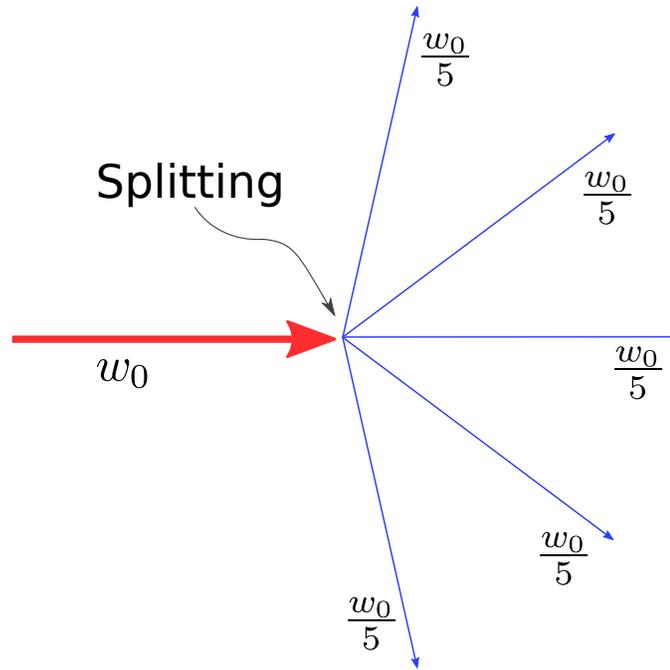


Figure 3.3: The particle splitting technique.

weighting, such that the final result is not biased toward the cloned particles, the new particles are “re-weighted” by w_0/n , where w_0 is the initial weight of the particle and n is the number of multiples created in the split.

The most basic form of splitting is called *uniform splitting*, while other forms have been developed, such as *selective brems splitting* and *directional brems splitting*, which have specific applications in efficient x-ray generation simulation [53].

3.3.6 Russian roulette

Russian roulette is a technique often used in conjunction with particle splitting [57]. Roulette consists of probabilistically terminating particles among a set of similar particles which are of little interest, whether due to their physical properties or because they are in a region of little interest. Just as with splitting, surviving

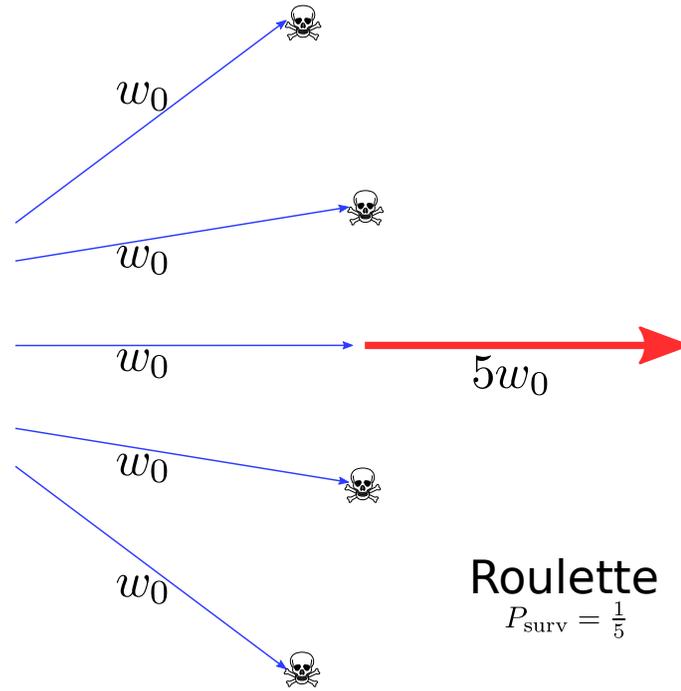


Figure 3.4: The Russian roulette technique.

particles are re-weighted to maintain statistical soundness. In the case of roulette, if a particle survives with probability P , its new weight will then be w_0/P . Roulette is fundamentally a time saving technique.

3.3.7 Interaction forcing

Interaction forcing is a technique to artificially create interactions of neutral particles, such as photons and neutrons, which interact only probabilistically, such that the mean free path (MFP) is described by Eq. 3.4 [56].

$$P(x) = \exp^{-x/\lambda} \quad (3.4)$$

where λ is the medium-dependent mean free path. In order to avoid wasting resources on the transport of particles that leave the region of interest without interacting,

particles are forced to interact within the region and re-weighted based on their initial probability of interacting. A distance is chosen, N_λ , which is the number of MFP's after which the particle will be forced to interact, based on the length of the region, Λ , in multiples of MFP's (Eq. 3.5).

$$N_\lambda = -\ln [1 - \eta(1 - \exp^{-\Lambda})] \quad (3.5)$$

The interacting particle's weight is then modified to

$$w' = w_0(1 - \exp^{-\Lambda}) \quad (3.6)$$

where w_0 is the particle's original weight.

3.4 Deterministic methods

A number of deterministic methods have been developed to calculate dose due to photon, electron, and charged hadron beams. Analytical models have been developed which have good agreement with experimental data, but have not proven flexible for treatment planning situations [58]. Another class of algorithms are kernel based algorithms, such as pencil beam algorithms, which are the current clinical standard for treatment planning. Pencil beam algorithms generally use dose kernels generated with Monte Carlo and apply various schemes to convolve and transform the kernels to estimate the final dose distribution [59, 60]. Kernel based methods tend to fail in situations with highly heterogeneous media [61]. This shortcoming is often cited as one of the major motivations for the clinical use of Monte Carlo simulations.

Chapter 4

Fast LET mapping via track repeating

4.1 Introduction

There has been recent interest in the idea of simultaneously optimizing particle dose distributions and linear energy transfer (LET) distributions. As with many radiation-related quantities, Monte Carlo techniques have been the most accurate for both dose and LET calculations, but at the cost of timely performance. In this chapter, the use of the so-called “track repeating” method is described for fast calculation of three dimensional LET maps. This method proves to be both comparably accurate and provide improved performance compared with the standard Monte Carlo method.

4.2 LET optimization

Recent work by Bassler, et al [62], Grassberger, et al [63], and Riofrío, et al [64], has examined the benefits and practicality of simultaneously optimizing treatment plans

with regards to absorbed dose and LET, in so-called “LET painting” techniques. To achieve the desired level of accuracy in LET calculations, Monte Carlo would be the preferred calculation method, but would likely lack the performance necessary for clinical time frames.

4.2.1 Oxygen enhancement effect and LET

The presence of oxygen in the environment of cells being irradiated can have a significant impact on radiobiological effect [13]. Hypoxic environments decrease the effectiveness of sparsely ionizing radiation, such as photons. The ratio of dose required for the same biological effect in the low oxygen versus normal oxygen environments is called the “oxygen enhancement ratio” (OER). OER decreases with higher LET radiation (i.e. high LET radiation is more effective) [13]. This is the basis presented by Bassler, et al, for LET painting and optimization.

4.2.2 RBE and LET

While proton therapy planning is generally performed assigning a relative biological effectiveness (RBE) of 1.1 to the entire proton dose volume, proton RBE is known to vary along the Bragg curve [34]. Although the relationship between LET and RBE for proton beams is not fully understood, there does exist a monotonic LET-RBE relationship [65]. This relationship is the basis presented by Grassberger, et al, for LET painting and optimization.

4.3 Track repeating

Track repeating, or *history repeating*, is a Monte Carlo-based radiation transport technique designed to have accuracy on par with full Monte Carlo, but with higher

performance. Initially used as a technique for electron transport in photon-electron simulations in radiation therapy physics [66, 67, 56], track repeating has more recently been adapted as a method to perform treatment planning absorbed dose calculations for charged hadron beams [68, 69, 70, 71].

4.3.1 Track repeating fundamentals

The basic idea of track repeating is to pre-calculate radiation “tracks” using full Monte Carlo for later use. More generally, the motivation is to gain performance by splitting the radiation transport simulation into two distinct parts:

1. A calculation independent of the particulars of the ultimate calculation of interest.
2. An on-demand calculation, using problem-specific parameters.

This concept is similar to Monte Carlo schemes in which the treatment hardware (e.g. particle accelerator and beam collimators) is simulated in advance to create phase space characterizations of the radiation kinematics. This information can then be generically used for patient specific calculations [72, 73, 74, 41].

The track repeating algorithm

The basic algorithm of track repeating is illustrated in three steps in Figure 4.1.

1. A database of tracks is created using a full-fledged Monte Carlo package.
2. Tracks from the database are overlaid on the scoring grid with the geometry of interest.
3. The quantity of interest is tallied.

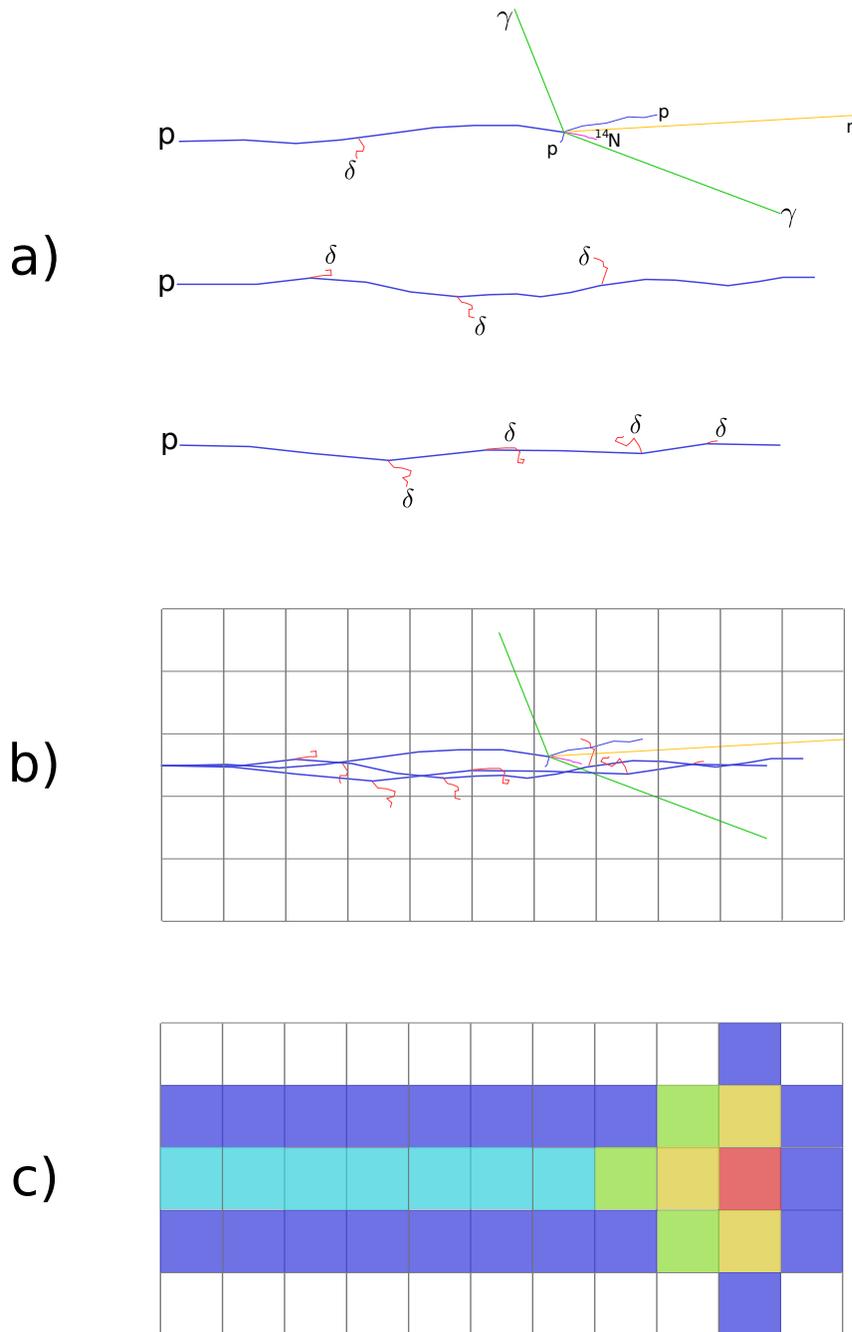


Figure 4.1: The basic track repeating concept. a) Generate a database of particle tracks using a full-fledged Monte Carlo package. b) Overlay tracks from database onto scoring grid. c) Accumulate quantity of interest.

A typical track will consist of series of “steps”. Each step contains information about the state of the particle during the step, such as kinetic energy, energy loss, energy deposited, angle of scatter, step length, and secondaries created [67, 70]. Each step is a summary of the many interactions that have taken place over the step length. Therein lies the primary performance advantage of track repeating over full Monte Carlo, as all of the individual interactions that the Monte Carlo carries out are contained within each step, summarized as simple parameters [67].

Radiation transport is performed by propagating the tracks step-by-step through a voxelized geometry. The end point of a step is determined from step data (e.g. step length and scatter angle) and the corresponding system coordinates are determined. If appropriate, energy lost by the particle is scored in the corresponding voxel. Secondary particles can be propagated in the same manner by spawning from the originating step in a primary track.

Various authors have dealt with propagation and energy deposition in media different from the generation medium by scaling the step length [67, 68, 70]. This strategy has dealt well with inhomogeneous media, but does not explicitly account for backscatter at media interfaces.

4.4 The LET track repeating algorithm

Linear energy transfer is the energy deposited by an incident particle to a medium per unit distance (see Section 2.5.6). For the case of charged hadron LET, the *unrestricted* version, L_∞ , is most appropriate.

$$L_\infty = \frac{dE_\infty}{dx} \quad (4.1)$$

To estimate LET in the voxel designated by \mathbf{x} , the *track-averaged* and *dose-averaged*

LET were calculated using Eq.'s 4.2 and 4.3.

$$L_{\text{track}}(\mathbf{x}) = \frac{\sum_i \Phi_i(\mathbf{x}) L_{\infty,i}(\mathbf{x})}{\sum_i \Phi_i(\mathbf{x})} \quad (4.2)$$

$$L_{\text{dose}}(\mathbf{x}) = \frac{\sum_i \Phi_i(\mathbf{x}) L_{\infty,i}^2(\mathbf{x})}{\sum_i \Phi_i(\mathbf{x}) L_{\infty,i}(\mathbf{x})} \quad (4.3)$$

where Φ_i is the particle fluence.

For an individual track, LET information is tallied by propagating the track inside the voxelized scoring grid (the phantom), one step at a time. At each step, the energy deposited, energy deposited divided by step length, and fluence (clearly one for a single particle) is recorded in the corresponding voxel, or bin. After all tracks have been propagated, the final LET values are then calculated using Eq.'s 4.2 and 4.3, where dose is the total energy inside a voxel divided by its mass (see Section 2.4.3).

4.4.1 Track generation with Geant4

The track databases of primary protons were generated with the Geant4.9.3 Monte Carlo package [46, 47, 48]. Geant4 is a set of C++ language libraries that allow the users full customization of the simulation, but also provides some reasonable preset physics “lists”, which set the interaction models used. The following physics lists were used to generate primary proton tracks:

- G4EmStandardPhysics_option3
- G4HadronElasticPhysics
- G4HadronInelasticQBBC
- G4IonBinaryCascadePhysics

- G4RadioactiveDecayPhysics

which are recommended for benchmarking against experimental proton depth-dose data.

In track generation each step contained the following information,

- The kinetic energy of the particle at the end of the step.
- The total energy deposited over the course of the step.
- The step length.
- The unit vector describing the direction of the step.

Each track was generated with a maximum step length of 0.1 mm and transport cut-off (high-pass) values equivalent to 1.0 mm for electrons, positrons, and photons. The tracks were generated in a 40 cm \times 40 cm \times 40 cm water volume, with the origin of the track centered on one face of the cubic volume. Each track was strictly monoenergetic at the entrance with all tracks having identical initial momenta.

Figures 4.2, 4.3, 4.4, and 4.5 show some of the properties of the tracks generated with Geant4. Figures 4.2 and 4.3 display the kinetic energy and energy loss per step of the primary protons respectively. It is apparent that energy loss first occurs in a nearly linear manner and then increases to a superlinear loss near the end of the primary range, which is coincident with the Bragg peak. While the majority of energy loss is due to Coulombic interactions with electrons, some nuclear interactions occur as well, although for 50 MeV protons, the probability of such interactions is less than 5% and lower at higher energies [68]. In Figure 4.2 dramatic energy loss by a few particles can be seen, which is due to inelastic nuclear interactions. These large energy losses are not reflected in Fig. 4.3, as the energy lost is transferred to secondary particles and not counted in the unrestricted stopping power of the primary particles.

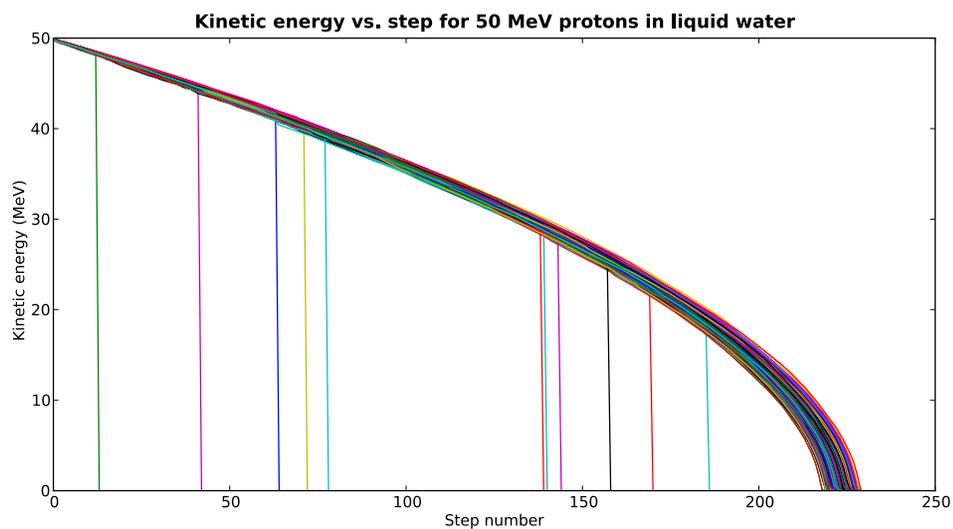


Figure 4.2: Kinetic energy versus step number for 200 proton tracks starting at 50 MeV in liquid water. Some tracks have abrupt drops to ≈ 0 MeV due to inelastic nuclear interactions.

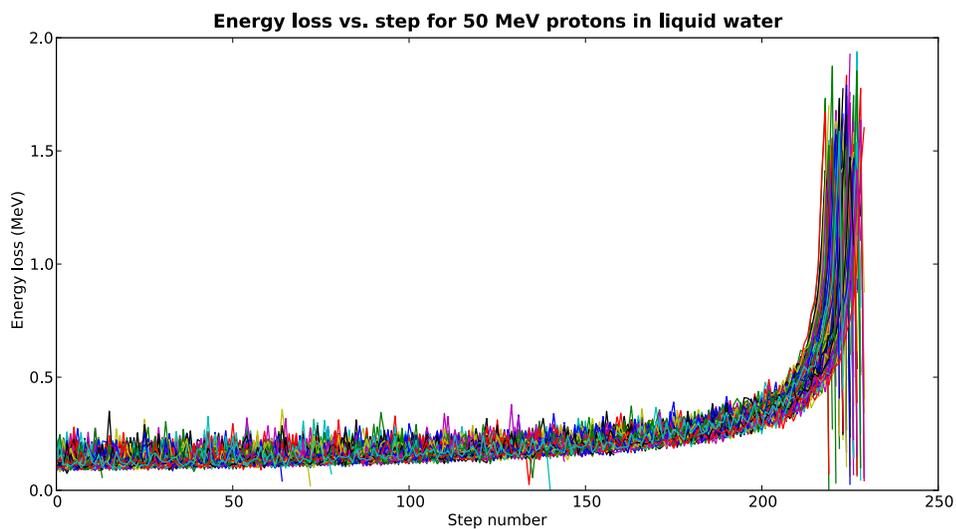


Figure 4.3: Energy loss versus step number for 200 proton tracks starting at 50 MeV in liquid water.

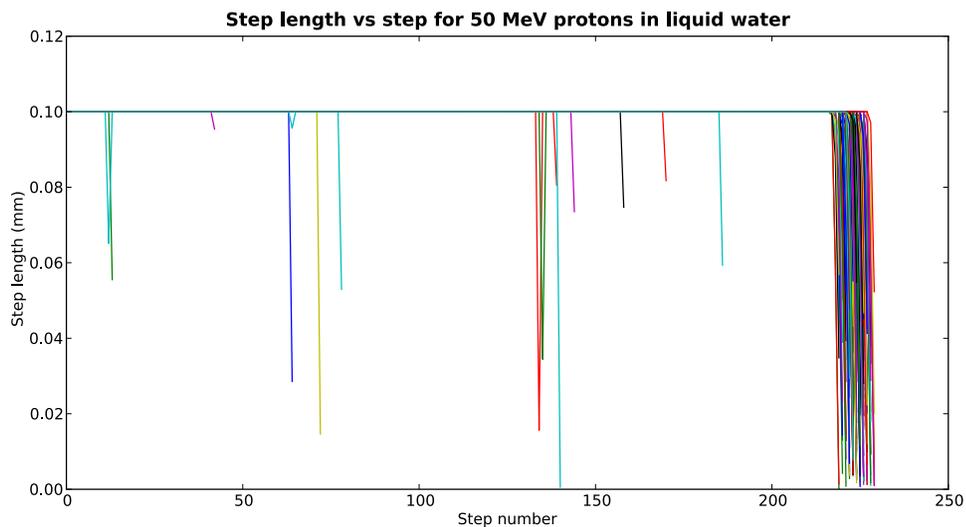


Figure 4.4: Step length versus step number for 200 proton tracks starting at 50 MeV in liquid water.

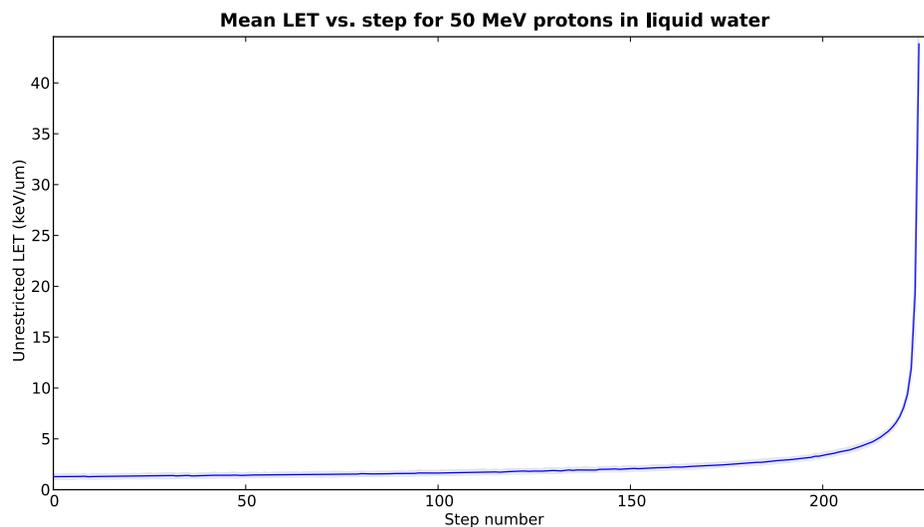


Figure 4.5: Mean LET versus step number for proton tracks starting at 50 MeV in liquid water. Only tracks with an equal numbers of steps were used to create the mean LET value. This value is very similar to LET versus depth, which requires track propagation.

Figure 4.4 displays the step length for the steps of 200 proton tracks generated with a 0.1 mm maximum step length parameter. It is clear from the figure that the step length does not greatly deviate from the enforced maximum until the kinetic energy of the proton has significantly decreased. At this point, the interaction cross sections have increased and with each subsequent interaction the probability of another collision further increases.

Figure 4.5 displays the mean unrestricted LET versus step for a set of approximately 225 50 MeV proton tracks. This LET value is very similar to LET versus depth since protons deviate very little from their initial trajectory. To calculate the true LET vs depth, the tracks must be propagated as described Sec. 4.3.1.

4.5 Benchmarking and performance

4.5.1 Benchmarks

A first benchmark was performed by comparing LET values for a simple geometry calculated with the track repeating algorithm against LET derived from well established stopping power values.

The track repeating values were calculated with a 50 MeV proton beam with a symmetric cross sectional Gaussian σ of 2 mm. The beam consisted of 10,000 tracks repeated 10X into a 40 mm \times 81 mm \times 81 mm water volume with a voxel size of 0.2 mm \times 1.0 mm \times 1.0 mm, with 40 mm being the dimension along the beam axis. The LET-depth curve was then calculated using Eq. 4.2 along the central voxel column, which should have been subject to charged particle equilibrium.

The benchmark LET-depth curve was generated with data from NIST's online *pstar* tool [6], which uses data that forms the bulk of ICRU Report 49 [75]. The curve was generated by iteratively looking up the stopping power for the current kinetic

energy of the proton, then subtracting the corresponding energy for the distance over an arbitrary step length.

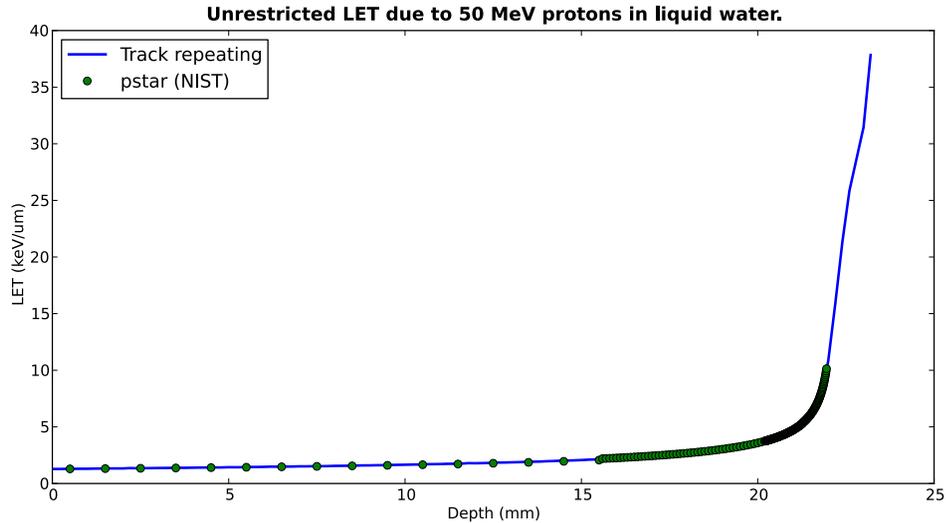


Figure 4.6: Unrestricted LET for 50 MeV protons in liquid water generated with track repeating and generated from data from NIST’s *pstar* tool [6].

Figure 4.6 shows a comparison of unrestricted LET of 50 MeV protons generated with the track repeating algorithm and the benchmark data derived from *pstar*. The striking similarity in values is likely due to Geant4 using the same underlying energy loss models and cross section data and models directly from ICRU 49 in part of the energy spectrum (2 MeV - 1 keV) [71]¹. While this results in a somewhat circular test, with Geant4 serving as an intermediary “black box”, it demonstrates that the values are largely preserved by the track repeating process.

A second comparison was of 3D LET distributions produced by full Monte Carlo with Geant4 and the track repeating algorithm. This comparison again used 50 MeV proton beams in liquid water. The beam consisted of 400k primary protons in Geant4 and 40k primary protons in track repeating, repeated 10 times. The cross sectional

¹NIST provides a detailed description of the models used to generate the stopping power data on its website [6].

Gaussian σ of the beam was 5 mm.

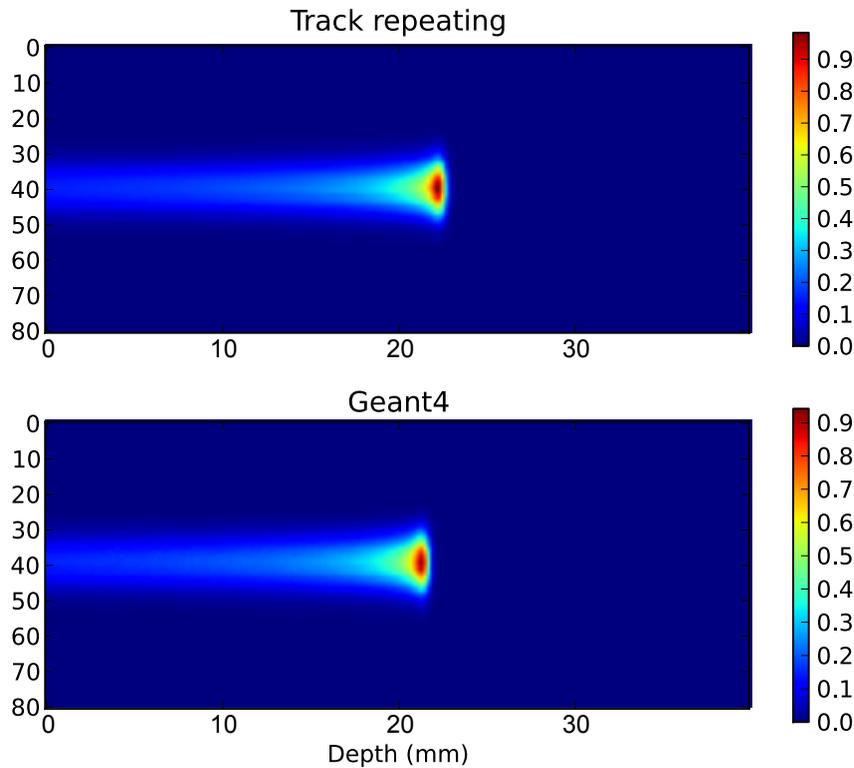


Figure 4.7: Absorbed dose of a 50 MeV proton beam in liquid water calculated with track repeating and Geant4.

Figures 4.7 and 4.8 show two dimensional cross sections of the dose and track-averaged LET distributions from 50 MeV protons generated with track repeating and Geant4. The distributions are substantially similar, especially with regards to range of the beam. Figure 4.9 shows the central voxel column of the same data, displaying the variation of the LET with depth. The values from track repeating match up well with Geant4.

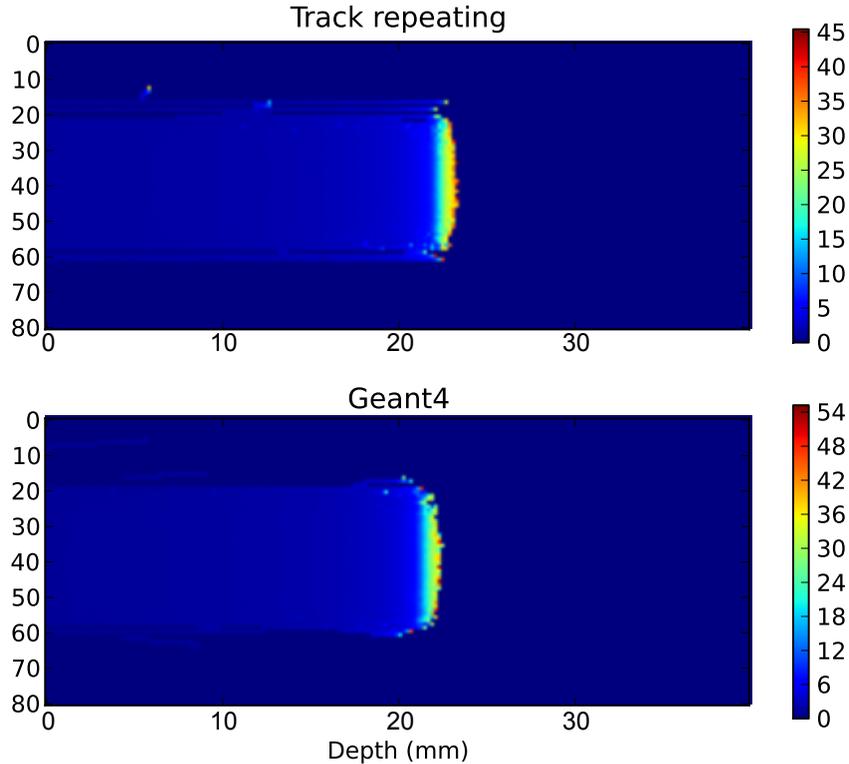


Figure 4.8: Track-averaged LET of a 50 MeV proton beam in liquid water calculated with track repeating and Geant4.

4.5.2 Performance

The track repeating algorithm was implemented in both the interpreted Python language [76] and the compiled C++ language [77, 78]. Track repeating performance was measured against a full Monte Carlo calculation with Geant4 of track-averaged LET for a 50 MeV proton beam in liquid water. The beam had a symmetric cross sectional Gaussian σ of 5 mm. Performance results are summarized in Table 4.1.

Even with the inherent slowness of Python, due to the language’s dynamic typing and interpreted nature, the LET calculation came out to 4.4 times as fast as the native Geant4 calculation. The C++ implementation was faster still at 15.2 times

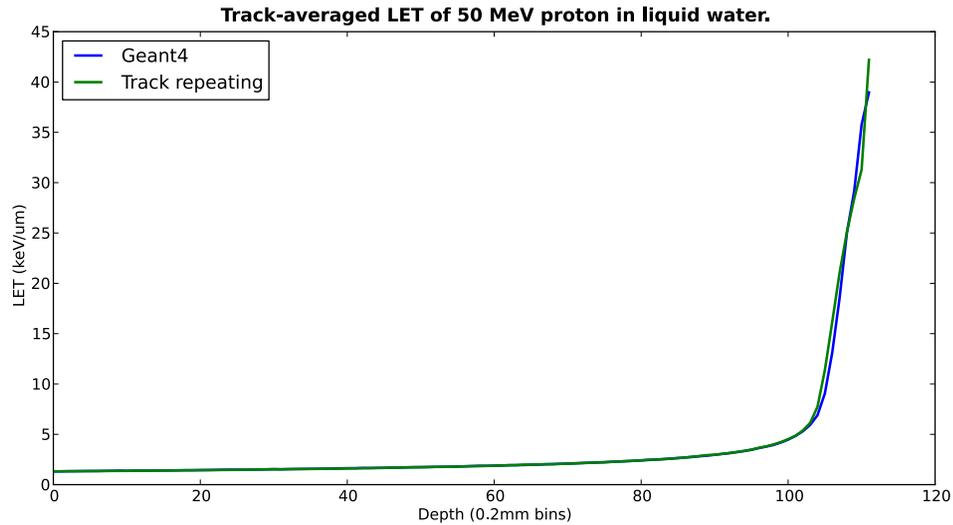


Figure 4.9: Track-averaged LET versus depth of a 50 MeV proton beam in liquid water calculated with track repeating and Geant4.

Algorithm	Geant4	Python TR	C++ TR
Mean time (sec)	207.0	46.6	13.6
Performance ratio	1.0	4.4	15.2

Table 4.1: Performance of LET calculations for 50 MeV proton beams in liquid water.

the calculation speed of Geant4.

4.6 Conclusion and outlook

The track repeating algorithm is designed to break up time intensive Monte Carlo simulation into two separate parts: a pre-processing, simulation independent part, where most of the “heavy lifting” is performed by full Monte Carlo, and a second “on-demand” part running calculation specific parameters. For dose calculations presented by other authors, this has proven to be both faster than Monte Carlo and highly accurate.

For the LET calculation presented here, the accuracy and performance promise of track repeating seems to have delivered. The fastest speed up over full Monte Carlo was the C++ version of the track repeating code, with a performance improvement of $\approx 15\times$ over Geant4. This performance is not as impressive as claims made in the literature of improvements up to $200\text{-}500\times$ [70], but we suspect that further improvements can be made via code optimizations.

This LET calculation method is a relatively straight forward novel use of the track repeating algorithm. In the following chapter, Chapter 5, track repeating is used to estimate peripheral dose due to particle beams. This is a significantly more complex simulation, for which the LET track repeating algorithm is a stepping-off point.

The current algorithm is somewhat limited, only calculating LET for primary particles and not adjusting for heterogeneous media. Using techniques from the next chapter, LET contributions from secondary particles could be taken into account. Transport through heterogeneous materials could be implemented using techniques described in the literature [67, 68, 70].

Acknowledgements

This project was greatly aided by discussions with Niels Bassler about LET and stopping power and Pablo Yepes and Iavor Veltchev regarding the track repeating algorithm. Shuang Luan wrote most of the Geant4 code used to generate LET.

Chapter 5

Fast secondary dose estimation via track repeating

5.1 Introduction

Of known radiation types, charged particle beams have the most desirable radiation dose distributions. Despite this, some scattered radiation is nonetheless present outside of the target volume. This peripheral, or secondary dose, is of concern not only because of the risk of acute effects, such as those which cause tissue death, but also risks such as radiogenic cancer. Accurately calculating the secondary dose in regions far from the target volume is a resource intensive task, due to low particle fluence and the effect of heterogeneities on particle ranges. In this chapter, the use of the so-called “track repeating” method is described for fast secondary dose estimation.

5.2 Secondary dose and second cancer in particle therapy

As with any medical therapy, side effects are a concern with charged hadron therapy. Compared with external photon therapy, charged hadrons have the potential for much better tumor dose conformity with lower dose to surrounding tissues. This brings the promise of using higher tumor doses as well as treating previously untreatable sites, such as many pediatric tumors, which were considered too risky due to potential long-term side effects. Both of these scenarios require careful consideration of potential long term side effects due to low dose exposure outside of the treatment region, since hadrons are fundamentally different than other therapies. This low dose exposure is termed secondary or peripheral dose. Secondary dose is of concern because it can lead to long term effects including radiation induced cancer and developmental problems in children [79, 13, 80]. Recently, the concern about secondary dose in hadron therapy was raised by Hall, in particular regarding dose from secondary neutrons in proton therapy [81]. This has led to a number of studies looking at the potential secondary dose from charged hadron therapy and x-ray therapy and attempting to quantify the associated cancer risk [82, 83, 84, 85].

5.2.1 Sources of secondary dose

Dose delivered to the target volume is referred to in this dissertation as primary dose, while dose that falls outside of the target volume is termed secondary dose. In charged hadron therapy, secondary dose is ultimately due to interactions of the incident particle with nuclei of the medium (i.e. tissue) [31]. Protons, for example, will collide with nuclei, knocking out neutrons, fragmenting nuclei into heavy ions, and scattering primary protons. Neutrons are especially noteworthy, as they can penetrate throughout the body. Fragmentation of medium nuclei can also produce

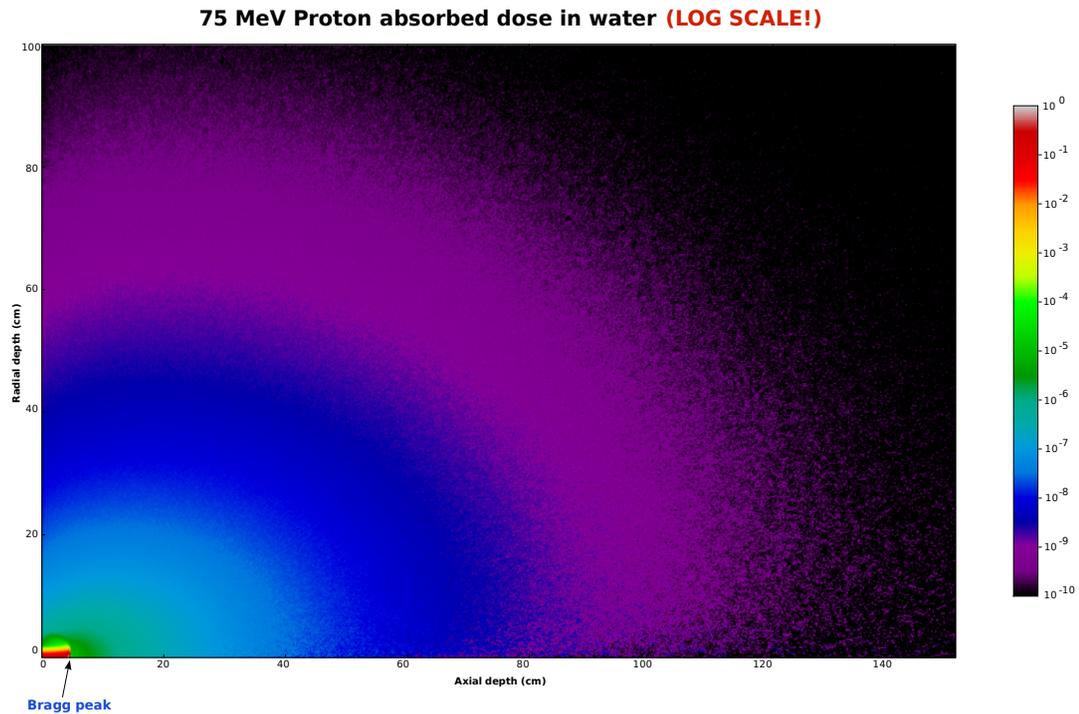


Figure 5.1: Absorbed dose in the peripheral region from 75 MeV protons in liquid water (**log scale**). Data generated with FLUKA.

long ranging gamma rays. Figure 5.1 shows the absorbed dose due to a 75 MeV proton beam plotted on a log scale. The dose extends beyond the 150 cm range shown in the figure, although at an even lower fraction of the maximum dose than what can be seen in the plot, which is normalized to one in the figure. If the incident beam is composed of compound nuclei, such as carbon-12, the primary particles themselves can fragment, leading to the distinctive forward-peaked dose tail seen in beams heavier than protons [32].

5.3 Estimating secondary dose

Generally, the published studies related to secondary dose in charged hadron therapy fall into three categories: disease specific estimates and case studies of individual patients, peripheral field estimates and measurements, and schemes to estimate patient-specific secondary dose [82, 86, 87, 88, 89, 90, 91, 92]. In the first two categories, the goals of the studies have been primarily to determine the relative safety of hadron therapy on the broad level and reassure the community of radiation therapy practitioners and potential patients. The third category, patient-specific estimate schemes, aims for more clinical relevance.

In nearly all of the published studies, calculations are typically performed by simulating a patient treatment using a Monte Carlo transport code such as Geant4 [46, 47, 48] or MCNPX [50] to estimate the secondary dose. This is considered the best method to produce accurate dose values, particularly after the particles have traversed media of various densities and potentially produced many secondary particles. The great drawback of most Monte Carlo implementations is the calculation time. A treatment region may be taking into consideration a volume as large as $\sim 30\text{k cm}^3$, whereas the volume included in a secondary dose calculation might be closer to $\sim 125\text{k cm}^3$. Thus, in order to produce statistics with acceptably low associated errors, runs of several days to weeks of CPU time are common [41, 93]. While this is acceptable for producing safety related results, these long times do not allow clinicians to have patient-specific information regarding secondary dose on which to act.

A previous study by the author investigated estimating secondary dose from proton, antiproton, and carbon ion beams using an analytical model fit to simulated data [94]. The advantage of this model was very fast calculation time once the initial Monte Carlo simulations were performed. The disadvantage was that the model only gave broad estimates and was unable to take geometric specifics and inhomogeneous

media into account.

The motivation behind a new track repeating based method of secondary dose estimation (SDE) is to provide estimates that are accurate while being fast enough that clinicians could use the information as the basis of treatment decisions. Current Monte Carlo based schemes, with time scales mentioned above, are not suitable for use in clinical settings.

5.4 SDE via track repeating

The basic track repeating algorithm is presented in Sec 4.3.1. For SDE the basic algorithm must be extended to take into account the many secondary particles produced, which then contribute to secondary dose. Additionally other speed up techniques are applied to reduce run time while maintaining overall accuracy (see Sec 5.5).

The SDE track repeating algorithm is similar to the basic one described in Fig 4.1, except that each track step contains additional information about secondaries produced. To propagate a secondary of an arbitrary energy, the algorithm must first identify the type of particle, then select the appropriate database. Because the secondary is very unlikely to have the same initial energy as the database tracks, a sub-section of a database track must be extracted with a maximum energy approximately equal to the initial energy of the secondary particle. To perform this selection, a binary search algorithm was implemented, which has a worst case performance of $O(\log_2 n)$, where n is the number of steps in the track [95]. The binary search was modified to have a “fitness” parameter, Δ , such that if the energy found were not within the fraction Δ of the sought after energy, the search moved to the next track. Δ was generally set to 0.1, which seemed to give the best performance (see Fig. 5.2, which displays run times for 20k primary protons with different Δ values). When choosing a sub-section of a track that meets the Δ criterion, the worst case perfor-

mance is then $O(m \log_2 n)$, where n is the number of steps in the track and m is the number of tracks in the database.

The track selected from the database of secondaries must also be aligned with the trajectory of the secondary before propagation. This is accomplished using a Householder transformation, which efficiently reflects the first step of the new track into alignment with the secondary's unit vector while rigidly transforming the remainder of the track [96]. This process is more efficient than using a similar rotational matrix procedure to align the secondary tracks.

The SDE algorithm was implemented in the interpreted Python language [76] using the numerical library NumPy [97].

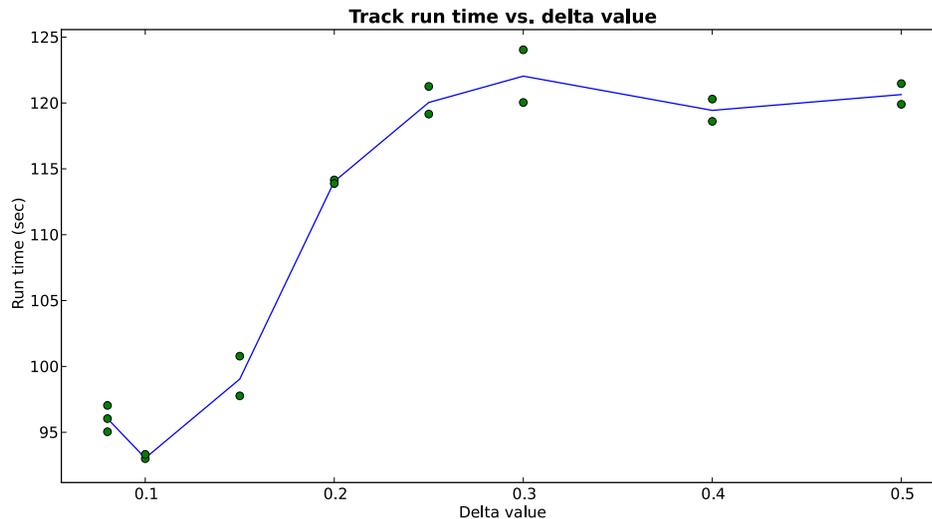


Figure 5.2: Run times for 20k primary protons using different Δ values in the binary search algorithm. Δ values below 0.08 caused recursion errors, since the algorithm had to continue to search for steps with energies that met such a strict criterion.

5.4.1 Databases

For SDE multiple track databases are necessary in addition to the database of primary particles. For proton SDE, additional databases are needed for electrons, positrons, neutrons, and photons. Heavy ions produced from nuclear fragmentation were not transported due to their low ranges (see Sec. 5.5.3).

The track databases were generated with the Geant4.9.3 Monte Carlo package [46, 47, 48]. Geant4 is a set of C++ language libraries that allow the users full customization of the simulation, but also provides some reasonable preset physics “lists”, which set the interaction models used. The following physics lists were used to generate all track databases:

- G4EmStandardPhysics_option3
- G4HadronElasticPhysics
- G4HadronInelasticQBBC
- G4RadioactiveDecayPhysics

which are recommended for benchmarking against experimental proton depth-dose data.

In track generation, each step contained the following information,

- The kinetic energy of the particle at the end of the step.
- The total energy deposited over the course of the step.
- The step length.
- The unit vector describing the direction of the step.
- A list of secondaries generated in the step, with the following information,

- The secondary particle type.
- The kinetic energy of secondary.
- The unit vector describing the initial trajectory of the secondary.

Each track was generated with a maximum step length of 0.1 mm and transport cut-off (high-pass) values equivalent to 10.0 mm for electrons and photons. The tracks were generated in a 40 cm \times 40 cm \times 40 cm water volume, with the origin of the track centered on one face of the cubic volume. Each track was strictly monoenergetic at the entrance with all tracks having identical initial momenta.

Proton database

Proton databases of several energies were generated as part of the investigation. The results presented in Sec. 5.6 were calculated with a primary proton database of 120 MeV protons. The secondary databases were generated based on statistics derived from a database of 250 MeV protons, which is approximately the maximum kinetic energy necessary for proton therapy. An unprocessed 120 MeV database of 20k protons was approximately 196 MB stored in an uncompressed flat file format (\approx 10KB per track).

Electron database

An initial energy of 31 MeV was selected for electrons, as no electrons with higher kinetic energies were observed to be produced in any subsequent reactions due to primary protons of 250 MeV. An unprocessed 31 MeV database of 2k electrons was approximately 26 MB stored in an uncompressed flat file format (\approx 13.3KB per track).

Neutron databases

Unlike charged particles, which are constantly losing energy as they pass through matter, neutral particles interact infrequently. This means that a neutron (above a few keV) will tend to have a few interactions resulting in major energy losses or potentially not interact at all (see Fig. 5.3)[29]. Because of this, the binary search method described above fares poorly for neutrons and photons, when compared with charged particles. Protons, electrons, and positrons tend to lose energy in a “smooth” manner, whereas neutrons and photon tend to have large step-downs. Using the binary search with a single initial energy database can then lead to secondaries that either do not have correct energies, or very long search times.

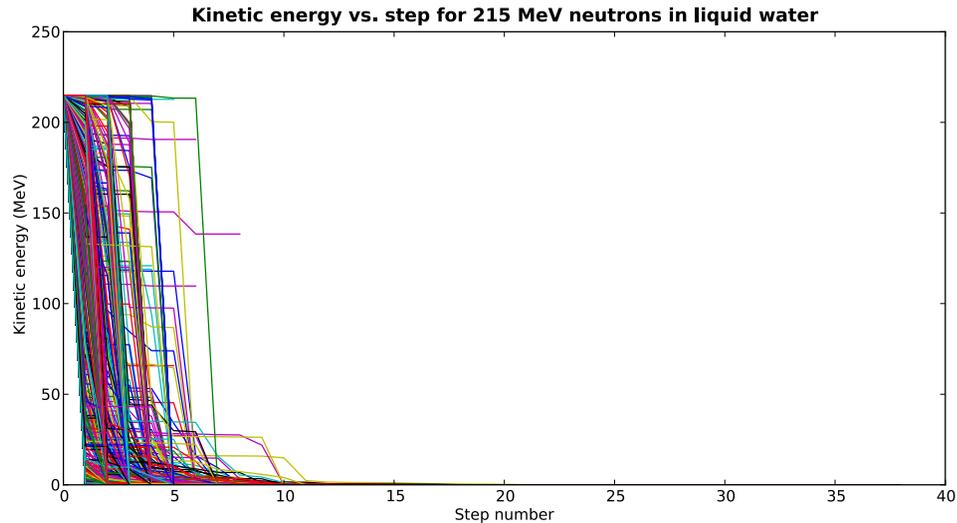


Figure 5.3: Kinetic energy versus step number for 215 MeV neutrons in liquid water phantom of 400 cm sides.

To overcome this issue databases of neutral secondary particles are broken up into several smaller databases. The initial energies are set such that the energies of secondaries are well covered based on the energy spectrum of observed secondaries. This is achieved by taking the frequency plot of secondary kinetic energies and ad-

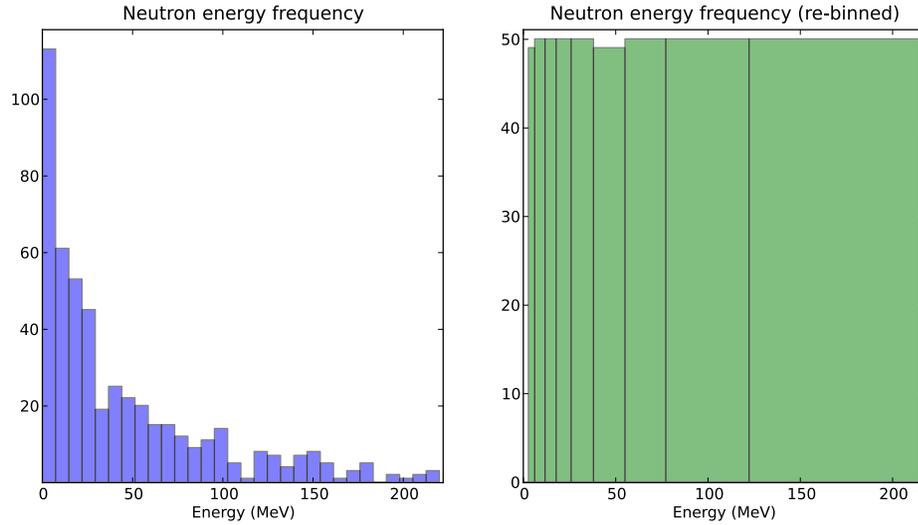


Figure 5.4: Frequency histograms of secondary neutrons produced by 250 MeV primary protons. The plot on the right has been re-binned, such that each bin has approximately equal numbers of neutrons.

justing the bin widths so that all bins have equal heights (see Fig. 5.4). The resulting energies for the ten neutron databases are found in Table 5.1.

Database	1	2	3	4	5	6	7	8	9	10
Energy (MeV)	3	6	12	18	26	38	55	77	122	215
Database size (MB)	11	9.4	7.6	6.7	5.6	4.3	3.6	3.2	2.8	2.5
Storage per track (kB)	10.5	9.6	7.7	6.8	5.6	4.3	3.6	3.2	2.8	2.5

Table 5.1: Energies and storage sizes of neutron track databases (10k particles each).

For energies listed above, there are approximately the same number of neutrons produced with initial kinetic energies in each range (i.e. there are approximately the same number with energies between 3-6 MeV as there are between 122-215 MeV). To select the appropriate track for a given neutron, a binary search is first performed to select a database, then the standard track search is performed. Each database contains 10k tracks.

Photon databases

Similar to neutrons, ten photon databases were generated to enable efficient secondary track searching. The database energies are listed in Table 5.2. Each database contains 10k tracks.

Database	1	2	3	4	5	6	7	8	9	10
Energy (MeV)	1.4	2.2	2.4	2.6	3.9	4.2	5.2	6.5	8.3	30.3
Database size (MB)	6.0	5.4	5.3	5.1	4.5	4.3	4.0	3.7	3.4	2.2
Storage per track (kB)	6.1	5.5	5.4	5.2	4.6	4.3	4.1	3.7	3.4	2.2

Table 5.2: Energies and storage sizes of photon track databases (10k particles each).

Positron database

An initial energy of 30 MeV was selected for positrons, as no positrons with higher kinetic energies were observed to be produced in any subsequent reactions due to primary protons of 250 MeV. An unprocessed 30 MeV database of 2k positrons was approximately 26 MB stored in an uncompressed flat file format (≈ 13.2 KB per track).

5.5 Performance improvement techniques

Several techniques were implemented beyond the standard track repeating algorithm in an attempt to improve the efficiency of the simulation.

5.5.1 Database filtering

One of the primary techniques was database filtering. The databases were filtered to discard tracks that would not contribute to secondary dose and would thus simply

add to calculation time without improving the result. This is somewhat analogous to the range rejection technique mentioned in Sec. 3.3.2. A key difference is that all filtering is performed in pre-processing, rather than at runtime, further enhancing performance over range rejection style techniques.

For neutron and photon databases, all tracks that did not interact within the generation phantom (40 cm × 40 cm × 40 cm) were filtered out. For high energies, this meant the majority of tracks were discarded.

For primary protons, two filtering schemes were tested. “Filter 1” discarded all tracks that did not produce any secondary particles. “Filter 2” discarded all tracks that did not produce either at least one neutron or gamma (i.e. the long range secondaries). “Filter 0” was used to mean no filter on primaries. A separate secondary proton database was used that was unfiltered.

5.5.2 Secondary splitting

Because secondary particle production tends to be rare¹, splitting was implemented for neutrons and gammas, which are the particles most likely to generate secondary dose events in the peripheral regions. See Sec. 3.3.5 for information on splitting. Each particle was split 10-fold, with a corresponding reduction of weight by $\frac{1}{10}$.

5.5.3 Transport cuts

Beyond the high-pass transport cut corresponding to 10.0 mm for electrons and photons used in generating the tracks, other cuts were used to reduce simulation time. The high-pass energy cuts are listed in Table 5.3. When a particle was discarded due to a cut, its energy was then deposited locally.

¹As mentioned in Chp. 4, 50 MeV protons have $\approx 5\%$ chance of a nuclear interaction, a probability which only decreases with kinetic energy

Particle	proton	neutron	electron	positron	photon
Cut energy (MeV)	5.0	0	0.35	0	0.1
Approx. range (mm)	0.04	-	1.1	-	200

Table 5.3: High-pass transport cut values for track repeating. Range sources [6, 7].

5.6 Benchmarking and performance

Benchmarking was performed by calculating absorbed dose in the peripheral region due to a 120 MeV proton beam, as calculated by track repeating and Monte Carlo using Geant4 and FLUKA [43, 44, 45]. A beam with a symmetric cross sectional Gaussian σ of 5 mm was incident on a water phantom of dimensions 8.1 cm \times 8.1 cm \times 40 cm. The scoring grid had voxels of 1 mm³.

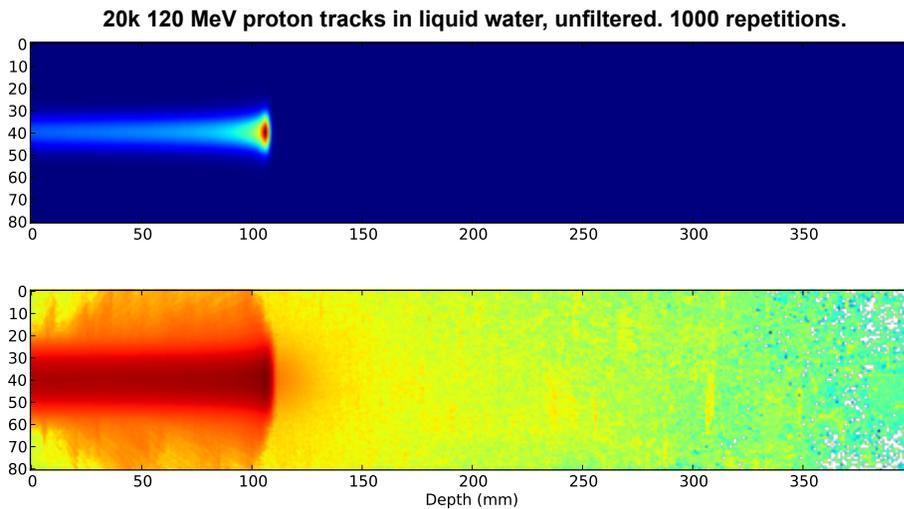


Figure 5.5: Absorbed dose from 20,000 unfiltered 120 MeV proton tracks repeated 1,000 times. Second plot is log scale.

Figures 5.5, 5.6, and 5.7 show the results of 20k 120 MeV tracks propagated with track repeating and repeated 1000 times. Figure 5.5 displays the dose of the primaries without any filtering, whereas Fig. 5.6 displays the dose from a simulation

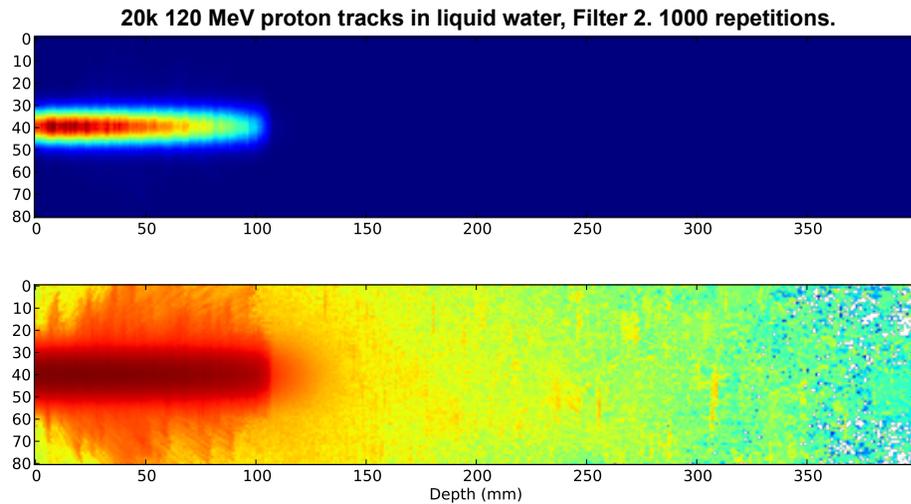


Figure 5.6: Absorbed dose from 20,000 filtered 120 MeV proton tracks repeated 1,000 times, using Filter 2. Second plot is log scale.

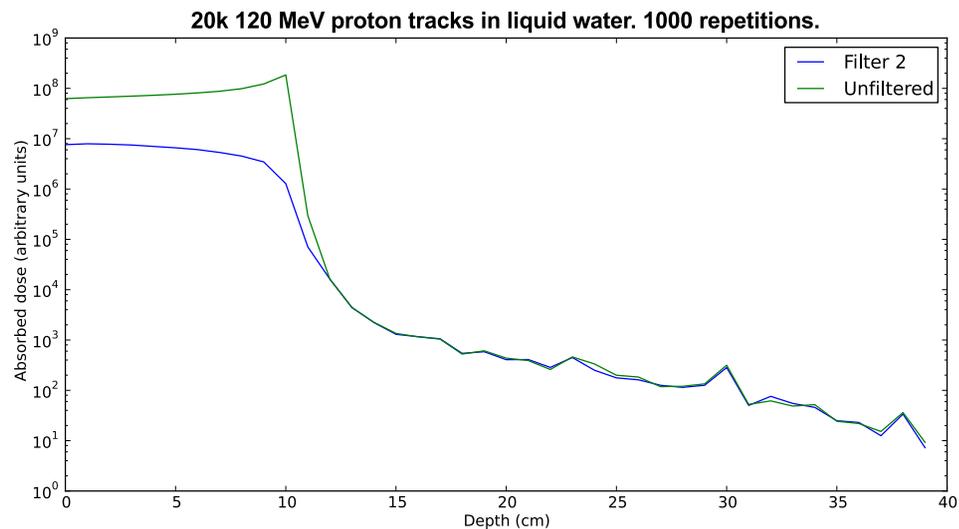


Figure 5.7: Depth-dose curves from 20,000 filtered and unfiltered 120 MeV proton tracks repeated 1,000 times (**log scale**).

using *Filter 2*, which discards primary tracks that do not produce either neutrons or photons (i.e. long range secondaries). While it is clear that the primary dose with

Filter 2 is very different than dose from the unfiltered tracks, Fig. 5.7 shows that dose from *Filter 2* and *Filter 0* is essentially identical in the peripheral region (in this case along the beam axis).

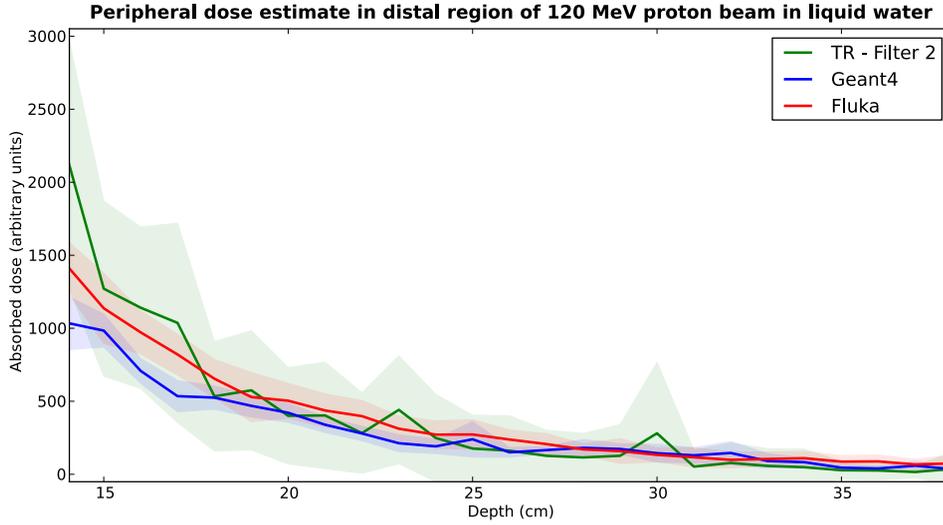


Figure 5.8: Depth-dose curves in the peripheral region from 120 MeV protons in liquid water. Curves generated by FLUKA, Geant4, and track repeating. Colored bands represent $\pm\sigma$ uncertainty.

Figures 5.8 and 5.9 show absorbed dose in the peripheral region due to 120 MeV protons calculated with track repeating (*Filter 2*), Geant4 and FLUKA. The secondary doses of Geant4 and FLUKA both fall within the error band of the track repeating data. The track repeating data was the result of 20k filtered tracks repeated $1000\times$, while the Geant4 and FLUKA data were there result of 55 million primaries each. The doses were scaled to assume 20M primary particles.

For the equivalent of 20M primaries the track repeating algorithm with *Filter 2* took ≈ 23.6 hours, Geant4 took ≈ 88.2 hours, and FLUKA took ≈ 8.1 hours on a single core of a 2007 quad-core Intel server processor at 3.16 GHz. Based on a rough estimate from the performance ratio of the C++ version of the LET track repeating algorithm to the Python version (≈ 3.4), a C++ port of the SDE track

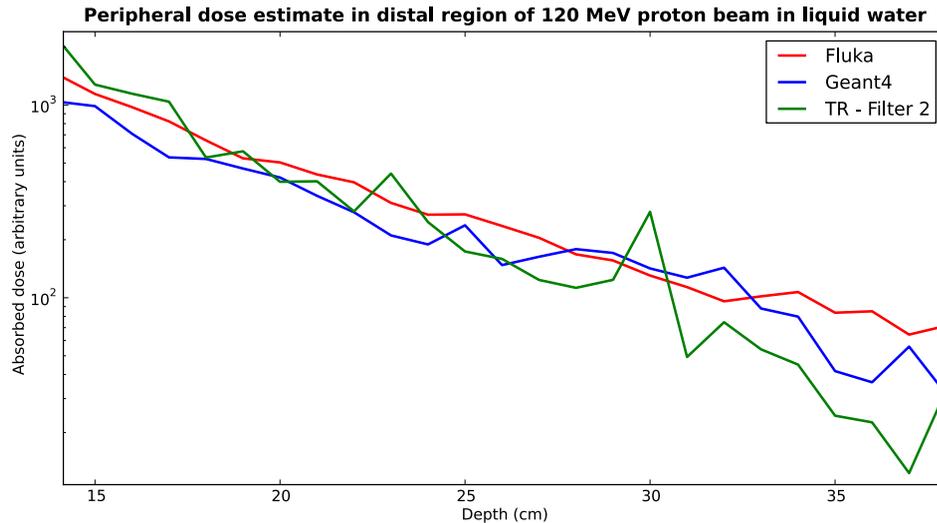


Figure 5.9: Depth-dose curves in the peripheral region from 120 MeV protons in liquid water. Curves generated by FLUKA, Geant4, and track repeating (**log scale**).

repeating algorithm is likely to reduce the above time from 23.6 hours to ≈ 6.5 -7.0 hours, making track repeating ≈ 1.1 - $1.3\times$ faster than FLUKA without further optimizations.

5.7 Conclusions and outlook

Estimating the secondary dose from charged hadron beams is a difficult problem because the volumes of interest are large and the particle fluences are low. Using Monte Carlo means very long run times for reasonable uncertainty levels. The track repeating algorithm presented here is an attempt to improve upon those run times while maintaining the accuracy of Monte Carlo.

The results of this study indicate that track repeating can achieve this goal to a degree. The accuracy of the dose estimates is reasonable and run times are reduced. Furthermore, track repeating can also take advantage of parallel processing, which

is largely the topic of the next chapter (Chap. 6).

The current algorithm is somewhat limited, only calculating SDE due to protons and not adjusting for heterogeneous media. The basic framework is well suited for use with other particles, such as carbon ions, antiprotons, or oxygen ions, and SDE with those should not involve deep changes in the structure. Transport through heterogeneous materials could be implemented using techniques described in the literature [67, 68, 70]. Performance of the track repeating algorithm could be improved using general performance optimization techniques. Besides being implemented in a relatively slow, interpreted language (Python), the track repeating implementation was not optimized for algorithmic performance, but instead focused on physics-based techniques. A C++ implementation of this algorithm is currently underway to increase performance.

Acknowledgements

This project was greatly aided by discussions with Niels Bassler and Michael Holscheiter about technical aspects of secondary dose, Monte Carlo, particle beams, and particle dosimetry. Pablo Yepes and Iavor Veltchev provided important insight regarding the track repeating algorithm. Shuang Luan lent his computer science knowledge in working through technical aspects of the track repeating algorithm. Dominic Maes is collaborating on the C++ implementation of this algorithm.

Chapter 6

Distributed radiation therapy simulations via cloud computing

6.1 Introduction

Dosimetric calculations are a crucial component in the accurate delivery and assessment of radiation therapy. While Monte Carlo techniques are widely seen as the gold standard of radiation dose calculations, they are only sparingly used clinically, in favor of faster, less resource intensive algorithms at the cost of dosimetric accuracy [98, 35]. The primary barrier to widespread adoption of Monte Carlo techniques has been the requirement of large computing resources to achieve clinically relevant run times, particularly in particle therapy applications. These resources, usually in the form of a computing cluster, require a sizable infrastructure investment as well as associated utility, maintenance, upgrade, and personnel costs. These costs make full Monte Carlo methods effectively unfeasible for routine clinical use, though some modified Monte Carlo methods for photon-electron calculations have made inroads recently [35]. The emerging cloud computing paradigm will remove this barrier, by making the necessary resources widely available in an economical way.

6.1.1 CPU intensive calculations in radiation therapy

As radiation therapy and diagnostic imaging techniques have become more complex, the associated physics calculations have become more resource intensive. These requirements have been largely met by the exponential increase in processor speed and RAM size, but sometimes outstrip the pace of computer technology even for conventional, deterministic calculation techniques. For example, TomoTherapy, Inc.'s TomoHD ships with a 14 node calculation cluster [99].

Non-deterministic algorithms, such as the Monte Carlo method, demand even greater computing resources than conventional algorithms, but generally offer superior dose calculation accuracy. This is particularly true for complex, heterogeneous treatment scenarios and particle therapy treatment planning, but Monte Carlo has not yet been put into routine clinical use due to long calculation times. For example, Paganetti et al. have reported times of more than 100 CPU hours to simulate proton beam treatment plans when using approximately 2×10^7 primary protons per field [41].

6.1.2 Cloud computing

Cloud computing is a name given to a set of technologies offered as services over the Internet [100]. Cloud service providers, such as Google Inc., Amazon Inc., and Microsoft Inc., offer online computing resources (i.e. CPU time, storage, software, etc.), which are scalable to the user's need. Pricing is usually based on a pay-as-you-go model, generally billed in hourly increments, and without set contract periods. This scheme allows cloud services to offer on-demand computing infrastructure, sized to fit the user's momentary needs. Cloud computing has become feasible because of the economies of scale afforded by the commodification of computer hardware, extensive availability of high bandwidth networks, and growth of free, open source

software, including entire operating systems, such as Linux, and virtual machine software.

For clinical usage the cloud computing paradigm has many potential advantages. Cloud resources can be scaled to meet patient and physics QA demand as it fluctuates on a daily basis. Typical computing clusters often face “bursty” usage: left under-utilized much of the day (and night) and over-queued at peak periods. The cloud paradigm is particularly well suited for one-off calculations, such as machine commissioning and vault shielding calculations, for which a very large cluster might be desirable, but expanding to even a small cluster would be prohibitively expensive for a single run. Additionally, hardware upgrade and maintenance is taken care of by the provider, rather than by the user.

Monte Carlo calculations are well suited to cloud style distributed computation by their fundamental nature, as the primary particle histories are completely independent of one another, requiring no communication between processes. This means that calculations, while parallel, need not maintain data or timing synchronization during execution. Previous investigations of scientific Monte Carlo calculations in a cloud environment were made in 2007 and 2009 by high energy physicists [101, 102]. The cloud computing model has not been applied to medical physics calculations until the present research.

An approach similar to cloud computing, called grid computing, has been utilized previously to perform distributed Monte Carlo calculations for radiation therapy [103, 104, 105]. Grid computing can be seen in many ways as the forerunner to cloud computing, having developed many of the distributed computing technologies and ideas upon which cloud computing is based [100]. The fundamental differences between the grid and the cloud lie in the virtual machine abstraction of the cloud and commercial implementation of the cloud infrastructure. The virtualization in the cloud abstracts away the particulars of the network and hardware architecture, so that they are transparent to the end user. Due to their commercial nature, cloud

EC2 Instance	Small	High CPU XL	High RAM Quad-XL
Compute Units*	1	20	26
USD/hr	0.085	0.68	2.0
USD/(hr \times CU*)	0.085	0.034	0.0769
Architecture	32-bit	64-bit	64-bit
Memory	1.7GB	7GB	68.4GB
Storage	160GB	1690GB	1690GB

Table 6.1: Amazon EC2 pricing for instances running Linux in US eastern region (as of April 2010). *An EC2 Compute Unit is equivalent to a 1.0-1.2 GHz 2007 Opteron or 2007 Xeon processor.

providers are likely to have infrastructures that are much larger than that of grid providers, with just a few Internet companies (including Google, Inc., Yahoo, Inc., Amazon, Inc., and Microsoft, Inc.) currently purchasing some 20% of all server computers [106]. The primary advantage of the cloud over the grid is the ubiquitous access to the cloud and the significantly lower barriers to entry by users, positioning the cloud to play a significantly larger role in medical physics computing.

6.2 The McCloud distributed processing framework

6.2.1 Software and infrastructure

To demonstrate the feasibility of performing medical physics calculations using the cloud computing paradigm, several “typical” physics calculations were performed. The calculations were carried out on Amazon Web Services (AWS) Elastic Compute Cloud (EC2) [107, 108]. Several other vendors offer similar cloud services and would have been appropriate for this research as well. AWS’s EC2 offers several different

processor-RAM combinations at different hourly rates (see Table 6.1). Each processor is rated in terms of *EC2 Compute Units*, which AWS claims is equivalent to a 1.0 - 1.2 GHz 2007 Opteron or 2007 Xeon processor.

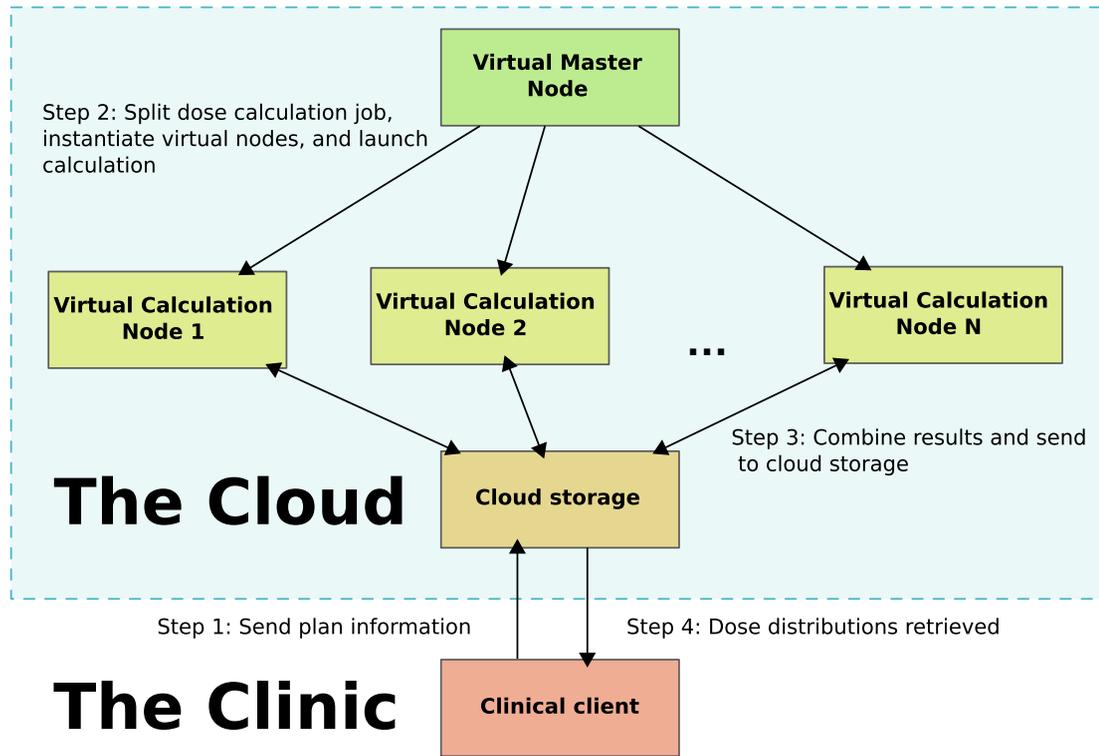


Figure 6.1: General scheme for calculating dose using the cloud paradigm.

All calculations reported in this chapter were carried out on the default EC2 Small instances with 1.7 GB of RAM, 1 virtual core, 160 GB of local disk storage, and a 32-bit architecture. Each EC2 instance runs inside a virtual machine on AWS servers. An operating system with user configured software is loaded using an Amazon Machine Image (AMI). AMI's can be chosen from a pre-configured set provided by AWS, acquired from third parties, or built by the user. From the user perspective, the boot-up of an instance using a pre-configured machine image is similar to starting a standard server computer. Once the instance is up and running, it will have a unique IP address and a domain name, allowing the user to log in.

A virtual cluster can be built by requesting multiple virtual nodes (instances). The size of the cluster can be scaled on demand (i.e. virtual nodes can be dynamically created and destroyed). Files were stored on the running EC2 instances and on the persistent Amazon Simple Storage Service (S3) to facilitate transfer to and from the cloud. From a user perspective, the AWS S3 storage service can be viewed as the counterpart of the underlying network file system (NFS) found in most cluster computing environments.

A custom distributed processing framework, dubbed `McCloud` (Monte Carlo in the Cloud) has been implemented using the Python programming language, the `boto` Python library (version 1.9b) to access AWS [109], and the secure shell protocol (SSH) for network communication. An AMI built from Fedora Linux was used for this research. Monte Carlo simulations were carried out with the FLUKA Monte Carlo package (version 2008.3b) [44, 45].

The general data flow of the cloud based calculations is as follows (see Fig. 6.1): (1) the client computer uploads the calculation parameter (input) file to the online storage. (2) The client requests N nodes for the calculation. (3) The input file is distributed to each node, given a unique, random seed, and the Monte Carlo calculation is carried out. (4) Once the dose calculations are completed, the resulting dose files are collected and combined into a single output file, which is returned to the client. The result combination step can occur on a single master node or in parallel.

6.2.2 Simulations

A variety of proof of concept studies were successfully carried out in the cloud. The studies included calculating depth-dose curves for 75 and 200 MeV proton pencil beams, a $10 \times 10 \text{ cm}^2$, 100 cm SSD¹ Co-60 beam, a $10 \times 10 \text{ cm}^2$, 100 cm SSD 10 MeV

¹Source-to-surface distance.

electron beam, and a simple broad-beam $3 \times 3 \text{ cm}^2$, single angle proton plan using a voxel phantom of heterogeneous tissues based on CT data (the Zubal phantom [110], see Figure 6.2).

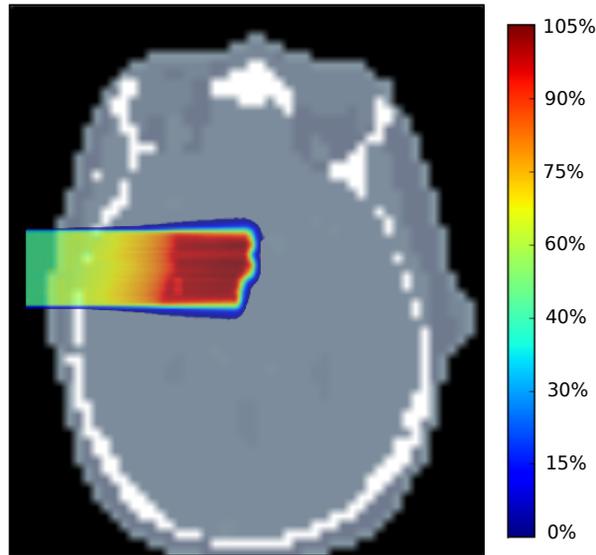


Figure 6.2: A simple broad-beam proton therapy plan calculated on 130 virtual nodes on AWS’s EC2 cloud service using a voxel based phantom.

6.3 Results

Performance data was generated with depth-dose calculations of 75 and 200 MeV proton beams. The depth-dose calculations were performed in a $40 \times 40 \times 40 \text{ cm}^3$ water phantom. To determine the dependence of total run time on the number of virtual nodes, a total of 1.4×10^7 primary protons were simulated on between 8 and 200 virtual nodes. The number of primaries was evenly divided between each node. Each data point (number of nodes) was run 3 - 4 times. The dependence of run time on cluster size is shown in Figure 6.3. For 75 MeV protons, average runs times to simulate 1.4×10^7 primaries ranged from 83 min 0 s on 8 nodes to 5 min 36 s on

200 nodes. For 200 MeV protons, average run times ranged from 222 min 54 s on 8 nodes to 10 min 37 s on 200 nodes.

Reliability statistics of the virtual cluster were also collected during the experiments. On average 1.5% of nodes (i.e. 3 nodes in a 200 node cluster) failed during the computation. To deal with node failures, a computation is declared successful if the number of nodes completing its individual computation reaches a certain threshold (e.g. 95%).

Cluster start-up times were also collected. The cluster initialization time was found to be relatively independent of the cluster size, and ranged from 3 min 22 s for an 8 node cluster to 3 min 55 s for a 200 node cluster.

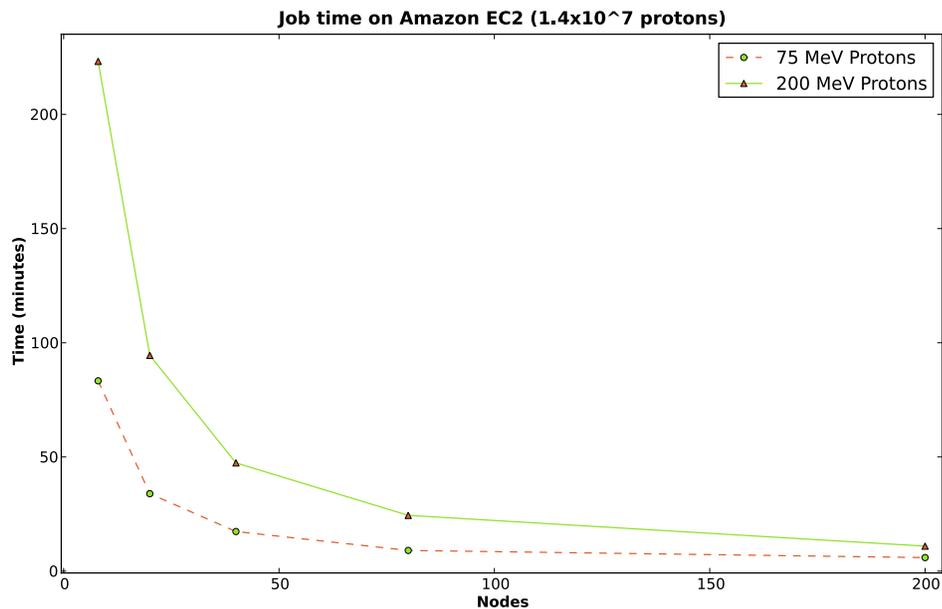


Figure 6.3: Calculation time versus number of virtual nodes for proton depth-dose curves on AWS's EC2 cloud service.

6.4 Discussion

The McCloud distributed processing framework was successfully used to perform proof-of-concept radiation therapy calculations using a cloud-based, virtual cluster. The expected, near $1/n$ speed-up was observed for clusters of n nodes. It should be noted that a perfect $1/n$ speed-up will never be achievable for an arbitrarily large number of nodes due to communication overhead. This fundamental asymptotic performance limit is described by Amdahl's Law applied to parallel processing [111].

Ultimately, the impetus behind using the cloud computing model as a computing cluster is the desire for large scale processing power without an associated large price tag. Costs associated with the cloud are generally only incurred on a usage basis, whereas in-house hardware incurs capital and maintenance costs. In order to compare the costs of the two models, the associated personnel, utility, equipment housing, insurance, vendor service, and other miscellaneous costs that might be associated with the in-house model are ignored. While this is unrealistic, it is very difficult to estimate the average costs of many of these categories (e.g. some departments may pay rent for their server space, while others need not). This set of assumptions also puts the in-house cluster in the best-case scenario.

An informal survey of various equipment purchasers within the University of New Mexico estimated the cost of a computing cluster at approximately 1000 USD per node plus approximately 200 USD in maintenance costs per year per node. For a 100 node cluster, the approximate cost over the expected 3 year lifetime of the cluster would be approximately 160,000 USD (or 53k USD per year).

The cost for using a cloud-based cluster is primarily determined by the number of CPU-hours used. The assumption used in this research was that each patient needs approximately 100 CPU-hours of cluster time. This assumption was not based on *typical* use of clinical Monte Carlo, because the types of calculations possible with an extremely large cluster are not typically performed in a clinical setting. Calculations

which might be possible with ubiquitous access to extremely large clusters, such as treatment planning with iterative Monte Carlo optimization, are the subject of future research. Assuming that a typical patient throughput for a clinic is 1000 patients per year, the CPU time necessary would be approximately 100,000 CPU-hours. AWS's current cheapest offering allows for 1 EC2 Compute Unit to be purchased for 0.034 USD (see Table 6.1). This comes out to approximately 0.10 USD per 3 GHz CPU-hour, with the assumption that 3 GHz is approximately the speed of a contemporary high end CPU core. Thus the CPU costs come out to approximately 10,000 USD per year. Additional costs are added by storing data in and transferring data to and from the cloud. A conservative estimate of data transfer and storage for this application puts these costs at less than 5% of the annual CPU costs. The resulting total annual cost using the AWS services was thus estimated at less than 10,500 USD. This was less than 20% of the estimated annual in-house cluster cost.

A major concern when transferring or storing personal or patient data on networked hardware is the security and privacy of the data. In many countries this is specifically regulated by the government. For the type of application investigated in this research, the technological aspects of securing the data are relatively well understood [112, 113, 114, 115]. Additionally, compared to full medical record or medical image storage, radiation treatment planning and similar calculations do not require as much personally identifiable information to be sent over the network (i.e. names can be replaced with hash tags). Adoption of cloud-based software for uses such as PACS and record storage, which would require compliance with the applicable regulations, seems to indicate some level of end user trust in current security implementations [116, 117].

Another technology that has gained widespread popularity in medical physics computation in recent years is the so-called GPGPU [118, 119, 120, 121, 122]. The GPU, which stands for *graphics processing unit*, is a specialized microprocessor that offloads and accelerates graphics rendering in a graphics card. The primary advan-

tage of GPUs is their relatively large number of cores, offering parallel hardware at low cost. GPGPU, which stands for *general purpose computing on GPU*, is the technique of using a GPU to perform computation in applications other than graphics [123]. GPGPU has been successfully demonstrated as a basis for radiation therapy dose calculations, including electron-photon Monte Carlo [124, 125].

While cloud computing is most often associated with CPU-based computing, it is in fact a general model. Several researchers and vendors have proposed or deployed clouds with some mix of CPU and GPU resources to provide the best possible speed-ups [126, 108]. Where appropriate, GPU clouds could potentially give medical physicists “the best of both worlds”. Choosing the best distributed computing architectures and taking advantage of their respective strengths will be a central question in the future of medical physics computing.

6.5 Conclusion and outlook

The emerging cloud computing paradigm appears to provide very interesting opportunities for computing in medical physics. We successfully demonstrated a proof-of-concept distributed calculation framework, which utilizes an on-demand, virtual computing cluster running on a commercial cloud computing service. The on-demand nature, ease of access, and pay-as-you-go pricing model yield the promise of providing clinics and researchers access to unprecedented amounts of computing power for medical physics calculations in the coming years.

The McCloud framework is currently designed to run with only one specific Monte Carlo engine, but demonstrates the conceptual basis for powerful distributed calculations using cloud computing. Future work will include: (1) Studying the performance of McCloud in larger clusters to understand the full potential and limitations of McCloud framework. (2) Investigating security policies to guarantee data in-

tegrity and patient privacy. (3) Building a web portal for researchers to upload their Monte Carlo calculations. (4) Extending `McCloud` to other Monte Carlo engines. (5) Testing `McCloud` with other cloud providers, architectures, operating systems, and performance parameters. (6) Exploring the potential of cloud-based GPU clusters for medical physics calculations. (7) Investigating the use of cloud computing for tightly coupled applications such as inverse treatment planning.

Preliminary results of this project were presented at the ICCR XVI conference in May, 2010 in Amsterdam [127]. Subsequent results were presented at the AAPM annual meetings in July, 2010 and July, 2011 [128, 129]. Additionally, an early version of this chapter was posted to arXiv.org with preliminary results in September, 2010 [130]. These results were the first known implementation of cloud-based calculations for radiation physics, although other authors have subsequently published results building on our reports.

Acknowledgements

This project was greatly helped by theoretical input from Shuang Luan and Dorian Arnold and programming contributions from Christian Romano, Shuang Luan, and Amy Reynaud. Vivek Bansal provided his radiology expertise in assigning organs and tissues to materials used with the Zubal voxel phantom FLUKA simulations. Use of Amazon Web Services was partially covered by two research grants provided by Amazon.com, LLC through the AWS in Education grant program.

Chapter 7

Conclusions

The topic of this dissertation was the development of methods to quickly compute various quantities of interest pertaining to charged hadron therapy. Three novel approaches were presented. Two applied the track repeating method to new applications: calculation of linear energy transfer (LET) distributions and estimation of secondary dose outside of the target volume. The third novel method was the use of the emerging cloud computing paradigm as the basis for large-scale, distributed radiation therapy calculations. This method is applicable to any radiation quantity of interest. Each of these methods provided the ability to calculate the quantity under consideration with performance gains relative to the standard methods.

As is typical of most research projects, the research presented here leaves some questions unanswered and sets direction for further investigations. Several speed improvement techniques are under consideration for the C++ implementation of the track repeating algorithm for secondary dose estimation. Chief among those is a new search technique over the secondary databases. It is hoped that these improvements will provide performance gains in-line with the general promise of track repeating. The cloud computing project has a number of potential future directions, including implementing track repeating in a cloud environment, calculating with GPU clouds,

and extending the current framework to other Monte Carlo engines. Each of these possibilities has the potential to enable faster, more accurate calculations for radiation therapy applications.

References

- [1] N. Bassler, I. Kantemiris, P. Karaiskos, J. Engelke, M. H. Holzscheiter, and J. B. Petersen, “Comparison of optimized single and multifield irradiation plans of antiproton, proton and carbon ion beams.,” *Radiotherapy and Oncology : Journal of the European Society for Therapeutic Radiology and Oncology*, vol. 95, pp. 87–93, Apr. 2010.
- [2] M. R. Raju, *Heavy Particle Radiotherapy*. Academic Press, 1 ed., 1980.
- [3] A. R. Smith, “Proton therapy,” *Physics in Medicine and Biology*, vol. 51, pp. R491–504, July 2006.
- [4] A. R. Smith, “Vision 2020: Proton therapy,” *Medical Physics*, vol. 36, no. 2, p. 556, 2009.
- [5] ICRP, “ICRP Publication 103: Recommendations of the International Commission on Radiological Protection,” *Annals of the ICRP*, vol. 37, no. 2-4, pp. 1–332, 2007.
- [6] NIST - Star. <http://physics.nist.gov/PhysRefData/Star/Text/programs.html>, 2011.
- [7] J. Persliden and G. A. Carlsson, “Energy imparted to water slabs by photons in the energy range 5-300 keV. Calculations using a Monte Carlo photon transport model,” *Physics in Medicine and Biology*, vol. 29, pp. 1075–1088, Sept. 1984.
- [8] World Health Organization. <http://www.who.int/cancer/>, 2011.
- [9] National Cancer Institute. <http://www.cancer.gov>, 2011.
- [10] W. Roentgen, “On a New Kind of Rays,” *Nature*, vol. 53, pp. 274–276, Jan. 1896.
- [11] V. Despeignes, “Observation concernant un cas de cancer de l’estomac traité par les rayons Röntgen,” *Lyon Médical*, vol. 82, pp. 428–430, 1896.

- [12] L. Freund, “Ein mit Röntgenstrahlen behandelter Fall von Naevus pigmentosus piliferus,” *Wiener Medizinische Wochenschrift*, vol. 47, no. 10, pp. 428–434, 1897.
- [13] E. J. Hall and A. Giaccia, *Radiobiology for the Radiologist*. Lippincott Williams & Wilkins, 6 ed., 2006.
- [14] PTCOG, “Hadron Therapy Patient Statistics,” Tech. Rep. May, 2011.
- [15] R. R. Wilson, “Radiological use of fast protons.,” *Radiology*, vol. 47, pp. 487–91, Nov. 1946.
- [16] W. H. Bragg and R. D. Kleeman, “On the ionization curves of radium,” *Philosophical Magazine and Journal of Science*, vol. 8, no. 6, pp. 726–738, 1904.
- [17] J. H. Lawrence, C. A. Tobias, J. L. Born, J. A. Linfoot, and G. J. D’Angio, “Heavy Particles in Experimental Medicine and Therapy,” *Journal of the American Medical Association*, vol. 196, no. 2, pp. 152–156, 1966.
- [18] AD-4/ACE. <http://users-phys.au.dk/hknudsen/ad4homepage.html>, 2011.
- [19] M. H. Holzscheiter, T. Ichioka, K. S. Iwamoto, S. Kovacevic, P. Møller, J. Petersen, V. Popovic, D. Scepanovic, and J. B., “Status Report on AD-4 Biological Effectiveness of Antiprotons,” tech. rep., CERN, 2004.
- [20] M. H. Holzscheiter, N. Bassler, N. Agazaryan, G. Beyer, E. Blackmore, J. J. DeMarco, M. Doser, R. E. Durand, O. Hartley, K. S. Iwamoto, H. V. Knudsen, R. Landua, C. Maggiore, W. H. McBride, S. P. Møller, J. Petersen, L. D. Skarsgard, J. B. Smathers, T. D. Solberg, U. I. Uggerhøj, S. Vranjes, H. R. Withers, M. Wong, and B. G. Wouters, “The biological effectiveness of antiproton irradiation.,” *Radiotherapy and Oncology : Journal of the European Society for Therapeutic Radiology and Oncology*, vol. 81, pp. 233–42, Dec. 2006.
- [21] N. Bassler, H. Knudsen, S. Moller, J. Petersen, D. Rahbek, and U. Uggerhoj, “Bubble detector measurements of a mixed radiation field from antiproton annihilation,” *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, vol. 251, pp. 269–273, Sept. 2006.
- [22] N. Bassler, *Experimental studies relevant for antiproton cancer therapy*. PhD thesis, University of Aarhus, 2006.

- [23] N. Bassler, J. Hansen, H. Palmans, M. Holzscheiter, and S. Kovacevic, “The antiproton depthdose curve measured with alanine detectors,” *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, vol. 266, pp. 929–936, Mar. 2008.
- [24] N. Bassler, M. H. Holzscheiter, O. Jäkel, H. V. Knudsen, and S. Kovacevic, “The antiproton depth-dose curve in water.,” *Physics in Medicine and Biology*, vol. 53, pp. 793–805, Mar. 2008.
- [25] N. Bassler, J. Alsner, G. Beyer, J. J. DeMarco, M. Doser, D. Hajdukovic, O. Hartley, K. S. Iwamoto, O. Jäkel, H. V. Knudsen, S. Kovacevic, S. P. Møller, J. Overgaard, J. B. Petersen, T. D. Solberg, B. S. Sørensen, S. Vranjes, B. G. Wouters, and M. H. Holzscheiter, “Antiproton radiotherapy.,” *Radiotherapy and Oncology : Journal of the European Society for Therapeutic Radiology and Oncology*, vol. 86, pp. 14–9, Jan. 2008.
- [26] N. Bassler, M. H. Holzscheiter, and J. B. Petersen, “Neutron fluence in antiproton radiotherapy, measurements and simulations.,” *Acta oncologica (Stockholm, Sweden)*, vol. 49, pp. 1149–59, Oct. 2010.
- [27] Optivus Proton Therapy, Inc. <http://www.optivus.com/pdf/historypbt.pdf>, 2011.
- [28] National Institute of Radiological Sciences. http://www.nirs.go.jp/ENG/research/charged_particle/, 2011.
- [29] F. H. Attix, *Introduction to Radiological Physics and Radiation Dosimetry*. Wiley-VCH, 1986.
- [30] H. Bethe, “Zur Theorie des Durchgangs schneller Korpuskularstrahlen durch Materie,” *Annalen der Physik*, vol. 397, no. 3, pp. 325–400, 1930.
- [31] W. T. Chu, B. a. Ludewigt, and T. R. Renner, “Instrumentation for treatment of cancer using proton and light-ion beams,” *Review of Scientific Instruments*, vol. 64, no. 8, p. 2055, 1993.
- [32] G. Kraft, “Tumor Therapy with Heavy Charged Particles,” *Progress in Particle and Nuclear Physics*, vol. 45, pp. S473–S544, 2000.
- [33] J. J. Wilkens and U. Oelfke, “Analytical linear energy transfer calculations for proton therapy,” *Medical Physics*, vol. 30, no. 5, p. 806, 2003.
- [34] H. Suit, T. DeLaney, S. Goldberg, H. Paganetti, B. Clasio, L. Gerweck, A. Niemierko, E. Hall, J. Flanz, J. Hallman, and A. Trofimov, “Proton vs carbon ion beams in the definitive radiation treatment of cancer patients.,”

- Radiotherapy and Oncology : Journal of the European Society for Therapeutic Radiology and Oncology*, vol. 95, pp. 3–22, Apr. 2010.
- [35] I. J. Chetty, B. Curran, J. E. Cygler, J. J. DeMarco, G. Ezzell, B. a. Faddegon, I. Kawrakow, P. J. Keall, H. Liu, C.-M. C. Ma, D. W. O. Rogers, J. Seuntjens, D. Sheikh-Bagheri, and J. V. Siebers, “Report of the AAPM Task Group No. 105: Issues associated with clinical implementation of Monte Carlo-based photon and electron external beam treatment planning,” *Medical Physics*, vol. 34, no. 12, p. 4818, 2007.
- [36] N. Metropolis, “The Beginning of the Monte Carlo Method,” *Los Alamos Science*, vol. 15, pp. 125–130, 1987.
- [37] N. Metropolis and S. Ulam, “The Monte Carlo Method,” *Journal of the American Statistical Association*, vol. 44, pp. 335–341, Nov. 1949.
- [38] F. B. Brown, “Fundamentals of Monte Carlo Particle Transport,” tech. rep., Los Alamos National Laboratory, Los Alamos, 2005.
- [39] M. H. Kalos and P. A. Whitlock, *Monte Carlo Methods*. Wiley-VCH, 2008.
- [40] F. M. Khan, *The Physics of Radiation Therapy*. Lippincott Williams & Wilkins, 2009.
- [41] H. Paganetti, H. Jiang, K. Parodi, R. Slopsma, and M. Engelsman, “Clinical implementation of full Monte Carlo dose calculation in proton beam therapy,” *Physics in Medicine and Biology*, vol. 53, pp. 4825–53, Sept. 2008.
- [42] EGSnrc. <http://irs.inms.nrc.ca/software/egsnrc/>, 2011.
- [43] FLUKA. <http://www.fluka.org/>, 2011.
- [44] A. Fassò, A. Ferrari, J. Ranft, and P. Sala, “FLUKA: a multi-particle transport code,” tech. rep., CERN-2005-10, INFN/TC_05/11, SLAC-R-773, 2005.
- [45] G. Battistoni, F. Cerutti, A. Fassò, A. Ferrari, S. Muraro, J. Ranft, S. Roesler, and P. R. Sala, “The FLUKA code: description and benchmarking,” in *Proceedings of the Hadronic Shower Simulation Workshop 2006* (M. Albrow and R. Raja, eds.), vol. 896 of *AIP Conference Proceedings*, pp. 31–49, AIP, 2007.
- [46] Geant4. <http://www.geant4.org/geant4/>, 2011.
- [47] S. Agostinelli, J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce, M. Asai, D. Axen, S. Banerjee, G. Barrant, F. Behner, L. Bellagamba, J. Boudreau, L. Broglia, A. Brunengo, H. Burkhardt, S. Chauvie, J. Chuma, R. Chytracsek, G. Cooperman, G. Cosmo, P. Degtyarenko, A. Dell’Acqua,

- G. Depaola, D. Dietrich, R. Enami, A. Feliciello, C. Ferguson, H. Fesefeldt, G. Folger, F. Foppiano, A. Forti, S. Garelli, S. Giani, R. Giannitrapani, D. Gibin, J. J. G. Cadenas, I. Gonzalez, G. G. Abril, G. Greeniaus, W. Greiner, V. Grichine, A. Grossheim, S. Guatelli, P. Gumplinger, R. Hamatsu, K. Hashimoto, H. Hasui, A. Heikkinen, A. Howard, V. Ivanchenko, A. Johnson, F. W. Jones, J. Kallenbach, N. Kanaya, M. Kawabata, Y. Kawabata, M. Kawaguti, S. Kelner, P. Kent, A. Kimura, T. Kodama, R. Kokoulin, M. Kossov, H. Kurashige, E. Lamanna, T. Lampn, V. Lara, V. Lefebure, F. Lei, M. Liendl, W. Lockman, F. Longo, S. Magni, M. Maire, E. Medernach, K. Minamimoto, P. M. de Freitas, Y. Morita, K. Murakami, M. Nagamatu, R. Nartallo, P. Nieminen, T. Nishimura, K. Ohtsubo, M. Okamura, S. O’Neale, Y. Oohata, K. Paech, J. Perl, A. Pfeiffer, M. G. Pia, F. Ranjard, A. Rybin, S. Sadilov, E. D. Salvo, G. Santin, T. Sasaki, N. Savvas, Y. Sawada, S. Scherer, S. Sei, V. Sirotenko, D. Smith, N. Starkov, H. Stoecker, J. Sulkimo, M. Takahata, S. Tanaka, E. Tcherniaev, E. S. Tehrani, M. Tropeano, P. Truscott, H. Uno, L. Urban, P. Urban, M. Verderi, A. Walkden, W. Wander, H. Weber, J. P. Wellisch, T. Wenaus, D. C. Williams, D. Wright, T. Yamada, H. Yoshida, and D. Zschesche, “G4—a simulation toolkit,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 506, no. 3, pp. 250 – 303, 2003.
- [48] J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce Dubois, M. Asai, G. Barrand, R. Capra, S. Chauvie, R. Chytrcek, G. Cirrone, G. Cooperman, G. Cosmo, G. Cuttone, G. Daquino, M. Donszelmann, M. Dressel, G. Folger, F. Foppiano, J. Generowicz, V. Grichine, S. Guatelli, P. Gumplinger, A. Heikkinen, I. Hrivnacova, A. Howard, S. Incerti, V. Ivanchenko, T. Johnson, F. Jones, T. Koi, R. Kokoulin, M. Kossov, H. Kurashige, V. Lara, S. Larsson, F. Lei, O. Link, F. Longo, M. Maire, A. Mantero, B. Mascialino, I. McLaren, P. Mendez Lorenzo, K. Minamimoto, K. Murakami, P. Nieminen, L. Pandola, S. Parlati, L. Peralta, J. Perl, A. Pfeiffer, M. Pia, A. Ribon, P. Rodrigues, G. Russo, S. Sadilov, G. Santin, T. Sasaki, D. Smith, N. Starkov, S. Tanaka, E. Tcherniaev, B. Tome, A. Trindade, P. Truscott, L. Urban, M. Verderi, A. Walkden, J. Wellisch, D. Williams, D. Wright, and H. Yoshida, “Geant4 developments and applications,” *IEEE Transactions on Nuclear Science*, vol. 53, pp. 270–278, Feb. 2006.
- [49] MCNP. <http://mcnp.lanl.gov/>, 2011.
- [50] MCNPX. <http://mcnpx.lanl.gov/>, 2011.
- [51] SHIELD-HIT. <http://www.inr.ru/shield/>, 2011.
- [52] Penelope. http://www.maestro-research.org/monte_carlo.htm, 2011.

- [53] D. Sheikh-Bagheri, I. Kawrakow, B. Walters, and D. W. Rogers, "Monte Carlo Simulations: Efficiency Improvement Techniques and Statistical Considerations," in *2006 AAPM Summer School*, 2006.
- [54] A. F. Bielajew and D. Rogers, "Presta: The parameter reduced electron-step transport algorithm for electron monte carlo transport," *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, vol. 18, no. 1-6, pp. 165–181, 1986.
- [55] A. Nahum, "Condensed-history Monte-Carlo simulation for charged particles: what can it do for us?," *Radiation and Environmental Biophysics*, vol. 38, pp. 163–73, Sept. 1999.
- [56] I. Kawrakow and M. Fippel, "Investigation of variance reduction techniques for Monte Carlo photon dose calculation using XVMC," *Physics in Medicine and Biology*, vol. 45, pp. 2163–2183, Aug. 2000.
- [57] L. L. Carter and E. Cashwell, *Particle Transport Simulation with the Monte Carlo Method*. ERDA Critical Review Series, Oak Ridge: National Technical Information Service, 1975.
- [58] T. Bortfeld, "An analytical approximation of the Bragg curve for therapeutic proton beams.," *Medical physics*, vol. 24, pp. 2024–33, Dec. 1997.
- [59] L. Hong, M. Goitein, M. Bucciolini, R. Comiskey, B. Gottschalk, S. Rosenthal, C. Serago, and M. Urie, "A pencil beam algorithm for proton dose calculations," *Physics in Medicine and Biology*, vol. 41, pp. 1305–1330, Aug. 1996.
- [60] B. Schaffner, E. Pedroni, and A. Lomax, "Dose calculation models for proton treatment planning using a dynamic beam delivery system: an attempt to include density heterogeneity effects in the analytical dose calculation," *Physics in Medicine and Biology*, vol. 44, pp. 27–41, Jan. 1999.
- [61] M. Soukup, M. Fippel, and M. Alber, "A pencil beam algorithm for intensity modulated proton therapy derived from Monte Carlo simulations.," *Physics in Medicine and Biology*, vol. 50, pp. 5089–104, Nov. 2005.
- [62] N. Bassler, O. Jäkel, C. S. Søndergaard, and J. B. Petersen, "Dose- and LET-painting with particle therapy.," *Acta oncologica (Stockholm, Sweden)*, vol. 49, pp. 1170–6, Oct. 2010.
- [63] C. Grassberger, A. Trofimov, A. Lomax, and H. Paganetti, "Variations in linear energy transfer within clinical proton therapy fields and the potential for biological treatment planning.," *International journal of radiation oncology, biology, physics*, vol. 80, pp. 1559–66, Aug. 2011.

- [64] D. Riofrio, R. Keyes, D. Maes, and S. Luan, "Simultaneous Optimization of Dose and LET in Proton Therapy Using Voronoi Partitions," *Medical Physics*, vol. 38, p. 3739, June 2011.
- [65] C. Grassberger and H. Paganetti, "Elevated LET components in clinical proton beams.," *Physics in Medicine and Biology*, vol. 56, pp. 6677–91, Oct. 2011.
- [66] I. Kawrakow, M. Fippel, and K. Friedrich, "3D electron dose calculation using a Voxel based Monte Carlo algorithm (VMC)," *Medical Physics*, vol. 23, p. 445, Apr. 1996.
- [67] P. J. Keall and P. W. Hoban, "Superposition dose calculation incorporating Monte Carlo generated electron track kernels," *Medical Physics*, vol. 23, p. 479, Apr. 1996.
- [68] J. S. Li, B. Shahine, E. Fourkal, and C.-M. Ma, "A particle track-repeating algorithm for proton beam dose calculation.," *Physics in Medicine and Biology*, vol. 50, pp. 1001–10, Mar. 2005.
- [69] P. Yepes, S. Randeniya, P. J. Taddei, and W. D. Newhauser, "A Track-Repeating Algorithm for Fast Monte Carlo Dose Calculations of Proton Radiotherapy.," *Nuclear technology*, vol. 168, pp. 736–740, Dec. 2009.
- [70] P. Yepes, S. Randeniya, P. J. Taddei, and W. D. Newhauser, "Monte Carlo fast dose calculator for proton radiotherapy: application to a voxelized geometry representing a patient with prostate cancer.," *Physics in Medicine and Biology*, vol. 54, pp. N21–8, Jan. 2009.
- [71] P. P. Yepes, D. Mirkovic, and P. J. Taddei, "A GPU implementation of a track-repeating algorithm for proton radiotherapy dose calculations.," *Physics in Medicine and Biology*, vol. 55, pp. 7107–7120, Dec. 2010.
- [72] C. M. Ma, B. A. Faddegon, D. W. O. Rogers, and T. R. Mackie, "Accurate characterization of Monte Carlo calculated electron beams for radiotherapy," *Medical Physics*, vol. 24, p. 401, Mar. 1997.
- [73] A. E. Schach von Wittenau, L. J. Cox, P. M. Bergstrom, W. P. Chandler, C. L. H. Siantar, and R. Mohan, "Correlated histogram representation of Monte Carlo derived medical accelerator photon-output phase space," *Medical Physics*, vol. 26, p. 1196, July 1999.
- [74] M. K. Fix, M. Stampanoni, P. Manser, E. J. Born, R. Mini, and P. Rügsegger, "A multiple source model for 6 MV photon beam dose calculations using Monte Carlo," *Physics in Medicine and Biology*, vol. 46, pp. 1407–1427, May 2001.

- [75] ICRU, “ICRU Report 49, Stopping Powers and Ranges for Protons and Alpha Particles,” tech. rep., International Commission on Radiation Units and Measurements, 1993.
- [76] Python. <http://www.python.org/>, 2011.
- [77] B. Stroustrup, *C++ Programming Language, The*. Addison-Wesley Professional, 3 ed., 1997.
- [78] Stroustrup: C++. <http://www2.research.att.com/~bs/C++.html>, 2011.
- [79] E. Hall and C. Wu, “Radiation-induced second cancers: the impact of 3D-CRT and IMRT,” *International Journal of Radiation OncologyBiologyPhysics*, vol. 56, pp. 83–88, May 2003.
- [80] P. M. Monteleone and A. T. Meadows, “Late effects of childhood cancer and treatment,” *eMedicine*, 2006.
- [81] E. J. Hall, “Intensity-modulated radiation therapy, protons, and the risk of second cancers.” *International journal of radiation oncology, biology, physics*, vol. 65, pp. 1–7, May 2006.
- [82] B. S. Athar and H. Paganetti, “Neutron equivalent doses and associated lifetime cancer incidence risks for head & neck and spinal proton therapy.” *Physics in Medicine and Biology*, vol. 54, pp. 4907–26, Aug. 2009.
- [83] X. G. Xu, B. Bednarz, and H. Paganetti, “A review of dosimetry studies on external-beam radiation treatment with respect to second cancer induction.” *Physics in Medicine and Biology*, vol. 53, pp. R193–241, July 2008.
- [84] B. Bednarz, C. Hancox, and X. G. Xu, “Calculated organ doses from selected prostate treatment plans using Monte Carlo simulations and an anatomically realistic computational phantom.” *Physics in Medicine and Biology*, vol. 54, pp. 5271–86, Sept. 2009.
- [85] J. D. Ruben, S. Davis, C. Evans, P. Jones, F. Gagliardi, M. Haynes, and A. Hunter, “The effect of intensity-modulated radiotherapy on radiation-induced second malignancies.” *International journal of radiation oncology, biology, physics*, vol. 70, pp. 1530–6, Apr. 2008.
- [86] W. D. Newhauser, J. D. Fontenot, A. Mahajan, D. Kornguth, M. Stovall, Y. Zheng, P. J. Taddei, D. Mirkovic, R. Mohan, J. D. Cox, and S. Woo, “The risk of developing a second cancer after receiving craniospinal proton irradiation,” *Physics in Medicine and Biology*, vol. 54, pp. 2277–2291, Apr. 2009.

- [87] S. Yonai, N. Matsufuji, T. Kanai, Y. Matsui, K. Matsushita, H. Yamashita, M. Numano, T. Sakae, T. Terunuma, T. Nishio, R. Kohno, and T. Akagi, "Measurement of neutron ambient dose equivalent in passive carbon-ion and proton radiotherapies," *Medical Physics*, vol. 35, no. 11, p. 4782, 2008.
- [88] S. Yonai, Y. Kase, N. Matsufuji, T. Kanai, T. Nishio, M. Namba, and W. Yamashita, "Measurement of absorbed dose, quality factor, and dose equivalent in water phantom outside of the irradiation field in passive carbon-ion and proton radiotherapies," *Medical Physics*, vol. 37, no. 8, p. 4046, 2010.
- [89] C. Zacharatou Jarlskog, C. Lee, W. E. Bolch, X. G. Xu, and H. Paganetti, "Assessment of organ-specific neutron equivalent doses in proton therapy using computational whole-body age-dependent voxel phantoms.," *Physics in Medicine and Biology*, vol. 53, pp. 693–717, Feb. 2008.
- [90] I. Alziar, G. Bonniaud, D. Couanet, J. B. Ruard, C. Vicente, G. Giordana, O. Ben-Harrath, J. C. Diaz, P. Grandjean, H. Kafrouni, J. Chavaudra, D. Lefkopoulos, F. de Vathaire, and I. Diallo, "Individual radiation therapy patient whole-body phantoms for peripheral dose evaluations: method and specific software.," *Physics in Medicine and Biology*, vol. 54, pp. N375–83, Sept. 2009.
- [91] Y. H. Na, B. Zhang, J. Zhang, P. F. Caracappa, and X. G. Xu, "Deformable adult human phantoms for radiation protection dosimetry: anthropometric data representing size distributions of adult worker populations and software algorithms.," *Physics in Medicine and Biology*, vol. 55, pp. 3789–811, July 2010.
- [92] M. Hultqvist and I. Gudowska, "Secondary absorbed doses from light ion irradiation in anthropomorphic phantoms representing an adult male and a 10 year old child.," *Physics in Medicine and Biology*, vol. 55, pp. 6633–6653, Oct. 2010.
- [93] W. Newhauser. Personal communication, 2009.
- [94] R. W. Keyes, S. Luan, and M. Holzscheiter, "Antiproton Therapy: A Simplified Method to Characterize and Compare Dose From Peripheral Radiation Fields," *Medical Physics*, vol. 36, p. 2617, July 2009.
- [95] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*. The MIT Press, 3 ed., 2009.
- [96] A. S. Householder, "Unitary Triangularization of a Nonsymmetric Matrix," *Journal of the ACM*, vol. 5, pp. 339–342, Oct. 1958.

- [97] NumPy. <http://numpy.scipy.org/>, 2011.
- [98] D. W. O. Rogers, “Fifty years of Monte Carlo simulations for medical physics.,” *Physics in Medicine and Biology*, vol. 51, pp. R287–301, July 2006.
- [99] TomoTherapy, Inc. <http://tomotherapy.com>, 2010.
- [100] I. Foster, Y. Zhao, I. Raicu, and S. Lu, “Cloud computing and grid computing 360-degree compared,” in *Grid Computing Environments Workshop, 2008. GCE '08*, pp. 1–10, 12-16 2008.
- [101] K. Keahey and T. Freeman, “Science clouds: Early experiences in cloud computing for scientific applications,” in *Cloud Computing and Its Applications 2008*, October 2008.
- [102] M. Sevier and T. Fifield, “Belle Monte-Carlo production on the Amazon EC2 cloud,” in *Victorian Partnership for Advanced Computing 2009*, April 2009.
- [103] A. Gomez, “Remote Radiotherapy Planning: The eIMRT Project,” *Stud. Health. Technol. Inform.*, vol. 120, pp. 330–5, 2006.
- [104] A. Gomez, C. Fernandez Sanchez, J. Mourino Gallego, J. Lopez Cacheiro, F. Gonzalez Castano, D. Rodriguez-Silva, L. Dominguez Carrera, D. Gonzalez Martinez, J. Pena Garcia, F. Gomez Rodriguez, D. Gonzalez Castano, and M. Pombar Camean, “Monte Carlo verification of IMRT treatment plans on grid,” *Stud. Health. Technol. Inform.*, vol. 126, pp. 105–14, 2007.
- [105] P. Downes, G. Yaikhom, J. Giddy, D. Walker, E. Spezi, and D. Lewis, “High-performance computing for Monte Carlo radiotherapy calculations,” *Phil. Trans. R. Soc. A*, vol. 367, no. 1897, pp. 2607–17, 2009.
- [106] R. Waters, “How many computers does the world need?.” <http://blogs.ft.com/techblog/2009/03/how-many-computers-does-the-world-need>, 2009.
- [107] Amazon Web Services. <http://aws.amazon.com/>, 2010.
- [108] Amazon Elastic Compute Cloud. <http://aws.amazon.com/ec2/>, 2010.
- [109] boto. <http://code.google.com/p/boto/>, 2010.
- [110] I. G. Zubal, C. R. Harrell, E. O. Smith, Z. Rattner, G. Gindi, and P. B. Hoffer, “Computerized three-dimensional segmented human anatomy,” *Med. Phys.*, vol. 21, no. 2, pp. 299–302, 1994.

- [111] J. L. Hennessy and D. A. Patterson, *Computer Architecture: A Quantitative Approach (The Morgan Kaufmann Series in Computer Architecture and Design)*. Morgan Kaufmann, May 2002.
- [112] D. R. Stinson, *Cryptography: Theory and Practice*. CRC Press, 1995.
- [113] M. Bishop, *Computer Security: Art and Science*. Addison-Wesley, 2002.
- [114] M. Y. Rhee, *Internet Security: Cryptographic Principles, Algorithms and Protocols*. Wiley, 2003.
- [115] M. Bishop, *Introduction to Computer Security*. Addison-Wesley, 2004.
- [116] J. Philbin, F. Prior, and P. Nagy, “Will the next generation of PACS be sitting on a cloud?,” *J. Digit. Imaging.*, pp. 1–5, 2010.
- [117] K. P. Andriole and R. Khorasani, “Cloud computing: What is it and could it be useful?,” *J. Am. Coll. Radiol.*, vol. 7, no. 4, pp. 252 – 254, 2010.
- [118] M. de Greef, J. Crezee, J. C. van Eijk, R. Pool, and A. Bel, “Accelerated ray tracing for radiotherapy dose calculations on a GPU,” *Med. Phys.*, vol. 36, pp. 4095–4102, 2009.
- [119] X. Gu, D. J. Choi, C. Men, H. Pan, A. Majumdar, and S. B. Jiang, “GPU-based ultra-fast dose calculation using a finite size pencil beam model,” *Phys. Med. Biol.*, vol. 54, pp. 6287–6297, 2009.
- [120] C. Men, X. Gu, D. J. Choi, A. Majumdar, Z. Zheng, K. Mueller, and S. B. Jiang, “GPU-based ultrafast IMRT plan optimization,” *Phys. Med. Biol.*, vol. 54, pp. 6565–6573, 2009.
- [121] R. Jacques, J. Wong, R. Taylor, and T. McNutt, “Real-time dose computation: GPU-accelerated source modeling and superposition/convolution,” *Medical Physics*, vol. 38, no. 1, pp. 294–305, 2011.
- [122] S. Hissoiny, B. Ozell, and P. Despres, “A convolution-superposition dose calculation engine for GPUs,” *Medical Physics*, vol. 37, no. 3, pp. 1029–1037, 2010.
- [123] J. D. Owens, D. Luebke, N. Govindaraju, M. Harris, J. Kruger, A. E. Lefohn, and T. J. Purcell, “A survey of general-purpose computation on graphics hardware,” *Comput. Graph. Forum*, vol. 26, pp. 80–113, 2007.
- [124] X. Jia, X. Gu, J. Sempau, D. Choi, A. Majumdar, and S. B. Jiang, “Development of a GPU-based monte carlo dose calculation code for coupled electron-photon transport,” *Physics in Medicine and Biology*, vol. 55, no. 11, p. 3077, 2010.

- [125] S. Hissoiny, B. Ozell, H. Bouchard, and P. Despres, “GPUMCD: A new GPU-oriented monte carlo dose calculation platform,” *Medical Physics*, vol. 38, no. 2, pp. 754–764, 2011.
- [126] G. Giunta, R. Montella, G. Agrillo, and G. Coviello, “A GPGPU transparent virtualization component for high performance computing clouds,” in *Euro-Par 2010 - Parallel Processing* (P. D’Ambra, M. Guarracino, and D. Talia, eds.), vol. 6271 of *Lecture Notes in Computer Science*, pp. 379–391, Springer Berlin / Heidelberg, 2010.
- [127] R. W. Keyes, C. Romano, D. Arnold, and S. Luan, “Cloud Computing as a Monte Carlo Cluster for Radiation Therapy,” in *Proceedings of the XVth ICCR*, (Amsterdam), 2010.
- [128] R. Keyes, C. Romano, D. Arnold, and S. Luan, “Medical Physics Calculations in the Cloud: A New Paradigm for Clinical Computing,” *Medical Physics*, vol. 37, no. 6, p. 3272, 2010.
- [129] R. W. Keyes, D. Arnold, A. Reynaud, and S. Luan, “McCloud: Toward 10 Million Monte Carlo Primaries in 5 Minutes for Clinical Use,” *Medical Physics*, vol. 38, p. 3648, June 2011.
- [130] R. W. Keyes, C. Romano, D. Arnold, and S. Luan, “Radiation therapy calculations using an on-demand virtual cluster via cloud computing,” *arXiv.org*, p. 12, Sept. 2010.