A COARSE MESH TRANSPORT METHOD FOR PHOTONS AND ELECTRONS IN 3-D

A Dissertation
Presented to
The Academic Faculty

By

Robert M. Hayward

In Partial Fulfillment
Of the Requirements for the Degree
Doctor of Philosophy in Nuclear and Radiological Engineering

Georgia Institute of Technology
May 2013

Copyright © 2013 by Robert M. Hayward

A COARSE MESH TRANSPORT METHOD FOR PHOTONS AND ELECTRONS IN 3-D

Approved by:

Dr. Farzad Rahnema, Advisor
George W. Woodruff School of
Mechanical Engineering
Georgia Institute of Technology

Dr. Dingkang Zhang
George W. Woodruff School of
Mechanical Engineering
Georgia Institute of Technology

Dr. Chris Wang
George W. Woodruff School of
Mechanical Engineering
Georgia Institute of Technology

Dr. Thomas Morely
School of Mathematics
Georgia Institute of Technology

Dr. Eric Elder
Department of Radiation Oncology
Emory University School of Medicine

Date Approved: March 27, 2013

The materials provided herein are proprietary in nature and are subject to copyrights by the Georgia Tech Research Corporation and/or employees of the Georgia Institute of Technology. Any materials which may be patented under United States patent law are owned by the Georgia Tech Research Corporation. Any reproduction, transmission, creation of a derivative work, or other use, in whole or in part, is prohibited without prior written permission.

To my wife,

Erin,

for her love and devotion.

ACKNOWLEDGEMENTS

I would like to express my deep gratitude to my entire committee for their seemingly endless patience during the preparation of this manuscript. I want to thank Dr. Rahnema and Dr. Zhang for valuable discussions, guidance, and insight. I would also like to thank Dr. Elder for providing anonymized CT scan data. Although I never met him, I would like to acknowledge the work of the late John D. Hunter, creator of the `matplotlib` package that was used to draw nearly every figure and diagram in this work, enabling a level of graphical eloquence that would have otherwise been difficult to achieve. Finally, I would like to thank my parents for their support and encouragement along the way.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

SUMMARY

A hybrid stochastic-deterministic method, COMET-PE, is developed for dose calculation in radiotherapy. Fast, accurate dose calculation is a key component of successful radiotherapy treatment. To calculate dose, COMET-PE solves the coupled Boltzmann Transport Equations for photons and electrons. The method uses a deterministic iteration to compose response functions that are pre-computed using Monte Carlo. Thus, COMET-PE takes advantage of Monte Carlo physics without incurring the computational costs typically required for statistical convergence. This work extends the method to 3-D problems with realistic source distributions. Additionally, the performance of the deterministic solver is improved, taking advantage of both shared-memory and distributed-memory parallelism to enhance efficiency. To verify the method's accuracy, it is compared with the DOSXYZnrc (Monte Carlo) method using three different benchmark problems: a heterogeneous slab phantom, a water phantom, and a CT-based lung phantom. For the slab phantom, all errors are less than 1.5% of the maximum dose or less than 3% of local dose. For both the water phantom and the lung phantom, over 97% of voxels receiving greater than 10% of the maximum dose pass a 2% (relative error) / 2 mm (distance-to-agreement) test. Timing comparisons show that COMET-PE is roughly 10-30 times faster than DOSXYZnrc. Thus, the new method provides a fast, accurate alternative to Monte Carlo for dose calculation in radiotherapy treatment planning.

CHAPTER I

INTRODUCTION

1.1 Motivation

An ideal cancer therapy would target and kill cancerous cells perfectly, leaving normal tissue unharmed. Unfortunately, such a therapy does not yet exist. While radiation therapy is an effective means of killing cancerous cells, it can also harm normal tissue, producing undesirable side effects. In fact, the radiation dose that can be delivered to a tumor is limited by the dose that will be delivered to the surrounding healthy tissue. Ideally, all dose would be deposited in the tumor itself, leaving surrounding tissue untouched, but this ideal has not yet been realized. This presents at least two problems:

1. The location of the tumor must be known with high accuracy.

2. Dose must be delivered accurately to the tumor's location.

There have been large advances in imaging technologies such as computed tomography (CT), magnetic resonance imaging (MRI), positron emission tomography (PET), single photon emission computed tomography (SPECT), and ultrasound. Furthermore, technology such as image fusion allow multi-modality imaging. Put together, all of this allows physicians to locate a tumor with greater accuracy than ever before.

These advances in imaging coupled with CT-based treatment planning and image-guided radiation therapy (IGRT) have resulted in treatments that targets tumors with greater accuracy and precision than ever before. The goal of escalating dose to the tumor while sparing normal tissue thus leads to treatment plans with high dose gradients immediately surrounding the tumor. Often this can be located in a region with highly heterogeneous tissue composition such as the lung. These situations demand accurate dose calculation.

Monte Carlo methods set the standard for accurate dose calculation [14]. This is because Monte Carlo methods can model particle transport with great detail, in terms of both physics and geometry. This great accuracy, however, comes at the price of precision. Because Monte Carlo calculations are stochastic in nature, all of the results they produce have inherent statistical uncertainty. The cost of reducing this uncertainty is increased computational effort. Unfortunately, the effort required for clinically acceptable precision is often impractically large. This is especially true when one would like to calculate the dose for a large irradiated volume or with high spatial resolution. Until computing power is cheap enough to enable the clinical ubiquity of Monte Carlo, alternative methods for dose calculation are required.

1.2 Objectives

The goal of this work is to develop one such method. The Coarse Mesh Transport (COMET) method has already been shown to be capable of Monte Carlo accuracy when applied to problems in nuclear reactor physics [55]. A variant of the method known as COMET-PE (for photon and electron transport) has also proven to be both accurate and efficient when calculating energy deposition in 2-D for a restricted set of source distributions [9, 54]. This work extends COMET-PE to allow photon and electron transport calculations in 3-D with realistic source distributions. This development includes a ground-up re-write of the codebase that introduces explicit analytical transport for uncollided photons, a new set of orthogonal basis functions to represent interface fluence, and improvements to the efficiency of the deterministic solver, which can now take full advantage of both shared-memory and distributed-memory parallel architectures. The newly implemented method is benchmarked by comparison with a Monte Carlo method using three benchmark problems: a slab phantom containing water, aluminum, and lung; a water phantom typical of accelerator calibration, and a phantom derived from CT data of a lung treatment.

CHAPTER II

LITERATURE REVIEW

Dose calculation methods can be divided into two broad categories: those that are based on transport theory and those that are not. Methods based on transport theory, whether stochastic or deterministic, provide a rigorous treatment of the underlying physics that is valid regardless of the presence of material heterogeneity or strong fluence gradients. The majority of methods in clinical use, however, are not based on transport theory. Instead these methods make empirically based approximations or simplifications in order to achieve short calculation times. The most advanced of these methods are those based on convolution and superposition.

2.1 Methods based on Convolution and Superposition

In the most general sense convolution-superposition methods calculate dose using a relationship of the form

$$D(\mathbf{r}) = \int_V \mathrm{d}\mathbf{r}' \int_{4\pi} \mathrm{d}\hat{\Omega} \int_0^\infty \mathrm{d}E\, G(\mathbf{r}, \mathbf{r}', \hat{\Omega}, E)\, \psi(\mathbf{r}', \hat{\Omega}, E) \tag{1}$$

where

$\mathbf{r}, \mathbf{r}'$ = position vectors

$\hat{\Omega}$ = direction of primary photons

E = primary photon energy

V = the patient volume

D = dose

ψ = primary photon fluence.

The kernel G is a Green's function derived from the transport equation that gives the dose at a point $\mathbf{r}$ that results from a unit primary photon fluence at the phase space location $(\mathbf{r}', \hat{\Omega}, E)$. In theory, this relationship is exact. However, the Green's function G is patient-specific; it depends on the patient's geometry and material composition. This means that G cannot be pre-computed, yet evaluating G at a single phase space point $(\mathbf{r}', \hat{\Omega}, E)$ requires solving the transport equation. In other words, calculating dose using (1) is more difficult than solving the transport equation directly. In practice several different approximate approaches are taken to overcome this difficulty. Most commonly [30, 11], the dose is approximated by a superposition integral of the form

$$D(\mathbf{r}) = \int_V \mathrm{d}\mathbf{r}' \, k_s(\mathbf{r}, \mathbf{r}') \, \Phi(\mathbf{r}') \tag{2}$$

or by a convolution integral of the form

$$D(\mathbf{r}) = \int_V \mathrm{d}\mathbf{r}' \, k_c(\mathbf{r} - \mathbf{r}') \, \Phi(\mathbf{r}'). \tag{3}$$

The superposition kernel k_s is more general and requires fewer approximations than the spatially invariant convolution kernel k_c, but this comes at the cost of additional computation.

To derive a spatially invariant kernel k_c for the convolution method, one must assume that the spatial dependence of the fluence is separable from the energy-angle dependence, so that $\psi(\mathbf{r}, \hat{\Omega}, E) = \Phi(\mathbf{r}) f(\hat{\Omega}, E)$. By additionally assuming a homogeneous phantom of infinite extent, one can evaluate a convolution kernel or *dose spread array* k_c such that

$$k_c(\mathbf{r}) = \int_{4\pi} \mathrm{d}\hat{\Omega} \int_0^\infty \mathrm{d}E \, G(\mathbf{r}, \mathbf{0}, \hat{\Omega}, E) \, f(\hat{\Omega}, E).$$

The kernel k_c can be calculated analytically for low energy photon beams, where one can assume that electrons deposit energy locally [11]. When this is not the case, k_c can be evaluated using Monte Carlo methods [29] . Because (3) is a convolution, it can be evaluated as a multiplication in Fourier space, specifically using a 3-D fast Fourier transform (FFT), which is extremely efficient. For a phantom with N^3 voxels, the calculation can be carried

out with $O(N^3 \log N)$ arithmetic operations, compared with $O(N^6)$ for the general case. Unfortunately, the assumptions involved in deriving the spatially invariant kernel limit its accuracy. Specifically, the assumption of fluence separability means that pure convolution methods cannot account for either the beam's angular divergence or for variations in the beam's energy spectrum as the beam passes through matter. Since low energy photons are attenuated more than high energy photons, the beam's energy spectrum changes as a function of depth due to the phantom itself, and the shape of the flattening filter causes variation with lateral distance from the beam's central axis. Perhaps more important than the assumptions about separability, however, is the assumption that the phantom is both homogeneous and infinite. As a result of this assumption, the convolution method is invalid near the phantom boundary or in the presence of material heterogeneities such as lung or bone.

At the cost of an $O(N^6)$ calculation, one can gain increased accuracy by using a more general superposition method. With a superposition method the separability constraints can be relaxed. This allows the angular divergence of the source to be handled by tilting the kernel as a function of position relative to the source. Similarly, the kernel can be made to vary as a function of depth to account for hardening of the beam spectrum. The effect of heterogeneities on the transport of secondary particles can be approximated using O'Connor's theorem [35] and evaluating the kernels in terms of radiological coordinates–a product of length and density (g/cm^2)–rather than the usual spatial coordinates (cm). This correction is based on the fact that interaction cross sections are directly proportional to density in different media of the same atomic composition, such as lung and water. This correction is not exact, but it provides an improvement over the assumption of homogeneous media. For media which differ in effective atomic number, such as water and bone, electron density may be used instead of mass density to scale the superposition kernel [30]. Such a correction provides only a crude approximation especially when photoelectric absorption, pair production, and electron transport–all of which have significant dependence on atomic

number–are important. This limitation is common to all superposition-style methods.

To improve upon the $O(N^6)$ scaling of the methods described so far, Ahnesjö introduced the Collapsed Cone Convolution (CCC) method [3] and the later the Pencil Beam Kernel (PBK) method [5]. The Pencil Beam Kernel method uses a 2-D kernel of the form

$$k_{PBK}(r,z) = \frac{A_z e^{-a_z r}}{r} + \frac{B_z e^{-b_z r}}{r} \tag{4}$$

where r is the cylindrical radius from the pencil beam axis, z is the depth, and A_z, a_z, B_z, and b_z are z-dependent parameters found by fit to Monte Carlo calculation. A_z and a_z characterize the relatively long-range secondary photon transport while B_Z and b_z characterize the transport of electrons and low-energy photons. This is a good approximation for first-scatter photons, but difficult to justify theoretically for electrons or multiply scattered photons. Furthermore, the Pencil Beam Kernel method does not accurately account for lateral heterogeneity in the distribution of particles from the point of primary interaction. The Collapsed Cone Convolution (CCC) is more sophisticated in this regard. Rather than a 2-D pencil kernel, the method uses a 3-D kernel

$$k_{CCC}(r,z,\Theta) = \frac{A_{z\Theta} e^{-a_{z\Theta} r}}{r^2} + \frac{B_{z\Theta} e^{-b_{z\Theta} r}}{r^2} \tag{5}$$

where r is the spherical radius from the primary interaction point, z is depth, Θ is the scattering angle, and $A_{z\Theta}$, $a_{z\Theta}$, $B_{z\Theta}$, and $b_{z\Theta}$ are parameters that depend on z and Θ. To avoid the $O(N^6)$ scaling, dose is deposited along a set of discrete rays leaving each primary interaction point. As with the discrete ordinates approximation in transport theory, this can lead to ray effects at large distances from the point of primary interaction. However, ray effects are not usually observed in practice thanks to the smoothness of the TERMA distribution. Unlike the pencil beam method, CCC can approximate the effects of heterogeneity on secondary particle transport by scaling its dose kernel in terms of radiological coordinates. The main drawback associated with this method is similar to that of the pencil beam kernel: The empirical approximation of electron "attenuation" is not rigorously justified.

The Analytical Anisotropic Algorithm (AAA) is another superposition method that uses a simplified kernel to achieve fast calculation times [48]. The kernel is separated into a 1-D kernel that acts along the direction of incident fluence, and a set of 2-D kernels that spread dose laterally, perpendicular to the direction of incident fluence. The cylindrical geometry of the 2-D kernels is similar to that of the Pencil Beam Kernel method. Unlike PBK however, the lateral kernels in the Analytical Anisotropic Algorithm are applied over a set of discrete directions and scaled to account for lateral heterogeneity. As with CCC and other methods that use kernels scaled by local electron density to correct for heterogeneity, AAA cannot rigorously model the transport of secondary particles, especially electrons, in heterogeneous media.

This fundamental limitation stems from the original formulation of these methods in terms of transport theory Green's functions. The exact Green's function is inherently dependent on the geometry and material composition of each unique patient. In order to make superposition methods practical, however, dose kernels are pre-computed based on calculations in homogeneous media. When using kernels calculated this way, any correction for heterogeneity is approximate at best.

Another limitation of the convolution-superposition methods is that they treat the angular dependence of the primary fluence implicitly. This means that when a treatment plan has multiple fields, the entire calculation must be repeated for each field. This can lead to long calculation times for complicated plans with the extreme example being rotational or arc therapy. Naqvi et al. have suggested that the superposition operation can be implemented using a Monte Carlo method to avoid this problem [33]. While this does prevent the calculation time from scaling with the number of fields, it introduces statistical uncertainty on top of the systematic errors that are inherent in convolution-superposition methods. Since the removal of this uncertainty is governed by the Central Limit Theorem, long calculation times will result if high precision is required. Thus, this approach may be impractical for the same reasons as pure Monte Carlo methods. For more on the class of

superposition-convolution methods the reader is referred to the review by Ahnesjö [4] and to the American Association of Physics in Medicine's (AAPM) Report 85 [36].

2.2 Transport Theory Based Methods for Dose Calculation

Transport theory provides the most accurate way to calculate dose because it is formulated in terms of the fundamental processes involved in the interaction of radiation with matter. Transport theory remains valid regardless of material heterogeneity or charged particle equilibrium [17]. Stochastic (or Monte Carlo) methods are widely considered the "gold standard" for accuracy in dose calculation. Deterministic methods, though less widely used, provide an alternative for accurate dose calculations. Both deterministic and stochastic methods yield the same solution in the asymptotic limit, but there are vast differences between the two approaches. Stochastic methods simulate one particle at a time. By observing a large number of particle histories, one attempts to determine the average or *expected* behavior. The resulting computation is straightforward, but statistical uncertainty pervades the results. Deterministic methods, on the other hand, attempt to solve for the mean distribution of all particles directly. In some sense, this is equivalent to following all possible particle histories simultaneously. This requires discretizing the particle phase space. Instead of statistical uncertainty, deterministic methods suffer from systematic discretization errors. Using either approach, one can trade increased computational effort for a higher quality solution. With stochastic methods, simulating more particle histories decreases statistical uncertainty. With deterministic methods, a finer discretization reduces systematic errors.

2.2.1 Stochastic Methods for Dose Calculation

The Monte Carlo method is perhaps the most intuitive method for radiation transport. The Monte Carlo method is a stochastic method because particle histories are simulated one at a time using pseudo-random numbers to sample known statistical interaction probabilities.

Thus each history follows a unique path through the problem phase space. By accumulating tally data from enough histories, one can estimate integral quantities, such as energy deposition. The Monte Carlo method was originally invented to solve neutron transport problems arising from work related to the Manhattan project, and it was one of the first methods to be implemented on an electronic computer. Indeed the ENIAC, the world's first electronic computer, was programmed to run Monte Carlo simulations [32]. Stochastic methods are unique in their ability to model complex geometries and detailed physics without approximation. Recognizing the advantages of this approach, the medical physics community has employed stochastic techniques for decades [37, 6, 38].

There are two main categories of Monte Carlo codes that are relevant to the present work: general-purpose codes and treatment-planning codes. General purpose codes used in Medical Physics include MCNP [47], GEANT [2], EGSnrc [24], and PENELOPE [40]. These codes focus on geometric flexibility and physical correctness, making them applicable to the solution of many different transport problems. This class of Monte Carlo codes is capable of achieving unparalleled accuracy in the transport of photons and electrons. EGSnrc, in particular, is used throughout this work as a standard for accuracy. In contrast to general-purpose codes, treatment-planning codes are aimed at making Monte Carlo computations practical (i.e. fast enough) for routine clinical use. In order to accomplish this, these codes make heavy use of variance reduction and of optimizations that trade generality or physical detail for speed. Examples of such codes are VMC++ [23], PEREGRINE [20], and DPM [41]. The report of AAPM Task Group 105 [14] provides extensive discussion on the clinical use of Monte Carlo methods.

The precision of stochastic methods is governed by the central limit theorem, which states that, given a large number of histories N, the standard error of an integral estimator decreases as $N^{-1/2}$ [44]. Thus, to reduce the error by one half, one must run four times as many histories. When used for dose calculation, one must also pay a price for increased dose resolution. This is because as the volume of each cell in the dose grid decreases, the

probability that any given history will contribute to that cell decreases. Thus, to maintain statistical precision while increasing resolution, one must also increase the number of histories, and hence, the computational effort. Furthermore, it is difficult to be confident in the dose at any single point. Indeed, clinicians implementing Monte Carlo treatment planning systems are advised to modify routine clinical procedures, such as prescribing dose to a single point, in order to avoid sensitivity to statistical fluctuations [14].

2.2.2 Deterministic Methods for Dose Calculation

Unlike stochastic methods, deterministic methods do not suffer from statistical uncertainty. These methods, although widely used in the modeling of nuclear reactors, are relatively new to the Medical Physics community. Deterministic methods solve for the fluence distribution throughout the problem volume by discretizing the phase space. Since the transport equation is linear, this discretization results in a system of linear equations that can be solved to find the fluence. Typically the discretization involves the discrete ordinates approximation in the angular variable, the multigroup approximation in the energy variable, and either a finite difference or a finite element discretization in the spatial variable. Furthermore, the Boltzmann scattering operator is often approximated via a truncated expansion in Legendre polynomials. [26]. The discrete ordinates approximation is equivalent to requiring the transport equation to hold only for a finite set of discrete solid angles. When the number of angles is too small, the error is manifest in the form of artificial oscillations, called ray effects, in the spatial component of the solution. The most common approach to mitigate ray effects, aside from increasing the number of discrete angles, is to use a first-collision source. That is, primary particles are transported via ray-tracing to the point of their first collision before the discrete ordinates approximation is applied.

Several deterministic transport theory methods have been applied to problems in Medical Physics although applications of coupled photon-electron transport for dose calculation are rare. Daskalov et al. show that DANTSYS, a discrete ordinates method, can be used

to calculate dose for ^{192}Ir and ^{125}I brachytherapy problems in 2-D cylindrical R-Z geometry [16, 15]. Due to the presence of charged particle equilibrium, only photon transport is required for these calculations. The method is accurate within 5% of the Monte Carlo reference solutions except for a region near the z-axis, where ray effects occur. The method is two orders of magnitude faster than a reference Monte Carlo calculation but the use of a stochastic first-collision source limits performance. Williams et al. demonstrate a 3-D photon transport method, based on the method of characteristics, with the geometric flexibility to model both the accelerator head and patient/phantom geometry and to achieve results within uncertainty of a Monte Carlo reference solution [52]. Lacking the capability to perform electron transport however, the method is unable to calculate dose. Yuan et al. also demonstrate KERMA calculation using a 3-D photon-only solver, resulting in good agreement (within 2-3%) of EGSnrc except for errors of ~30% along the beam edge [53]. Boman et al. use the finite element method (FEM) to discretize space, angle, and energy variables for the coupled photon-electron transport equations, and they show that this can be applied to solve a simple 2-D test problem [10]. Tervo et al. build on this work to develop a method to solve the inverse treatment planning problem–calculating the source distribution that produces a desired dose distribution [46]. They demonstrate the inverse method using a small 2-D problem, but they state that solving a problem in 3-D would be "too computationally expensive."

Gifford et al. describe the use of a finite-element multigroup discrete-ordinates method, called Attila, for dose calculation in radiotherapy applications [19]. Attila was originally developed by Transpire Inc. for applications in nuclear engineering and solves the transport equation over an unstructured tetrahedral mesh. For external beam photon problems, Attila models electron transport with a Boltzmann-CSD equation. In this model, the continuous slowing down (CSD) operator is used to treat small frequent energy losses, and the Boltzmann scattering operator is used to treat large-angle or large-energy-loss scattering [25]. To mitigate ray effects, Attila employs a first-collision source method. Gifford et

al.'s results show that Attila agrees within 5% of the Monte Carlo code MCNPX [21] for a brachytherapy benchmark and within 2.2% of EGS4 [34] for the external beam benchmark published by Rogers and Mohan [39]. Vassiliev et al. test Attila's ability to calculate dose for two CT-based benchmarks: a prostate case and a head and neck case [49]. To create the benchmarks, each CT Hounsfield number is mapped to one of four materials with fixed densities. Each phantom is illuminated with several open 6 MV fields. For the head and neck case comparison with EGSnrc shows that of voxels receiving greater than 10% of the maximum dose, 98.6% pass the 3%/3 mm criterion. For the head and neck case, 99% of voxels passes the 3%/3 mm criterion. Calculation times for these cases are 16.1 minutes and 19.6 minutes on a single 2.2 GHz AMD Opteron.

Based on the success of Attila, Transpire Inc. developed Acuros, a re-write of Attila designed specifically for radiotherapy dose calculation [50]. One of the biggest differences between Attila and Acuros is that Acuros uses discontinuous finite elements defined on a cartesian grid rather than the unstructured tetrahedral elements employed by Attila. Vassiliev et al. evaluate Acuros using a benchmark similar to the Rogers and Mohan slab phantom although bone (ρ =1.85 g/cm^3) was substituted for aluminum (ρ =2.7 g/cm^3), and larger field sizes ($\geq$ 2.5 x 2.5 cm^2 compared to the original 1.5 x 1.5 cm^2) were used. They find that Acuros agrees with EGSnrc to within 2.3% when excluding the initial 2 cm build-up region, where distance-to-agreement is reported as less than 0.3 mm. Using a breast treatment plan with an anthropomorphic phantom, they find that, among voxels receiving greater than 10% of the maximum dose, 98.7% of voxels pass the 2%/1 mm criterion and that 99.9% of voxels pass the 2%/2 mm criterion. Bush et al. provide further comparison between Acuros and Monte Carlo [12]. They report larger discrepancies than in previous work, but it is difficult to determine the exact cause since the Acuros and Monte Carlo calculations differ in both the source model and in the CT-to-material assignment.

In general, the accuracy and efficiency of deterministic methods depends on the discretization employed and–in the case of charged particle transport–on the physical model

used. A finer discretization provides a better approximation, but it also increases the size of the linear system, which requires more computational effort to solve. Thus, as with stochastic methods, there is a trade-off between the quality of the result and the required computational resources. While the multigroup Boltzmann-CSD and Boltzmann-Fokker-Plank equations can be used to model electron transport [27, 25], they cannot provide physical detail at the level of a continuous-energy condensed-history Monte Carlo method such as EGSnrc [24].

2.3 The Coarse Mesh Transport Method (COMET)

COMET combines the accuracy and physical detail of Monte Carlo with the precision of deterministic transport theory. The method solves the transport equation by coupling together stochastically generated response functions using a deterministic solver. While this idea is relatively new to medical physics, it has been extensively validated for neutron transport in nuclear reactor physics [55]. Zhang and Rahnema extended the method to handle coupled photon and electron transport in 2-D [54], and showed that it can achieve efficiency nearly three orders of magnitude higher than Monte Carlo when tested on stylized benchmark problems. The method was extensively benchmarked by Blackburn [9] for a large variety of test cases.

While accurate and efficient, the 2-D version of COMET-PE was limited to solve only a small class of problems. Aside from restriction to 2 spatial dimensions, the implementation was based around a simplistic source model that was difficult to adapt for realistic cases. The source was required to be mono-energetic, mono-directional, and spatially flat. The method did treat uncollided and collided photon fluence separately, however the implementation relied on response functions to transport the uncollided fluence, an approach that cannot be easily extended to more general source models. Additionally, Blackburn reported long calculation times for problems with a large number of meshes [9], a problem that is worse in 3 dimensions. This work addresses all of these issues and shows that

COMET-PE can attain accuracy comparable to Monte Carlo for strongly heterogeneous benchmarks with realistic source distributions in 3-D.

CHAPTER III

METHOD OVERVIEW

3.1 Introduction to COMET

The key idea underlying the COMET method is that of the response function. A response function describes how a single mesh responds to internal and external source distributions. A mesh is just a small (local) piece of a larger (global) problem. Response functions are general in the sense that they can be applied even when the source distributions are not known ahead of time. This is depicted in figure 1.

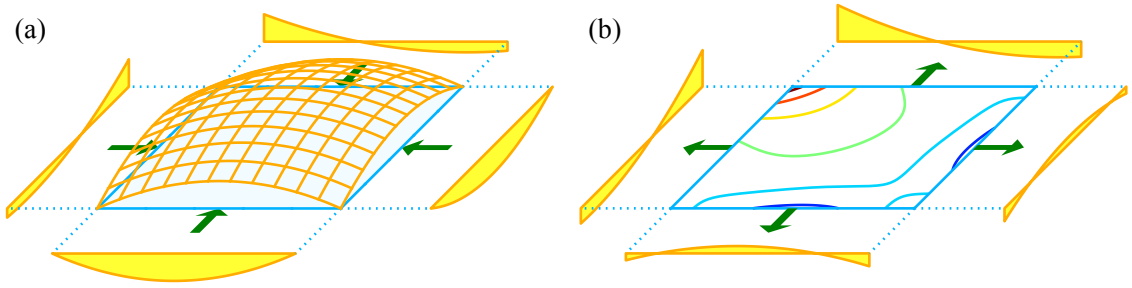

Figure 1: The action of a response function. Part (a) illustrates source distributions; the wireframe depicts an internal (volumetric) source, and the four solid curves on each side depict external (surface) sources. Part (b) shows the response to these sources. Contour lines within the mesh represent energy deposition, and the solid curves on each side represent outgoing fluence.

Each response function is a combination of solutions to many simpler problems. These simpler problems fall into two categories: surface-to-volume responses and volume-to-volume responses. Surface-to-volume responses are the result of imposing a unit incoming fluence (or surface source) on a single surface of a mesh. This process is repeated for each member of a set of basis functions as depicted in figure 2 on the next page. Similarly volume-to-volume responses are computed from a set of basis functions that correspond to volumetric source distributions as in figure 3 on page 17. In both the surface-to-volume

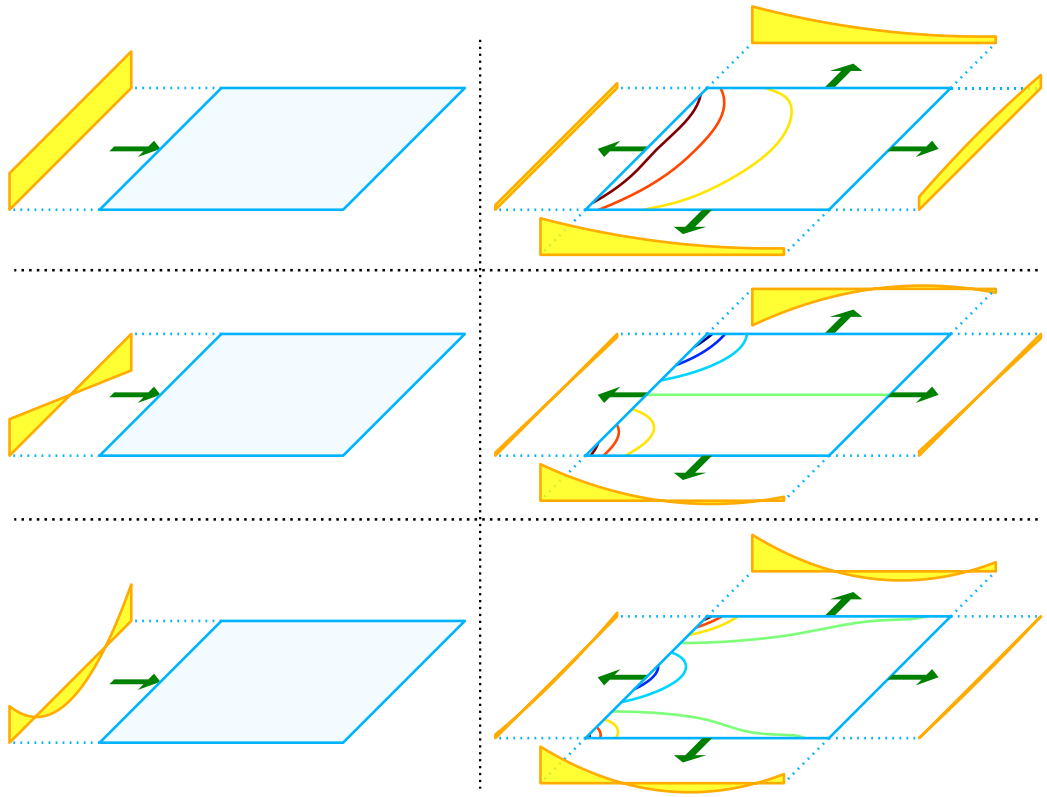

Figure 2: Surface-to-volume responses. The column on the left depicts surface sources (incoming fluence distributions) chosen from a basis set. The right column depicts the corresponding responses. Contour lines within the mesh represent energy deposition, and the solid curves represent outgoing fluence.

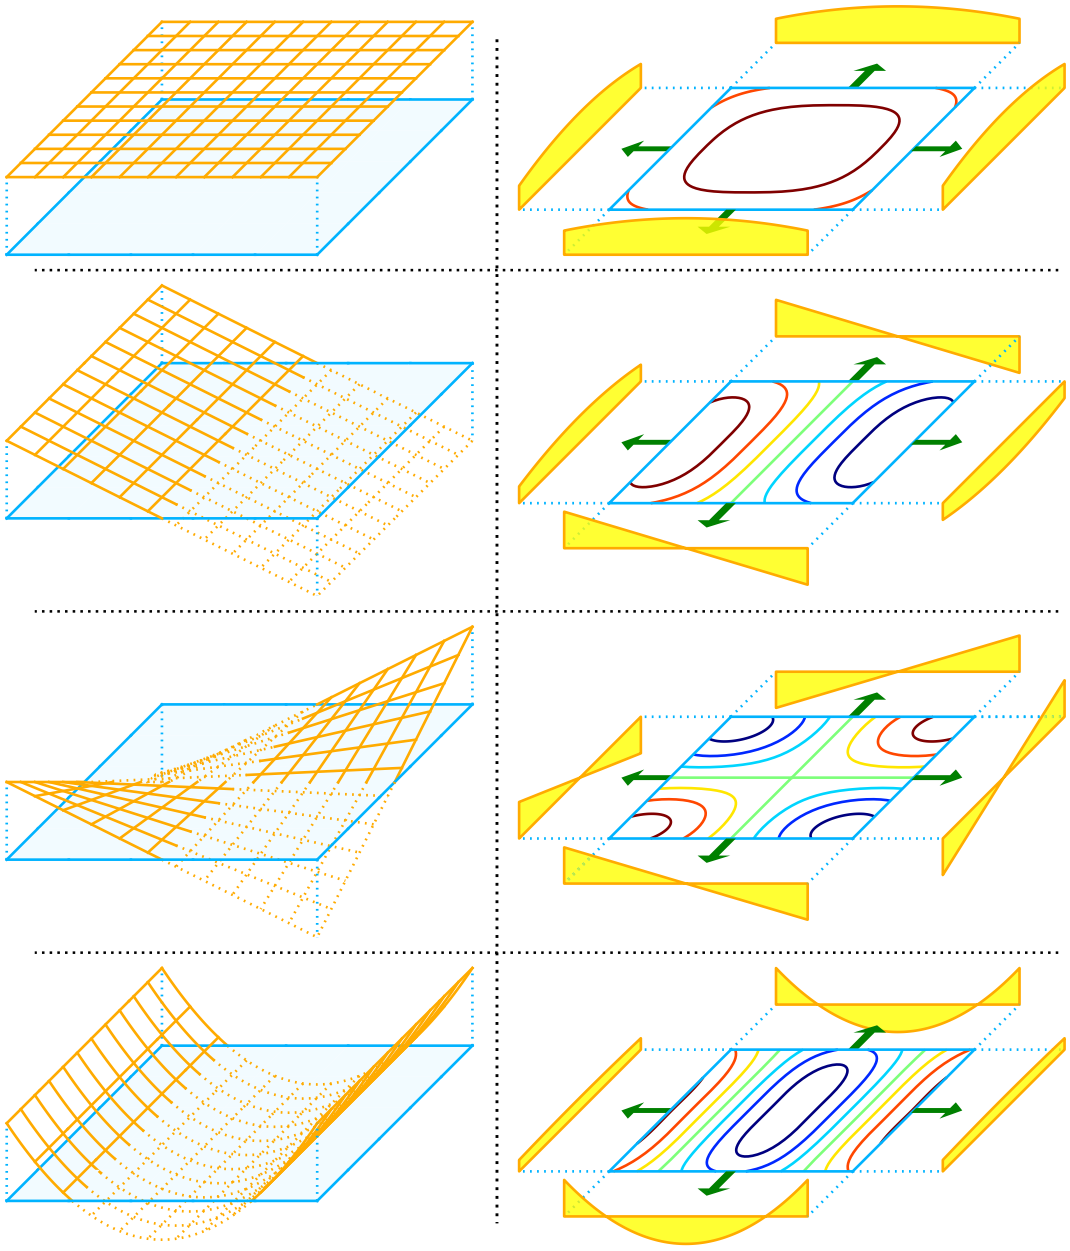

Figure 3: Volume-to-volume responses. The column on the left depicts volumetric sources chosen from a basis set. The right column depicts the corresponding responses. The contour lines within each mesh represent energy deposition, and the solid curves represent outgoing fluence.

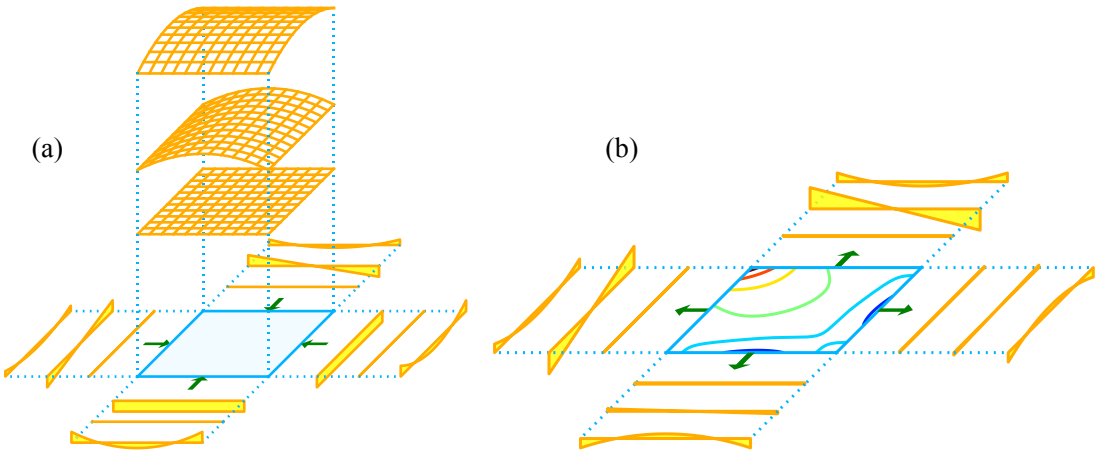

Figure 4: The response function as a superposition. Part (a) shows the expansion of each surface and volume source from figure 1(a). Part (b) shows the response; contour lines represent energy deposition, and the three solid curves on each side represent an expansion of the outgoing fluence (shown originally in figure 1(b)).

and volume-to-volume response problems, the sources are known ahead of time since they correspond to a fixed basis set. This allows the solutions to be pre-computed and stored in a library. Once pre-computation is complete, the response to arbitrary surface and volume sources can be found by superposition of the pre-computed responses as shown in figure 4. The same basis functions are used to expand both incoming and outgoing fluence on the mesh surfaces. This allows response functions to be easily connected to one another. By connecting response functions, one can solve larger (global) problems.

The first step in solving a global problem is calculating the response to the fixed source. In general, the source distribution may consist of an external source (as in external beam radiotherapy), an internal source (as in brachytherapy), or both. Furthermore, any external source distribution can be represented directly, as an incoming fluence, or indirectly, as the source resulting from the distribution of first-collisions. These options are depicted in figure 5 on the following page.

Once the source response has been computed, the global solution is found by iteratively refining the fluence distributions that connect the meshes. Initially, the fluence distributions

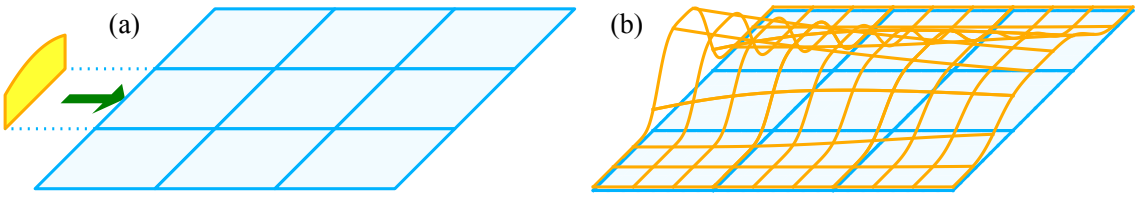

Figure 5: Possible source distributions for a hypothetical 9-mesh problem. An external source can be represented by (a) an incoming fluence distribution on an external surface or by (b) the volumetric distribution of first-collisions.

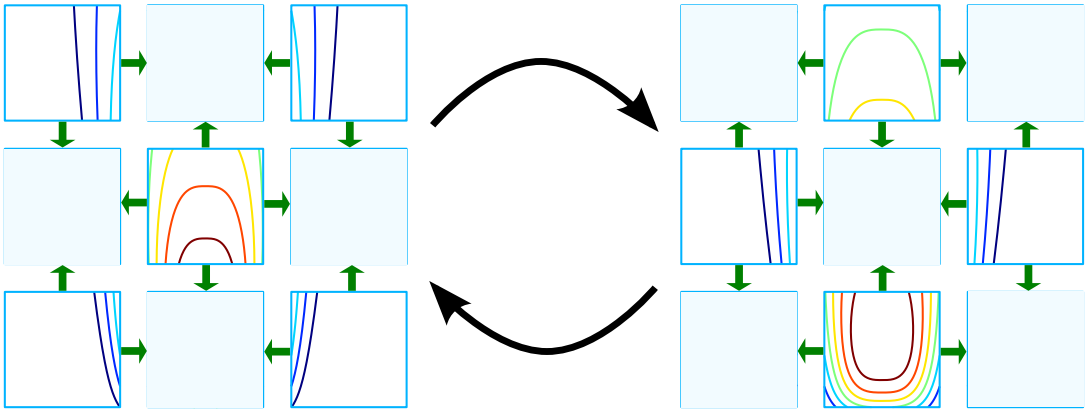

Figure 6: The fluence distribution is determined by iteratively applying the response functions to update the interface conditions. This diagram depicts a red-black Gauss-Seidel iteration in which the fluence updates occur in an alternating checkerboard pattern.

are discontinuous at the mesh interfaces. As the iteration (depicted in figure 6) progresses, the interface fluences converge to a continuous distribution. Once the solution is converged, the resulting energy distribution can be computed almost instantaneously. Given this conceptual basis, a more detailed mathematical description is now presented.

3.2 The Boltzmann Transport Equation

The COMET method solves the Boltzmann transport equation

$$\hat{\Omega}\cdot\nabla\psi+\sigma_t(\mathbf{r},E)\,\psi-\int_0^\infty dE'\int_{4\pi} d\hat{\Omega}'\sigma_s(\mathbf{r},\hat{\Omega},E;\hat{\Omega}',E')\,\psi=q(\mathbf{r},\hat{\Omega},E),\quad \mathbf{r}\in V \qquad (6)$$

with boundary condition

$$\psi(\mathbf{r}, \hat{\Omega}, E) = \gamma(\mathbf{r}, \hat{\Omega}, E), \quad \mathbf{r} \in \partial V \quad \text{and} \quad \hat{n} \cdot \hat{\Omega} < 0. \tag{7}$$

where $\psi(\mathbf{r}, \hat{\Omega}, E)$ is the differential energy fluence (cm^{-2}), which is a function of position $\mathbf{r}$ (cm), angle $\hat{\Omega}$, and energy E (MeV). Hereinafter, *fluence* shall be used to mean *differential energy fluence* unless otherwise noted. The spatial domain is a volume V, which has a boundary ∂V with outward unit normal vector $\hat{n}$. Fluence originates from either a source internal to V or from a source external to V. Internal sources are represented by a term on the right hand side of the transport equation q (with units of cm^{-3}), and external sources are represented by the boundary condition γ (cm^{-2}). The Boltzmann transport equation is a balance equation. The first term $\hat{\Omega} \cdot \nabla$ gives the net energy of particles streaming out of the differential phase space $\text{d}\mathbf{r}\text{d}\hat{\Omega}\text{d}E$. The second term describes the removal of particles from the phase space due to collision interactions, and σ_t (cm^{-1}) is the total cross section. The third term describes the distribution of particles emerging from scattering interactions, and σ_s ($\text{cm}^{-1}\text{MeV}^{-1}$) is the double differential scattering cross section. Equation (6) is a balance of energy sources and sinks within the differential phase space volume. The fluence in a given volume is uniquely determined by the material properties, which are given by σ_t and σ_s, and by the source distributions q and γ.

3.3 Spatial Decomposition and Orthogonal Function Expansion

The COMET method employs a divide-and-conquer approach to solve (6) and (7). This strategy is based on the idea that solutions of several small problems, called response functions, can be combined to solve a larger problem. Assume that volume V has been divided into non-overlapping meshes V_i such that $\bigcup_i V_i = V$. Let ∂V_i denote the boundary of mesh V_i, and let ∂V_{is} denote surface s of mesh i. Let $\hat{n}_i$ denote the outward unit normal on ∂V_i and let $\hat{n}_{is}$ denote the outward unit normal on ∂V_{is}. For simulating patient geometry, these meshes correspond to cubes arranged in a Cartesian grid. Let $\{Q^m\}$ be a set of orthogonal basis functions that can be used to represent an internal source within the volume V_i, and let

$\{\Gamma^m\}$ be a set of orthogonal basis functions that can be used to represent an external source impinging on the surface ∂V_{is}. Each member of each basis corresponds to a response function. The volume-to-volume response function R_i^m is the fluence distribution within mesh V_i that results from the internal source Q^m. Thus it satisfies

$$\hat{\Omega} \cdot \nabla R_i^m + \sigma_t(\mathbf{r}, E)\, R_i^m - \int_0^\infty \mathrm{d}E' \int_{4\pi} \mathrm{d}\hat{\Omega}' \sigma_s(\mathbf{r}, \hat{\Omega}, E; \hat{\Omega}', E')\, R_i^m = Q^m(\mathbf{r}, \hat{\Omega}, E), \quad \mathbf{r} \in V_i \tag{8}$$

with the vacuum boundary condition

$$R_i^m(\mathbf{r}, \hat{\Omega}, E) = 0, \quad \mathbf{r} \in \partial V_i \quad \text{and} \quad \hat{n}_i \cdot \hat{\Omega} < 0. \tag{9}$$

Similarly, the surface-to-volume response function R_{is}^m is the fluence distribution within mesh V_i that results from the external source Γ^m impinging on surface s. Thus it satisfies

$$\hat{\Omega} \cdot \nabla R_{is}^m + \sigma_t(\mathbf{r}, E)\, R_{is}^m - \int_0^\infty \mathrm{d}E' \int_{4\pi} \mathrm{d}\hat{\Omega}' \sigma_s(\mathbf{r}, \hat{\Omega}, E; \hat{\Omega}', E')\, R_{is}^m = 0, \quad \mathbf{r} \in V_i \tag{10}$$

with boundary condition

$$R_{is}^m(\mathbf{r}, \hat{\Omega}, E) = \begin{cases} \Gamma^m(\mathbf{r}, \hat{\Omega}, E), & \mathbf{r} \in \partial V_{is} \quad \text{and} \quad \hat{n}_{is} \cdot \hat{\Omega} < 0 \\ 0, & \mathbf{r} \in (\partial V_i - \partial V_{is}) \quad \text{and} \quad \hat{n}_i \cdot \hat{\Omega} < 0. \end{cases} \tag{11}$$

Both types of response functions are defined as if mesh i exists in a vacuum so that:

1. the only way a particle can enter the mesh is to be emitted by the source, and

2. once a particle leaves the mesh it cannot return.

Using the linearity of the transport equation, these response functions can be superimposed to reconstruct the solution

$$\psi = \sum_m^\infty h_i^m R_i^m + \sum_s \sum_m^\infty j_{is}^{m-} R_{is}^m \tag{12}$$

$$h_i^m = \langle Q^m, q \rangle_i, \quad j_{is}^{m-} = \langle \Gamma^m, \psi \rangle_{is}^-, \quad \mathbf{r} \in V_i. \tag{13}$$

Here, the volume and surface inner products are defined

$$\langle f, g \rangle_i = \int_{V_i} \mathrm{d}\mathbf{r} \int_{4\pi} \mathrm{d}\hat{\Omega} \int_0^\infty \mathrm{d}E \; f(\mathbf{r}, \hat{\Omega}, E) g(\mathbf{r}, \hat{\Omega}, E) \tag{14}$$

$$\langle f, g \rangle_{is}^{\pm} = \int_{\partial V_{is}} \mathrm{d}\mathbf{r} \int_{\pm \hat{n}_{is} \cdot \hat{\Omega} > 0} \mathrm{d}\hat{\Omega} \int_0^\infty \mathrm{d}E \left| \hat{n}_{is} \cdot \hat{\Omega} \right| f(\mathbf{r}, \hat{\Omega}, E) g(\mathbf{r}, \hat{\Omega}, E). \tag{15}$$

The $-$ superscript indicates incoming fluence and the $+$ superscript, outgoing fluence. The coefficients h_m^i and j_{is}^{m-} are referred to as expansion coefficients. When the basis sets $\{Q^m\}$ and $\{\Gamma^m\}$ are complete, equations (12) and (13) provide an exact solution to (6) and (7). For a practical implementation, however, at least two changes are required. First, the infinite series expansions in (12) must be truncated at finite order. Second, the incoming fluence expansion coefficients j_{is}^{m-} must be found using a method other than the relation given by (13) since the fluence ψ is not know a priori.

3.4 Solving for the Interface Fluence Expansion Coefficients

The fluence expansion coefficients can be computed by solving a linear system that arises from a fluence continuity constraint between neighboring meshes. If s is an external surface ($\partial V_{is} \subset \partial V$), the external boundary condition γ in (7) gives the incoming fluence. Otherwise s must be an internal surface ($\partial V_{is} \subsetneq \partial V$), and the continuity condition in (16) determines the incoming fluence

$$\psi_{is}^-(\mathbf{r}, \hat{\Omega}, E) = \psi_{\tilde{is}}^+(\mathbf{r}, \hat{\Omega}, E), \quad \mathbf{r} \in \partial V_{is} \quad \text{and} \quad \hat{n}_{is} \cdot \hat{\Omega} < 0 \tag{16}$$

where the $+$ and $-$ superscripts denote incoming and outgoing fluence respectively. The mesh index $\tilde{i}$ and the surface index $\tilde{s}$ define neighbors to surface s of mesh i such that $\partial V_{is} = \partial V_{\tilde{is}} = (\partial V_i \cap \partial V_{\tilde{i}})$. Equation (16) simply states that energy flows directly from one mesh into its neighbor. So that both the incoming fluence and the outgoing fluence have the same form, the outgoing fluence is expanded in terms of the same basis used for the incoming fluence. This yields outgoing fluence expansion coefficients j_{is}^{m+} . With interface

fluence defined this way, the interface continuity condition given by (16) is satisfied if and only if the coefficients satisfy

$$j_{is}^{m-} = j_{\widetilde{is}}^{m+}. \tag{17}$$

Using this relationship, one can write the outgoing fluence expansion coefficients j_{is}^{m+} as a sum of contributions from the source, from the boundary condition, and from neighboring meshes

$$j_{is}^{m+} = \sum_{m'} h_i^{m'} R_{i\to s}^{m'\to m} + \sum_{\substack{m',s' \\ \partial V_{is'} \subset \partial V}} \gamma_{is'}^{m'} R_{is'\to s}^{m'\to m} + \sum_{\substack{m',\ s' \\ \partial V_{is'} \subsetneq \partial V}} j_{\widetilde{is'}}^{m'+} R_{is'\to s}^{m'\to m}. \tag{18}$$

Where the volume-to-surface and surface-to-surface response coefficients are defined

$$\begin{aligned} R_{i\to s}^{m'\to m} &= \left\langle \Gamma^m, R_i^{m'} \right\rangle_{is}^{+} \\ R_{is'\to s}^{m'\to m} &= \left\langle \Gamma^m, R_{is'}^{m'} \right\rangle_{is}^{+}, \end{aligned} \tag{19}$$

and where the external boundary condition has also been expanded in coefficients $\gamma_{is'}^{m'} = \left\langle \Gamma^{m'}, \gamma \right\rangle_{is'}^{-}$. Equation(18) can be written in matrix form as

$$\mathbf{j} = \mathbf{R}_h \mathbf{h} + \mathbf{R}_\gamma \mathbf{g} + \mathbf{R}\mathbf{j}. \tag{20}$$

Because the vectors $\mathbf{h}$ and $\mathbf{g}$ are computed directly from the given source and external boundary conditions, the vector $\mathbf{b} = \mathbf{R}_h\mathbf{h} + \mathbf{R}_\gamma\mathbf{g}$ is defined to represent the sum of all (fixed) source contributions to the interface fluence. By grouping the terms involving $\mathbf{j}$ on the left hand side and defining $\mathbf{A} = \mathbf{I} - \mathbf{R}$, one finds the canonical linear system

$$\mathbf{A}\mathbf{j} = \mathbf{b}.$$

This system has a block-sparse structure, and by taking advantage of this, a careful implementation of the COMET matrix-vector product has speedup that scales linearly with the number of processors in a distributed-memory computer. In this work, the linear system is solved using a red-black Gauss-Seidel iteration [7].

3.5 Energy Deposition

Once the interface fluence is known, dose calculation is straightforward. As with the outgoing angular fluence, the distribution of energy deposited $D_i(\mathbf{r})$ can be written as a response to contributions from the source and from the incoming fluence

$$D_i(\mathbf{r}) = \sum_m h_i^m R_{i\to D}^m(\mathbf{r}) + \sum_{m,\ s} j_{is}^{m-} R_{is\to D}^m(\mathbf{r}). \tag{21}$$

Here the energy deposition response coefficients are defined

$$\begin{aligned} R_{i\to D}^m(\mathbf{r}) &= \int_0^\infty \mathrm{d}E \int_{4\pi} \mathrm{d}\hat{\Omega} \frac{\sigma_{ED}(\mathbf{r}, E)}{\rho(\mathbf{r})} R_i^m(\mathbf{r}, \hat{\Omega}, E) \\ R_{is\to D}^m(\mathbf{r}) &= \int_0^\infty \mathrm{d}E \int_{4\pi} \mathrm{d}\hat{\Omega} \frac{\sigma_{ED}(\mathbf{r}, E)}{\rho(\mathbf{r})} R_{is}^m(\mathbf{r}, \hat{\Omega}, E) \end{aligned} \tag{22}$$

where σ_{ED} is an energy deposition cross section (cm^{-1}) and ρ is the material density ($\mathrm{g/cm}^3$). The spatial dependence of $D_i(\mathbf{r})$ could be represented by expansion in a continuous basis, but in this work, a discrete basis is used. For example, the method is tested using mesh sizes of 5 mm and 10 mm, but in both cases, energy deposition is calculated at a resolution of 1 mm. COMET-PE provides a detailed distribution of the energy deposition within each mesh, whether using a continuous or a discrete basis.

3.6 Coupled Transport

The method works just as well for photons as it does for electrons and neutrons, or theoretically, for protons or heavy ions. Because the method constructs solutions via superposition, the only requirement is a linear transport process. This linearity can also be used to generalize the method for multi-species transport. Instead of solving a single transport equation, one must solve a system of transport equations: one for each particle type. These transport equations are coupled to one another through their source terms. For a given species ν, the source term can be written

$$q_\nu = q_{\nu 0} + \sum_{\nu'} q_{\nu\nu'}$$

where $q_{\nu 0}$ is the primary source of particles of type ν, and $q_{\nu\nu'}$ is the secondary source that results from collisions made by particles of type ν'. To accommodate this generalization, only a small change in notation is required: replace all occurrences of the index m with the pair of indices $m\nu$. For example, the surface-to-surface response coefficient $R_{is'\to s}^{m'\to m}$ becomes $R_{is'\to s}^{m'\nu'\to m\nu}$. In other words, adding multi-species transport is equivalent to adding an extra dimension to the expansion bases. For megavoltage photon therapy, the index ν will take on one of two values corresponding to either photons or electrons. Although positrons are also generated by pair production, explicit positron transport is not essential for dose calculation. Instead of including a transport equation for positrons, pair production and annihilation are replaced with a single interaction that results in two photons and two electrons. Since human tissue is composed of materials with relatively low atomic number, this approximation does not significantly perturb the dose distribution.

CHAPTER IV

ANALYTICAL FIRST COLLISION SOURCE

4.1 Analytical First-Collision Source

The highly anisotropic photon sources typical of medical physics applications–specifically external beam radiotherapy–necessitate high expansion orders. To overcome this challenge, the source term can be treated analytically. Let us decompose the fluence into an uncollided component ψ^u and a collided component ψ^c such that $\psi = \psi^u + \psi^c$. The uncollided fluence represents particles that have not experienced any collisions since being emitted from the source; the collided fluence represents everything else. Because the transport equation in ψ^u lacks scattering, it can be solved analytically using an integrating factor $\tau(\mathbf{r}, \mathbf{r}', E)$, known as the optical depth, that gives the distance between points $\mathbf{r}$ and $\mathbf{r}'$ in mean free paths. The uncollided fluence can be written

$$\psi^u(\mathbf{r}, \hat{\Omega}, E) = \int_V \mathrm{d}\mathbf{r}' \frac{q(\mathbf{r}', \hat{\Omega}, E)}{4\pi|\mathbf{r} - \mathbf{r}'|^2} \exp\left[-\tau(\mathbf{r}, \mathbf{r}', E)\right] \delta\left(\hat{\Omega} - \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|}\right) \tag{23}$$

where δ is the Dirac delta. The transport equation in ψ^c is identical to (6) except that the original source q is replaced by the first scatter source

$$q^c(\mathbf{r}, \hat{\Omega}, E) = \int_0^\infty \mathrm{d}E' \int_{4\pi} \mathrm{d}\hat{\Omega}' \sigma_s(\mathbf{r}, \hat{\Omega}, E; \hat{\Omega}', E')\, \psi^u(\mathbf{r}, \hat{\Omega}', E'). \tag{24}$$

The combination of (23) and (24) acts as a low pass filter, so that the first scatter source q^c has a broader distribution in space and angle than the original source q. This means it can be better approximated by a low order expansion. In the case of a point source, the uncollided fluence can be evaluated by ray-tracing. To use this technique more generally, an arbitrary source can be approximated by a sum of point sources; this is analogous to evaluating the integral in (23) via quadrature. To incorporate the new analytical source into

the COMET method, let us expand ψ^u in an orthogonal basis $\{\tilde{Q}^m\}$ such that

$$\psi^u = \sum_m^\infty \tilde{h}_i^m \tilde{Q}^m \quad \text{with} \quad \tilde{h}_i^m = \left\langle \tilde{Q}^m, \psi^u \right\rangle_i, \quad \mathbf{r} \in V_i. \tag{25}$$

By combining (24) and (25) one can derive the form of the first-scatter distribution Q^m that results from the uncollided fluence $\tilde{Q}^m$

$$Q^m(\mathbf{r}, \hat{\Omega}, E) = \int_0^\infty \mathrm{d}E' \int_{4\pi} \mathrm{d}\hat{\Omega}' \sigma_s(\mathbf{r}, \hat{\Omega}', E'; \hat{\Omega}, E)\, \tilde{Q}^m(\mathbf{r}, \hat{\Omega}', E').$$

The resulting basis $\{Q^m\}$ is not orthogonal, so the coefficients h_i^m cannot be used as they are defined in (13). Instead, they are replaced with the coefficients $\tilde{h}_i^m$ as defined in (25). These new coefficients can be used to synthesize the first-scatter source

$$q^c = \sum_m^\infty \tilde{h}_i^m Q^m.$$

The inner product in (25) implies that a 6-dimensional integral must be evaluated for each coefficient $\tilde{h}_i^m$. The space-angle delta in (23) obviates the angular integral, and the multigroup method [26] is used to handle the energy dependence of the uncollided fluence. More detail on the multigroup method is provided in the next section. This leaves a triple integral in the spatial coordinate over the mesh volume. These remaining integrals are evaluated using a purpose-built adaptive cubature. The cubature algorithm is inspired by the work of Berntsen et al. [8], but it is designed specifically to meet the requirements of the COMET-PE method in evaluating integrals given by the inner product in (25). The optical depth τ is evaluated using a variant of the Siddon algorithm [42].

4.2 Multigroup Transport Theory

In the case of a photon point source $q(\mathbf{r}, \hat{\Omega}, E) = \delta(\mathbf{r} - \mathbf{r}_0) q_0(\hat{\Omega}, E)$, (23) becomes

$$\psi^u(\mathbf{r}, \hat{\Omega}, E) = \frac{q_0(\hat{\Omega}, E)}{4\pi |\mathbf{r} - \mathbf{r}_0|^2} \exp\left[-\tau(\mathbf{r}, \mathbf{r}_0, E)\right] \delta\left(\hat{\Omega} - \frac{\mathbf{r} - \mathbf{r}_0}{|\mathbf{r} - \mathbf{r}_0|}\right). \tag{26}$$

in continuous energy. The multigroup energy approximation is introduced by integrating the uncollided fluence

$$\psi_g^u(\mathbf{r}, \hat{\Omega}) = \int_{E_{g+1}}^{E_g} \mathrm{d}E\ \psi^u(\mathbf{r}, \hat{\Omega}, E)$$

over some energy group structure $E_1 > E_2 > \cdots > E_{G+1}$. The result, ψ_g^u, is called the multigroup energy fluence. Integrating the right hand side of (26) gives

$$\psi_g^u(\mathbf{r}, \hat{\Omega}) = \frac{q_{0g}(\hat{\Omega})}{4\pi|\mathbf{r} - \mathbf{r}_0|^2} \exp\left[-\tau_g(\mathbf{r}, \mathbf{r}_0)\right] \delta\left(\hat{\Omega} - \frac{\mathbf{r} - \mathbf{r}_0}{|\mathbf{r} - \mathbf{r}_0|}\right). \tag{27}$$

The corresponding multigroup source, total cross section, and optical depth are defined

$$q_{0g}(\hat{\Omega}) = \int_{E_{g+1}}^{E_g} \mathrm{d}E \; q_0(\hat{\Omega}, E)$$

$$\sigma_{tg}(\mathbf{r}) = \frac{\int_{E_{g+1}}^{E_g} \mathrm{d}E \int_{4\pi} d\hat{\Omega}\, \sigma_t(\mathbf{r}, E) \; f(\mathbf{r}, \hat{\Omega}, E)}{\int_{E_{g+1}}^{E_g} \mathrm{d}E \int_{4\pi} d\hat{\Omega} \; f(\mathbf{r}, \hat{\Omega}, E)}$$

$$\tau_g(\mathbf{r}, \mathbf{r}_0) = \int_0^{|\mathbf{r}-\mathbf{r}_0|} \mathrm{d}s\, \sigma_{tg}\left(\mathbf{r}_0 + s\frac{\mathbf{r} - \mathbf{r}_0}{|\mathbf{r} - \mathbf{r}_0|}\right).$$

When the spectral weighting function $f(\mathbf{r}, \hat{\Omega}, E) = \psi^u(\mathbf{r}, \hat{\Omega}, E)$ is used, equation (27) is exact. Because the uncollided fluence is not known a priori, the problem-independent, flat energy spectrum $f(\mathbf{r}, \hat{\Omega}, E) = E$ is used.

The degree of approximation introduced by using multigroup cross sections can be controlled by selecting an appropriate group structure. A finer group structure will result in a more accurate calculation, but computational cost increases when the number of groups G increases. One would like to use the coarsest group structure that preserves accuracy. In general some sort of analysis or sensitivity study is required to choose an optimal group structure. Such an analysis is performed in section 6.5.1.

4.2.1 Cross Section Evaluation

A user code, called GETXS, is developed to extract continuous energy photon cross sections from EGSnrc. This ensures the consistency with response coefficients, which are also generated using an EGSnrc user code. Additionally, this procedure eliminates a possible source of discrepancy between COMET-PE and DOSXYZnrc, which is yet another user code of EGSnrc.

CHAPTER V

EXPANSION BASES AND PRE-COMPUTATION

5.1 Response Pre-computation

So far only the deterministic solution method has been discussed, but this method presupposes the existence of response functions. One would like to have these functions readily available for each mesh in the problem, so that the solution can be synthesized quickly. This implies that the response functions must be generated using only patient-independent knowledge. To accomplish this, a homogeneous material distribution is assumed within each mesh. Although COMET can handle mesh heterogeneity without approximation, the number of possible coarse mesh configurations is unmanageably large. It would be difficult to compute a library with response functions covering an exhaustive set of heterogeneous meshes. Homogeneous meshes, however, allow one to construct a patient volume one material at a time. Additionally, the use of homogenous meshes allows one to take advantage of the cube's symmetry, which increases the efficiency of both the response function generation and the iterative solution. One must pre-compute response functions for each unique material that it to be modeled. Voxel phantoms derived from computed tomography require a continuum of materials. Although such phantoms are not studied here, one should note that COMET can effectively represent a continuous variation in material density ρ by interpolation between response functions

$$R^{m'\to m}_{is'\to s}(\rho) = \frac{\rho-\rho_1}{\rho_2-\rho_1}R^{m'\to m}_{is'\to s}(\rho_1) + \frac{\rho_2-\rho}{\rho_2-\rho_1}R^{m'\to m}_{is'\to s}(\rho_2)$$

that have been pre-computed for a set of fixed densities $\{\rho_1, \rho_2, \ldots\}$. This type of interpolation is used routinely and effectively for nuclear reactor calculations [55].

The use of a pre-computed response library enables COMET-PE to be both accurate and efficient. The accuracy comes from the fact that pre-computation can take advantage of the

best transport theory methods available, whether stochastic or deterministic. The efficiency comes from the fact that this pre-computation need only be performed once. Since a single library can be used for all patients, the amortized cost is practically negligible. When Monte Carlo is used for response generation, this allows much higher precision than is practical in patient-specific calculations. Even without amortization, the pre-computations are efficient because the small–in terms of mean free path–meshes combined with vacuum boundary conditions result in problems that converge quickly (i.e. small number of histories for a stochastic method or small number of iterations for a deterministic method). Furthermore the pre-computation is trivially parallelizable.

5.2 Orthonormal Surface Basis

The choice of response basis $\{\Gamma^m\}$ is key to both the accuracy and the efficiency of the method. The only approximation that has been made so far is the replacement of infinite series by their partial sums. The method's accuracy is therefore closely tied to the ability of each basis to provide a good approximation with only a small number of terms. The choice of basis also determines the efficiency of the method. Assuming that the problem contains N meshes, the interface fluence expansion contains M terms, and the method requires K iterations for convergence, the number of floating point operations will be $\mathrm{O}\left(M^2NK\right)$.

Based on her experience with COMET-PE in 2-D, Blackburn [9] suggested that the surface expansion should be improved for both energy and angular variables. For the original 2-D implementation of COMET-PE, each phase space variable on the mesh surface–one spatial dimension, the cosine of the polar angle, the azimuthal angle, and energy–were all expanded in Legendre Polynomials [9]. In further development by Zhang and Rahnema, the azimuthal angle expansion was replaced with a double Legendre polynomial to allow a discontinuity between fluence traveling in the forward and backward directions, which enabled improved treatment of problems with forwardly peaked angular distributions [54]. Since the original development of COMET-PE, Zhang and Rahnema [55] have developed

an improved orthogonality condition:

$$\langle f, g \rangle_{is}^{\pm} = \int_{\partial V_{is}} \mathrm{d}\mathbf{r} \int_{\pm\hat{n}_{is}\cdot\hat{\Omega}>0} \mathrm{d}\hat{\Omega} \int_0^\infty \mathrm{d}E \left|\hat{n}_{is}\cdot\hat{\Omega}\right| f(\mathbf{r},\hat{\Omega},E) g(\mathbf{r},\hat{\Omega},E) \tag{28}$$

$$= \delta_{mm'} \tag{29}$$

where $\delta_{mm'}$ is the Kronecker delta. A basis that satisfies this condition conserves planar fluence (i.e. the net flow of energy across mesh interfaces) exactly regardless of truncation order. Additionally, when the 0^{th} order function is a constant, isotropic fluence is represented exactly by the 0^{th} order expansion. Neither of the expansions previously used for COMET-PE satisfy the new orthogonality condition. The basis used by Zhang and Rahnema for neutron transport [55] is a product of Legendre polynomials and Chebyshev polynomials. While this basis meets the new orthogonality requirement, it is not symmetric. Use of this basis in COMET-PE would prevent the underlying symmetry of a homogeneous mesh from being fully exploited. Instead, the following basis is used for the suface fluence expansion

$$\Gamma^{i_x i_y \ell k g}(\mathbf{r},\hat{\Omega},E) = P_{i_x}(x) P_{i_y}(y)\, H_{\ell k}(\hat{\Omega})\, \chi_g(E)$$

where P_{i_x} and P_{i_y} are Legendre polynomials of order i_x and i_y, and $H_{\ell k}$ is the hemispherical harmonic function of degree ℓ and order k with $|k| < \ell$. The hemispherical harmonic expansion allows the basis to satisfy the new orthonormality condition, and it is rotationally invariant about the surface normal. Thus, it does not break any of the symmetries of the cube. More detail on the hemispherical harmonics, including a novel recurrence relation, can be found in section 5.3. Equation (30) defines the energy binning function χ_g

$$\chi_g(E) = \begin{cases} E, & E_{g+1} \le E < E_g \\ 0, & \text{otherwise,} \end{cases} \tag{30}$$

$$g \in \{1,\ldots,G\}, \quad E_1 > E_2 > \cdots > E_{G+1}.$$

This is distinct from the multigroup approximation because all interactions within the mesh occur in continuous energy; particles are only "binned" as they cross a surface. Compared

with the Legendre Polynomial expansion, energy bins have the advantage of allowing the same library to be used for several beam energies without the artifacts induced by fitting a polynomial to a discontinuous function. The truncation of the series is specified by maximum expansion orders in space $M_{\mathbf{r}}$ and angle $M_{\hat{\Omega}}$ such that $i_x + i_y \leq M_{\mathbf{r}}$ and $\ell \leq M_{\hat{\Omega}}$. The additional requirement that $i_x + i_y + \ell \leq \max(M_{\mathbf{r}}, M_{\hat{\Omega}})$ removes high order cross terms.

5.3 Hemispherical Harmonics

Hemispherical harmonics presented here are due to Makhotkin [31]. (Note that this basis is distinct from that of Gautron et al. [18], which does not have the same orthonormality relation.) As defined here, the hemispherical harmonics differ only from Makhotkin's by a constant normalization factor. The basis is both orthonormal (see 5.3.1) and complete (see 5.3.2). Let the hemispherical harmonics be defined as

$$H_{\ell k}(\theta, \phi) = \begin{cases} \frac{2}{\sqrt{\pi}} J_\ell^{(-k)}(2\cos\theta - 1)\sin(-k\phi), & k < 0 \\ \frac{\sqrt{2}}{\sqrt{\pi}} J_\ell^{(0)}(2\cos\theta - 1), & k = 0 \\ \frac{2}{\sqrt{\pi}} J_\ell^{(k)}(2\cos\theta - 1)\cos(k\phi), & k > 0 \end{cases} \tag{31}$$

for $\ell \geq 0$ and $|k| \leq \ell$ where $J_\ell^{(k)}(x)$ is a normalized adjoint Jacobi function. The adjoint Jacobi functions play a role analogous to that of the associated Legendre polynomials used to define the spherical harmonics. The adjoint Jacobi functions can be defined by

$$J_\ell^{(k)}(x) = c_{\ell k}(1 - x^2)^{k/2} \frac{\mathrm{d}^k}{\mathrm{d}x^k} P_\ell^{(0,1)}(x). \tag{32}$$

where $c_{\ell k}$ is a constant factor chosen so that

$$\int_{-1}^{1} \mathrm{d}x\, (1 + x) \left(J_\ell^{(k)}(x) \right)^2 = 1. \tag{33}$$

In (32), $P_\ell^{(0,1)}(x)$ is a Jacobi polynomial. The Jacobi polynomials are orthogonal on the interval [-1,1] with respect to the weighting function $w_{(\alpha,\beta)}(x) = (1-x)^\alpha (1+x)^\beta$. The Jacobi polynomials generalize both the Legendre polynomials and the Chebyshev polynomials [45]. They can be defined by

$$P_n^{(\alpha,\beta)}(x) = \frac{(-1)^n}{2^n n!\, w_{(\alpha,\beta)}(x)} \frac{\mathrm{d}^n}{\mathrm{d}x^n} \left[w_{(\alpha,\beta)}(x)(1 - x^2)^n \right].$$

5.3.1 Orthonormality

As defined above, the hemispherical harmonics are orthonormal:

$$\int_{\mathcal{H}} \mathrm{d}\hat{\Omega}\ (\hat{\mathrm{n}} \cdot \hat{\Omega}) H_{\ell k}(\hat{\Omega}) H_{\ell' k'}(\hat{\Omega}) = \delta_{\ell \ell'} \delta_{k k'} \tag{34}$$

where $\mathcal{H}$ denotes the hemisphere (angular half-space) such that $\hat{\mathrm{n}} \cdot \hat{\Omega} > 0$ for some unit normal $\hat{\mathrm{n}}$, and δ is the Kronecker delta. Here the angular coordinates are chosen such that

$$\int_{\mathcal{H}} \mathrm{d}\hat{\Omega} \ = \int_0^{2\pi} \mathrm{d}\phi \int_0^{\pi/2} \mathrm{d}\theta\ \sin(\theta)$$

$$\hat{\mathrm{n}} \cdot \hat{\Omega} = \cos(\theta).$$

Note that this system is rotated 90 degrees relative to the coordinates used by Zhang and Rahnema.

5.3.2 Completeness

One can see that the functions defined in (31) and (32) differ only by a constant factor from those defined in equations 1 and 2 in reference [43]. Thus, Smelov's proof of completeness applies, and $H_{\ell k}$ provide an orthonormal basis for space $L^2\left(\mathcal{H}, (\hat{\mathrm{n}} \cdot \hat{\Omega})\, \mathrm{d}\hat{\Omega}\right)$ of square integrable functions on the hemisphere $\mathcal{H}$ with weight $(\hat{\mathrm{n}} \cdot \hat{\Omega})$.

5.4 A Recursion for the Adjoint Jacobi Functions

One can derive a three-term recursion relation that can be used to efficiently evaluate the adjoint Jacobi functions $J_\ell^{(k)}(x)$. The following relation regarding Jacobi polynomials comes from reference [45]

$$\frac{\mathrm{d}}{\mathrm{d}x} P_n^{(\alpha,\beta)}(x) = \frac{1}{2}(\alpha + \beta + n + 1) P_{n-1}^{(\alpha+1,\beta+1)}(x). \tag{35}$$

By applying (35) k times one can rewrite (32) as

$$J_\ell^{(k)}(x) = c'_{\ell k}(1 - x^2)^{k/2} P_{\ell-k}^{(k,k+1)}(x) \tag{36}$$

where $c'_{\ell k}$ is another constant factor. This transformation permits the use of several existing results for Jacobi polynomials [45]. The normalization for Jacobi polynomials is such that

$$\begin{aligned} h_n^{(\alpha,\beta)} &= \int_{-1}^{1} dx\, w_{(\alpha,\beta)}(x) \left(P_n^{(\alpha,\beta)}(x)\right)^2 \\ &= \frac{2^{\alpha+\beta+1}}{2n+\alpha+\beta+1} \frac{\Upsilon(n+\alpha+1)\Upsilon(n+\beta+1)}{\Upsilon(n+1)\Upsilon(n+\alpha+\beta+1)} \end{aligned}$$

where Υ is the gamma function, a generalization of the factorial. (The greek upsilon (Υ) is used here to avoid conflict with the notation for the surface expansion basis.) By the definition in (33)

$$\begin{aligned} 1 &= \int_{-1}^{1} \mathrm{d}x\, (1+x) \left(J_\ell^{(k)}(x)\right)^2 \\ &= (c'_{\ell k})^2 \int_{-1}^{1} \mathrm{d}x\, (1+x)(1-x^2)^k \left(P_{\ell-k}^{(k,k+1)}(x)\right)^2 \\ &= (c'_{\ell k})^2 \int_{-1}^{1} \mathrm{d}x\, w_{(k,k+1)}(x) \left(P_n^{(k,k+1)}(x)\right)^2 \\ &= (c'_{\ell k})^2\, h_{\ell-k}^{(k,k+1)}. \end{aligned}$$

Solving for $c'_{\ell k}$ yields

$$c'_{\ell k} = \sqrt{\frac{2\ell+2}{2^{2k+2}} \frac{\Upsilon(\ell-k+1)\Upsilon(\ell+k+2)}{\Upsilon(\ell+1)\Upsilon(\ell+2)}}. \tag{37}$$

The Jacobi polynomials have a 3-term recurrence of the form

$$P_n^{(\alpha,\beta)}(x) = (a_{n,0}^{(\alpha,\beta)} x + a_{n,1}^{(\alpha,\beta)}) P_{n-1}^{(\alpha,\beta)}(x) + a_{n,2}^{(\alpha,\beta)} P_{n-2}^{(\alpha,\beta)}(x) \tag{38}$$

with coefficients

$$a_{n,0}^{(\alpha,\beta)} = \frac{(2n+\alpha+\beta-1)(2n+\alpha+\beta)}{2n(n+\alpha+\beta)} \tag{39}$$

$$a_{n,1}^{(\alpha,\beta)} = \frac{(\alpha^2-\beta^2)(2n+\alpha+\beta-1)}{2n(n+\alpha+\beta)(2n+\alpha+\beta-2)} \tag{40}$$

$$a_{n,2}^{(\alpha,\beta)} = -\frac{(n+\alpha-1)(n+\beta-1)(2n+\alpha+\beta)}{n(n+\alpha+\beta)(2n+\alpha+\beta-2)}. \tag{41}$$

Using equations (36) to (41) one can derive a new 3-term recurrence for the orthonormal adjoint Jacobi functions

$$J_k^{(k)}(x) = \sqrt{\frac{(2k+1)!!}{2(2k)!!}}(1-x^2)^{k/2} \tag{42}$$

$$J_{k+1}^{(k)}(x) = \left(\hat{a}_{k+1,0}^{(k)} x + \hat{a}_{k+1,1}^{(k)}\right) J_k^{(k)}(x) \tag{43}$$

$$J_\ell^{(k)}(x) = \left(\hat{a}_{\ell,0}^{(k)} x + \hat{a}_{\ell,1}^{(k)}\right) J_{\ell-1}^{(k)}(x) + \hat{a}_{\ell,2}^{(k)} J_{\ell-2}^{(k)}(x) \tag{44}$$

where $n!! = n(n-2)(n-4)\cdots$ denotes the product of every other positive integer. The coefficients are given by

$$\hat{a}_{\ell,0}^{(k)} = (2\ell+1)\frac{1}{\sqrt{(\ell-k)(\ell+k+1)}} \tag{45}$$

$$\hat{a}_{\ell,1}^{(k)} = -\frac{2k+1}{2\ell-1}\frac{1}{\sqrt{(\ell-k)(\ell+k+1)}} \tag{46}$$

$$\hat{a}_{\ell,2}^{(k)} = -\frac{2\ell+1}{2\ell-1}\sqrt{\frac{(\ell+k)(\ell-k-1)}{(\ell-k)(\ell+k+1)}} \tag{47}$$

To the author's knowledge, the recursion given by equations (42) to (47) has not been published previously. It is simpler than the recursion given by equation 5 in reference [31], and unlike the recursion derived in reference [43], it has no singularity at $x = \pm 1$.

5.5 Uncollided Fluence Basis

For the uncollided fluence expansion, basis functions of the form

$$\tilde{Q}^{i_x i_y i_z klg}(\mathbf{r}, \hat{\Omega}, E) = P_{i_x}(x) P_{i_y}(y) P_{i_z}(z)\; Y_{\ell k}(\hat{\Omega})\, \chi_g(E)$$

are used where $Y_{\ell k}$ is the real spherical harmonic of degree ℓ and order k with $|k| < \ell$. As with the surface basis $\{\Gamma^m\}$, the high order cross-terms are eliminated upon truncation of the basis.

5.6 Spherical Harmonics

The the real spherical harmonic $Y_{\ell k}(\mu, \phi)$ are defined such that

$$Y_{\ell k}(\mu,\phi) = \begin{cases} \frac{1}{\sqrt{\pi}} P_\ell^k(\mu) cos(k\phi), & k > 0 \\ \frac{1}{\sqrt{2\pi}} P_\ell^k(\mu), & k = 0 \\ \frac{1}{\sqrt{\pi}} P_\ell^{-k}(\mu) sin(-k\phi), & k < 0 \end{cases}$$

The normalized associated Legendre Polynomials are defined

$$P_\ell^k(x) = (-1)^k \sqrt{\frac{(\ell-k)!}{(\ell+k)!}} (1-x^2)^{k/2} \frac{d^k}{dx^k} P_\ell(x)$$

for $0 \leq k \leq \ell$ with ℓ, k integers with the orthonormal Legendre polynomial

$$P_n(x) = \sqrt{\frac{2n+1}{2}} \frac{1}{2^n n!} \frac{d^n}{dx^n} \left[(x^2-1)^n \right]$$

For more, see Abramowitz and Stegun [1].

5.7 Response Function Generator

A user code for EGSnrc called RF_GENERATOR is developed to evaluate all of the responses, both volume-to-volume and surface-to-volume, used in this work. Because Monte Carlo is used to estimate the response coefficients, there is also a variance estimate associated with each response coefficient. To estimate the impact of this uncertainty on COMET-PE dose calculations, these variances are propagated through the solution of the linear system as described by Zhang and Rahnema [54, 55].

CHAPTER VI

WATER-ALUMINUM-LUNG INTERFACE PHANTOM

To test the accuracy of the method, a benchmark developed by Rogers and Mohan [39] is used. While the benchmark was originally intended for Monte Carlo methods, it challenges any method's ability to handle strong heterogeneity. It was used, for example, to test the deterministic code Attila [19]. The problem consists of a 30.5 x 39.5 x 30 cm^3 layered slab phantom irradiated by a photon point source with 100 m source to surface distance. This is depicted in figure 7 on the following page. The 18 MV spectrum comes from Monte Carlo simulation of a Varian Clinac 2100. The the benchmark specification also includes a 6 MV photon energy spectrum from the same accelerator [39]. These energy spectra are plotted in figure 8 on page 39. The beam is collimated to 1.5 x 1.5 cm^2 on the phantom surface. The phantom's material layout is given by table 1. The cross sections for aluminum, water, and lung tissue are plotted in figure 9. The material density changes by an order of magnitude across the aluminum-lung interface, and the beam's narrow collimation precludes lateral charged particle equilibrium. These factors combine to make this benchmark an extreme test of photon-electron transport.

Table 1: Material composition of the interface phantom.

Region*	Material	Density (g/cm^3)
$0 \leq d < 3$	Water	1.0
$3 \leq d < 5$	Aluminum	2.702
$5 \leq d < 12$	Lung Tissue	0.26
$12 \leq d < 30$	Water	1.0

* d=depth in cm

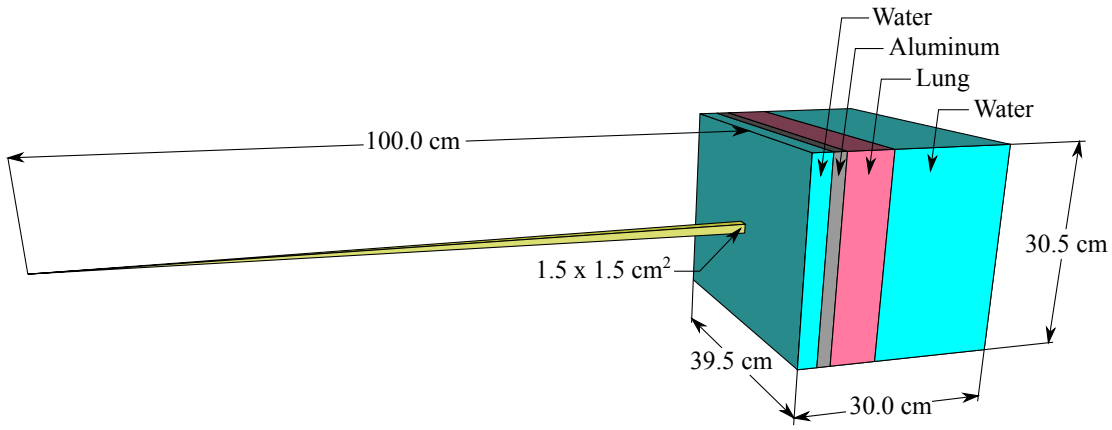

Figure 7: The ICCR benchmark problem.

6.1 Reference Solutions

Reference KERMA (kinetic energy released per unit mass), and dose distributions were calculated using the DOSXYZnrc user code [51] of the EGSnrc code system [24]. Dose was calculated with 1 billion source particles and the transport parameters listed in table 18 on page 89. KERMA was calculated using 16 billion source particles and all the same transport parameters with the exception of the electron energy cutoff, which was set to 100 MeV to disable electron transport. Energy deposition was tallied in 5 x 5 x 2 mm^3 voxels along the beam centerline. The dose reference solution was found to agree with the EGS4 solution provided by Rogers and Mohan [39] within 1% of the maximum dose or within 2.5% of local dose. The reference KERMA solution required 3.0 hours on 12 compute nodes and the reference dose solution required 1.0 hour on the same system. All calculations presented in this work were performed using a cluster with 16 compute nodes, each with 16 GB of memory and 2 quad-core (a total of 8 cores per node) AMD Opteron 2350s clocked at 2.0 GHz.

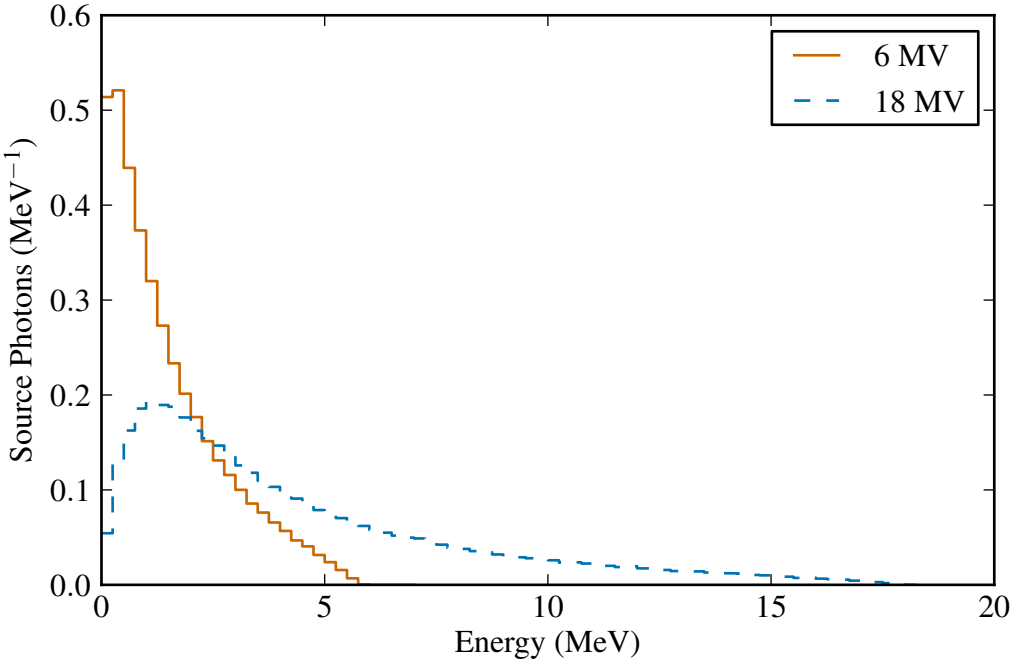

Figure 8: Photon source spectra for a Varian Clinac 2100

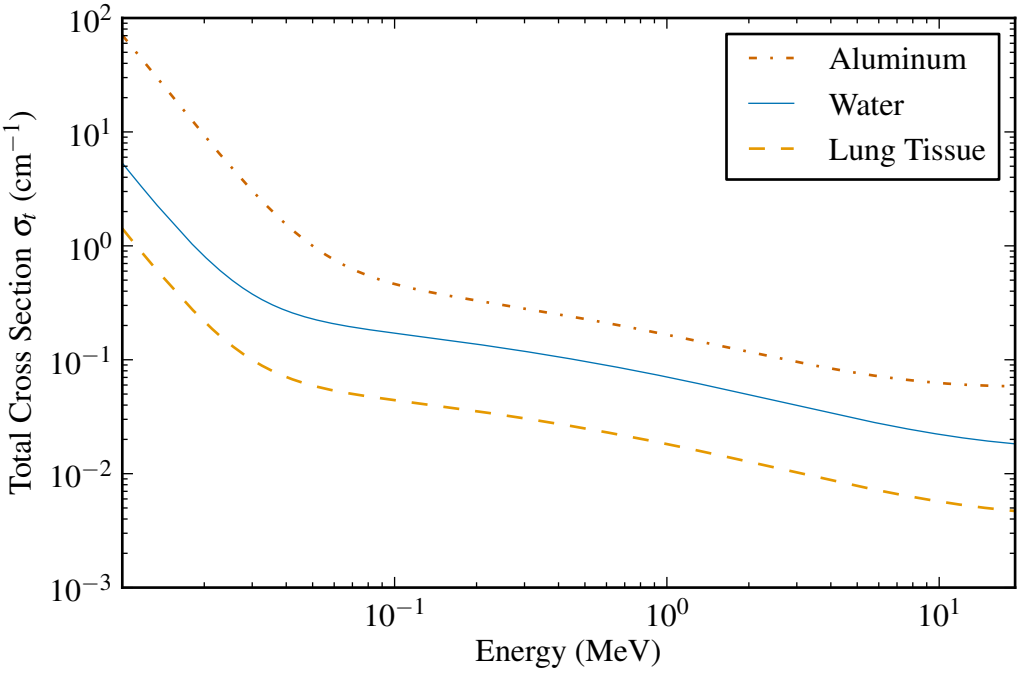

Figure 9: Total photon cross sections for aluminum, water, and lung tissue

6.2 Expansion order notation

In order to succinctly express the truncation order used for a particular calculation, the notation

$$(o_1 o_2 o_3 o_4 o_5 o_6)$$

is introduced where

o_1 = angular expansion order for uncollided fluence

o_2 = spatial expansion order for uncollided fluence

o_3 = angular expansion order for photon fluence

o_4 = spatial expansion order for photon fluence

o_5 = angular expansion order for electron fluence

o_6 = spatial expansion order for electron fluence.

For example, order=(917062) would denote an expansion of 9^{th} order in uncollided angle, 1^{st} order in uncollided space, 7^{th} order in photon angle, 0^{th} order in photon space, 6^{th} order in electron angle, and 2^{nd} order in electron space. Additionally, when discussing KERMA, the electron expansion orders are irrelevant, so they are omitted. Thus, (9264) means (9264...). Similarly, when discussing TERMA, neither of the surface expansions are relevant, so (92) is written to mean (92...).

6.3 Pre-computation

Response functions for water, aluminum, and lung are pre-computed using RF_GENERATOR, and GETXS is used to extract continuous energy cross sections from EGSnrc and to compute multigroup cross sections. All of the continuous energy cross sections and physics settings were the same as those used to for the reference solutions. The uncollided photon energy group structure was composed of 20 groups with widths varying between 40 keV and 2 MeV. This same energy structure was used to form bins for both the photon and electron

fluences; in the electron case there were only 19 bins because one was below the energy cutoff. More detail on this energy structure is given in section 6.5.1 below. Response functions were computed for two different mesh sizes: 5 mm and 10 mm. Responses were computed up to order (929172). For each source basis function, 10 million source particles were simulated. Dose and KERMA were tallied in 1 x 1 x 1 mm^3 voxels within each mesh. The pre-computation time for the 5 mm library was 4.7 hours and for the 10 mm library, 9.3 hours. Both libraries were run using 12 compute nodes (96 cores).

6.4 COMET-PE Solutions

Only the central 11 x 11 x 30 cm^3 column of the phantom was modeled in COMET-PE (an additional DOSXYZnrc calculation verified that this did not affect the solution). The solver was applied sequentially to the photon fluence system and then to the electron fluence system. Electron to photon feedback was found to be negligible, so iteration between the photon and electron solutions was not required. A tolerance of $\epsilon = 10^{-3}$ was used for convergence of the Gauss-Seidel iteration. Because the response coefficients were computed stochastically, they contribute uncertainty to the COMET-PE solution. Although mesh sizes of 5 mm and 10 mm were used, energy deposition was calculated on a finer 1 x 1 x 1 mm^3 grid for the entire phantom. To obtain a value for each of the 5 x 5 x 2 mm^3 voxels required by the benchmark, the output was integrated using nearest-neighbor interpolation.

6.5 Sensitivity Study

There are many input parameters that affect the accuracy and efficiency of COMET-PE solutions. The most important parameters are the energy group structure and the response expansion truncation orders. In this section, values for these parameters are chosen in an attempt to maximise both accuracy and efficiency. To assess accuracy, the water-aluminum-lung slab benchmark problem is used. The multigroup energy structure is assessed using an

analytical continuous energy solution to a 1-D version of the slab phantom. Building upon this, the new method is used to calculate total energy released per unit mass (TERMA) for the 3-D problem. This is verified against the analytical solution and against a Monte Carlo reference solution. Next, kinetic energy released per unit mass (KERMA) is calculated using various photon expansion truncation orders to determine appropriate parameters for photon transport. Finally, dose is calculated, testing the effect of the electron expansion truncation parameters. By stepping through the energy deposition process systematically– from TERMA to KERMA to dose–the approximation is built in layers, validating one at a time. This way, an exhaustive search of the parameter space is avoided. Table 19 on page 90 gives definitions for the quantities used to evaluate the accuracy and precision of the results.

6.5.1 Energy Group Structure

To analyze the effectiveness of different group structures, a 1-D analogue of the slab phantom is examined, considering only the uncollided photon fluence. Assuming an isotropic source $q_0(\hat{\Omega}, E) = q_0(E)$, one can integrate (26) and (27) over all angles to get the 1-D fluence distributions

$$\psi^u(x, E) = \frac{1}{x^2} q_0(E) \exp[-\tau(x, E)] \tag{48}$$

$$\psi^u_g(x) = \frac{1}{x^2} q_{0g} \exp[-\tau_g(x)] \tag{49}$$

where $x = |\mathbf{r} - \mathbf{r}_0|$ is the distance from the source. These equations form the basis of a simple test problem that can help decide on a multigroup energy structure. The material composition for this test is taken directly from the slab phantom as detailed in table 1 on page 37. The prospective group structures are tested against both beam energies provided with the benchmark (see figure 8 on page 39) to derive a group structure that works for both 6 MV and 18 MV beams. To compare group structures, both the total uncollided energy fluence Ψ (MeV/cm^2) and the total energy released per unit mass (TERMA) (MeV/g) are

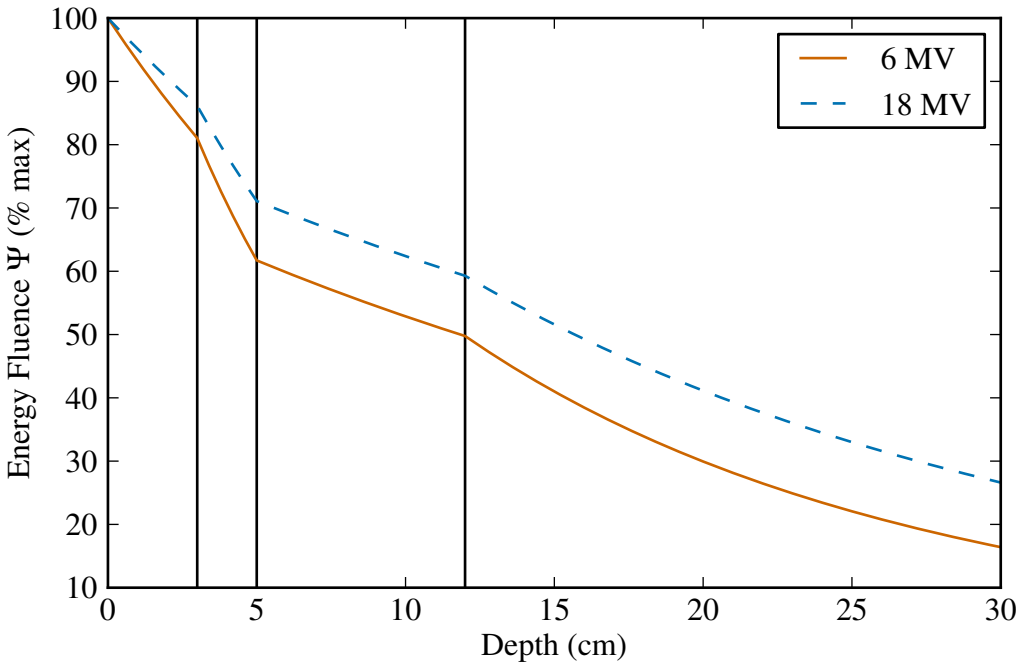

Figure 10: Total uncollided energy fluence as a function of depth. Vertical divisions indicate material interfaces.

examined. In continuous energy, these quantities are defined

$$\Psi(x) = \int \mathrm{d}E\ \psi^u(x, E) \tag{50}$$

$$\mathrm{TERMA}(x) = \int \mathrm{d}E\ \frac{\sigma_t(x, E)}{\rho(x)} \psi^u(x, E). \tag{51}$$

These integrals are evaluated using an 8,192 point composite trapezoidal quadrature to calculate continuous energy reference solutions. These reference solutions are plotted in figures 10 and 11 on pages 43–44. Using the multigroup approximation, these quantities take the form

$$\Psi_{MG}(x) = \sum_g \psi_g^u(x) \tag{52}$$

$$\mathrm{TERMA}_{MG}(x) = \sum_g \frac{\sigma_{tg}(x)}{\rho(x)} \psi_g{}^u(x). \tag{53}$$

The effectiveness of the multigroup structures is assessed by comparing values calculated with multigroup summations (52) and (53) to those calculated with continuous energy integrals (50) and (51). Four different group structures with 5, 10, 20, and 30 nested groups as

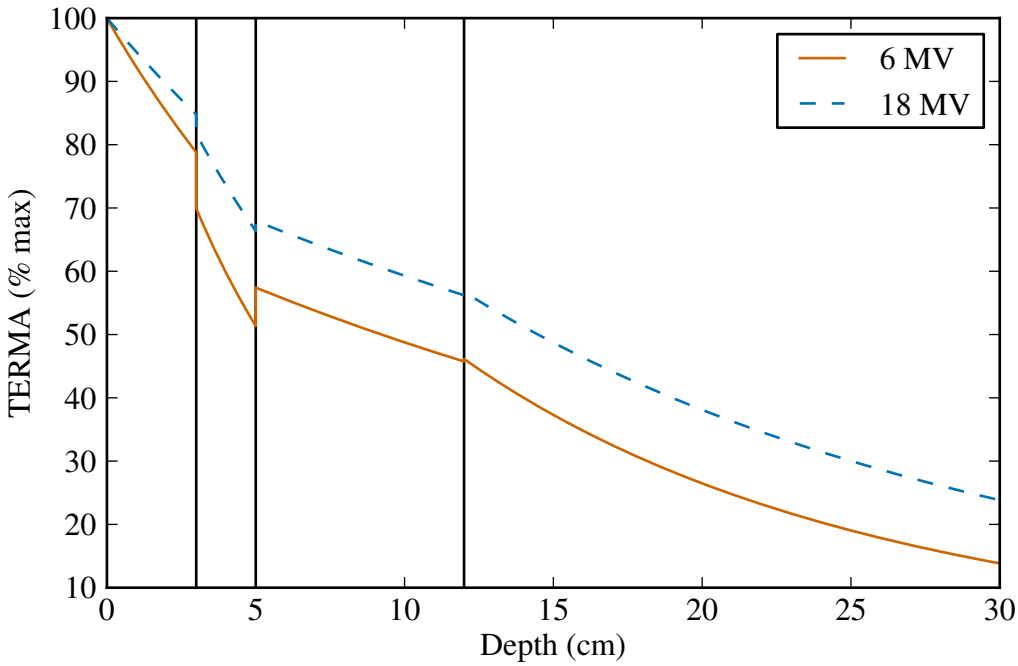

Figure 11: TERMA as a function of depth. Vertical divisions indicate material interfaces.

given in table 2 on the next page are tested. The effects of the different group structures on total uncollided energy fluence Ψ can be seen in figures 12 and 13 on page 46. The results for TERMA are shown in figures 14 and 15 on page 47. Tables 3 and 4 on page 48 provide a summary of these comparisons. In each case, the finer group structures lead to more accurate estimations. Thus, the 30 group structure is the most accurate and the 5 group structure is the least accurate. Because computational effort increases with the number of groups, the 20 group structure is used for the rest of this work. The same structure is also used to specify the energy bins for the photon and electron fluence expansions. It is more difficult to assess the impact of the bin structure since it is built in to any response function library at the time of pre-computation. Unlike the spatial and angular truncation order, it is more challenging to vary the bin structure at run time. Methods to optimize the choice of energy bin structure are an area of ongoing research.

Table 2: Energy group structures for uncollided photons. Each column gives the group boundaries $E_1, E_2, ..., E_{G+1}$ in MeV.

30 Groups	20 Groups	10 Groups	5 Groups
19.00	19.00	19.00	19.00
18.00	18.00		
17.00			
16.00	16.00	16.00	
15.00			
14.00	14.00		
13.00			
12.00	12.00	12.00	12.00
11.00			
10.00	10.00		
9.00			
8.00	8.00	8.00	
7.00			
6.00	6.00		
5.50	5.50	5.50	5.50
5.00	5.00		
4.50	4.50	4.50	
4.00			
3.50	3.50		
3.00	3.00	3.00	3.00
2.50	2.50		
2.00	2.00	2.00	
1.50	1.50		
1.25			
1.00	1.00	1.00	1.00
0.75			
0.50	0.50		
0.25	0.25	0.25	
0.15			
0.05	0.05		
0.01	0.01	0.01	0.01

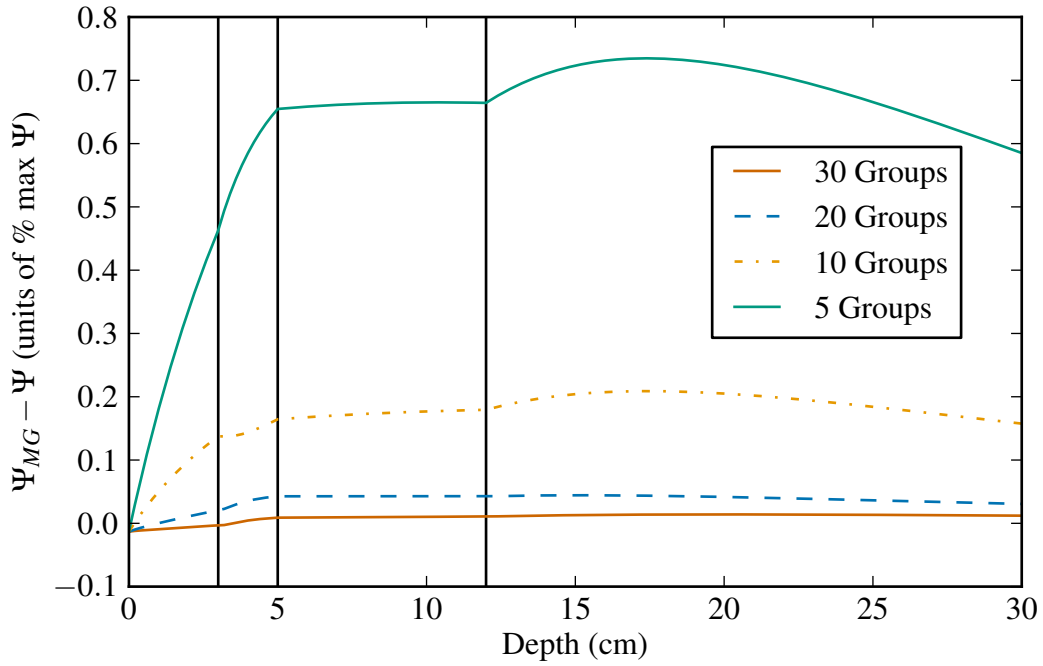

Figure 12: Effect of group structure on error in total uncollided energy fluence for the 6 MV source. Vertical divisions indicate material interfaces.

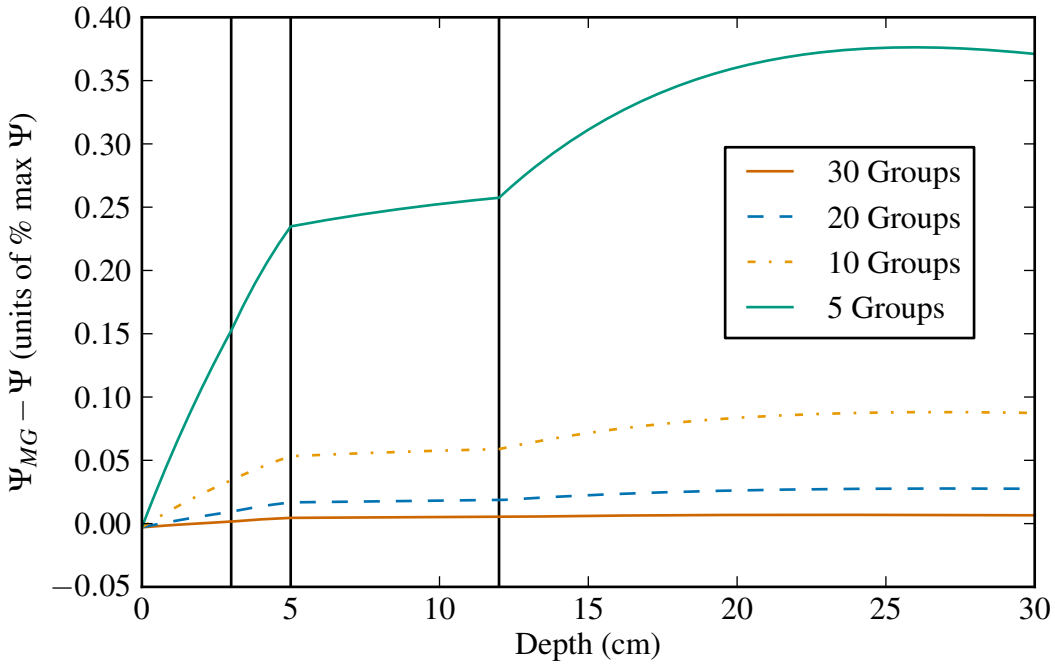

Figure 13: Effect of group structure on error in total uncollided energy fluence for the 18 MV source. Vertical divisions indicate material interfaces.

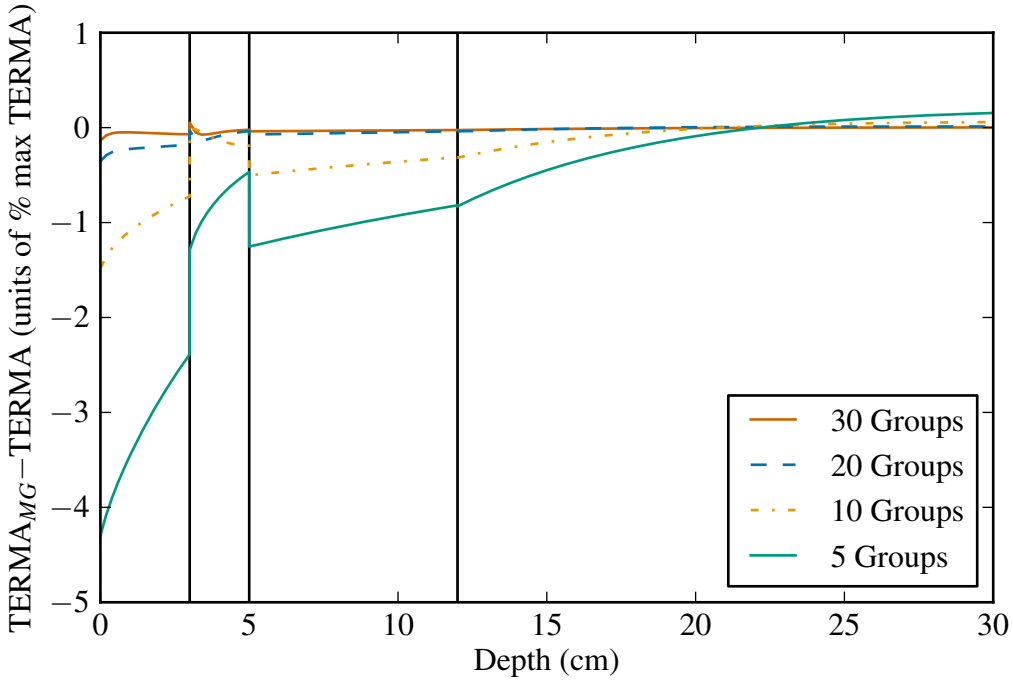

Figure 14: Effect of group structure on error in TERMA for the 6 MV source. Vertical divisions indicate material interfaces.

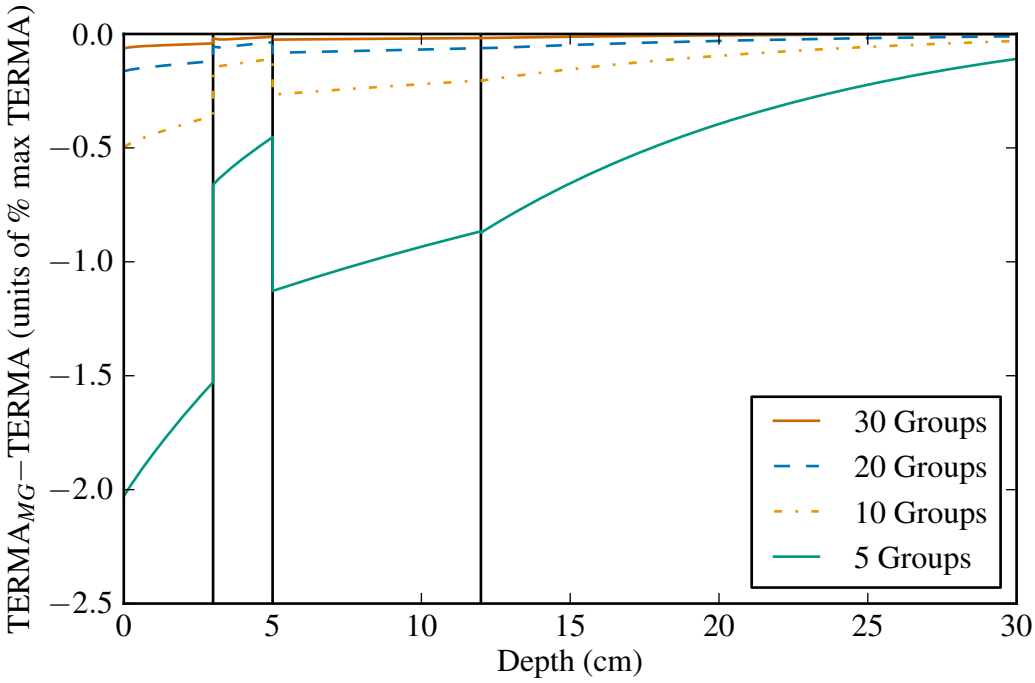

Figure 15: Effect of group structure on error in TERMA for the 18 MV source. Vertical divisions indicate material interfaces.

Table 3: Effects of group structure on total uncollided energy fluence and TERMA for the 6 MV source.

Number of Groups	Error in Ψ		Error in TERMA	
	Max (%)	RMS (%)	Max (%)	RMS (%)
5	0.74	0.65	4.30	1.20
10	0.21	0.18	1.48	0.40
20	0.04	0.04	0.35	0.08
30	0.01	0.01	0.14	0.03

Table 4: Effects of group structure on total uncollided energy fluence and TERMA for the 18 MV source

Number of Groups	Error in Ψ		Error in TERMA	
	Max (%)	RMS (%)	Max (%)	RMS (%)
5	0.38	0.30	2.03	0.83
10	0.09	0.07	0.50	0.20
20	0.03	0.02	0.16	0.06
30	0.01	0.01	0.06	0.02

6.5.2 Accuracy of Source Integration (TERMA)

It is desirable to test the new aspects of COMET-PE in isolation from the rest of the method. Namely, the multigroup cross sections, ray-tracing, adaptive cubature, and volume-to-volume response function generator should be tested. Unlike convolution-superposition methods, COMET-PE does not actually use TERMA when calculating KERMA or dose. However, TERMA calculation provides a good verification since the solution can be computed analytically. TERMA is computed for the benchmark problem using the COMET-PE method, and this is compared with the 1-D analytical result that is used to test multigroup energy structures above. For further verification, the analytical solution is compared to a result from DOSXYZnrc, which was modified to tally TERMA. One billion source histories are used for the DOSXYZnrc computation. Figure 16 shows a comparison between the the analytical and Monte Carlo results. The difference between these two solutions is attributable almost entirely to statistical uncertainty. Figure 17 on page 51 and 18 show a comparison between COMET-PE and the analytical results for 5 mm and 10 mm mesh

sizes. These figures also include the effect of varying the spatial order of the uncollided fluence expansion. The effect is larger for the 10 mm mesh size than it is for the 5 mm mesh size. When a 0^{th} order spatial expansion is used it creates a stair step pattern in the resulting TERMA distribution. This is more difficult to see with the 5 mm mesh size because the effect is smaller and because of the TERMA grid that is used as part of the benchmark. For both the 5 mm mesh size and the 10 mm mesh size, COMET-PE outputs TERMA on a 1 x 1 x 1 mm^3 grid; these values are then summed to give values on a coarser 5 x 5 x 2 mm^3 grid per the benchmark specification. Thus, with the 5 mm mesh size, more of the TERMA comparison points represent an average of values from 2 neighboring meshes. This obscures the stair step effect that would otherwise be more obvious. Table 5 summarizes the results of the comparison and table 6 gives information on the statistical uncertainty. The differences between COMET-PE and the analytical solution are caused by statistical uncertainty and systematic errors from three sources:

1. the multigroup approximation – see figure 15 on page 47
2. volume effects – the analytical solution is a 1-D average and the COMET-PE solutions are 3-D averages
3. expansion truncation – this is most obvious for the 10 mm mesh.

The volume effects also contribute to the difference between Monte Carlo and the analytical solution, but these differences are overwhelmed by statistical uncertainty.

6.5.3 Accuracy of Photon Transport (KERMA)

By calculating kinetic energy released per unit mass (KERMA), one can measure the accuracy of photon transport. To determine appropriate expansion truncation orders for the surface photon fluence, COMET-PE is compared with DOSXYZnrc for several settings of these parameters. The Monte Carlo reference solution has average relative uncertainty of 0.08% and maximum relative uncertainty of 0.11%.

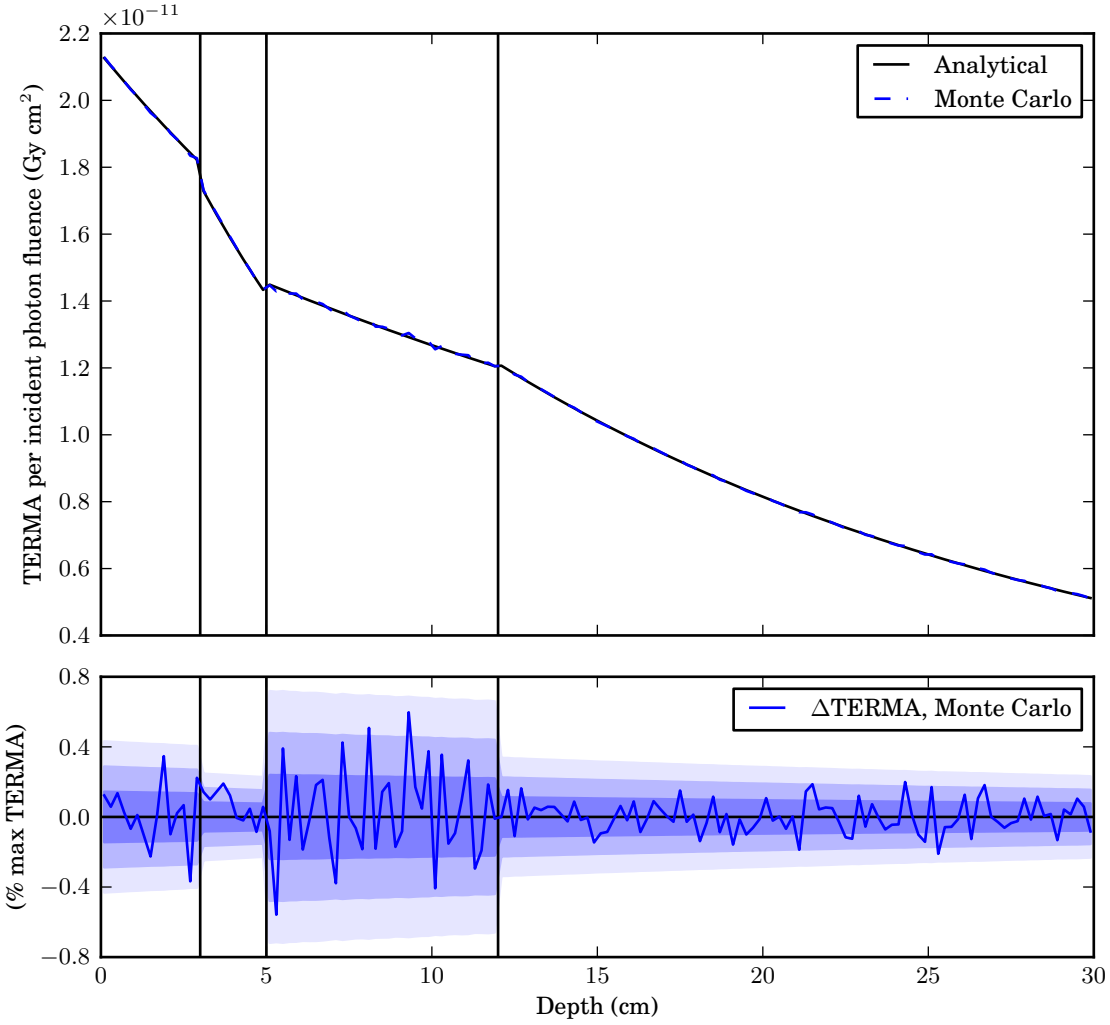

Figure 16: Monte Carlo compared with analytical TERMA. Shading indicates uncertainty within 1, 2, and 3 estimated standard deviations, and vertical divisions indicate material interfaces.

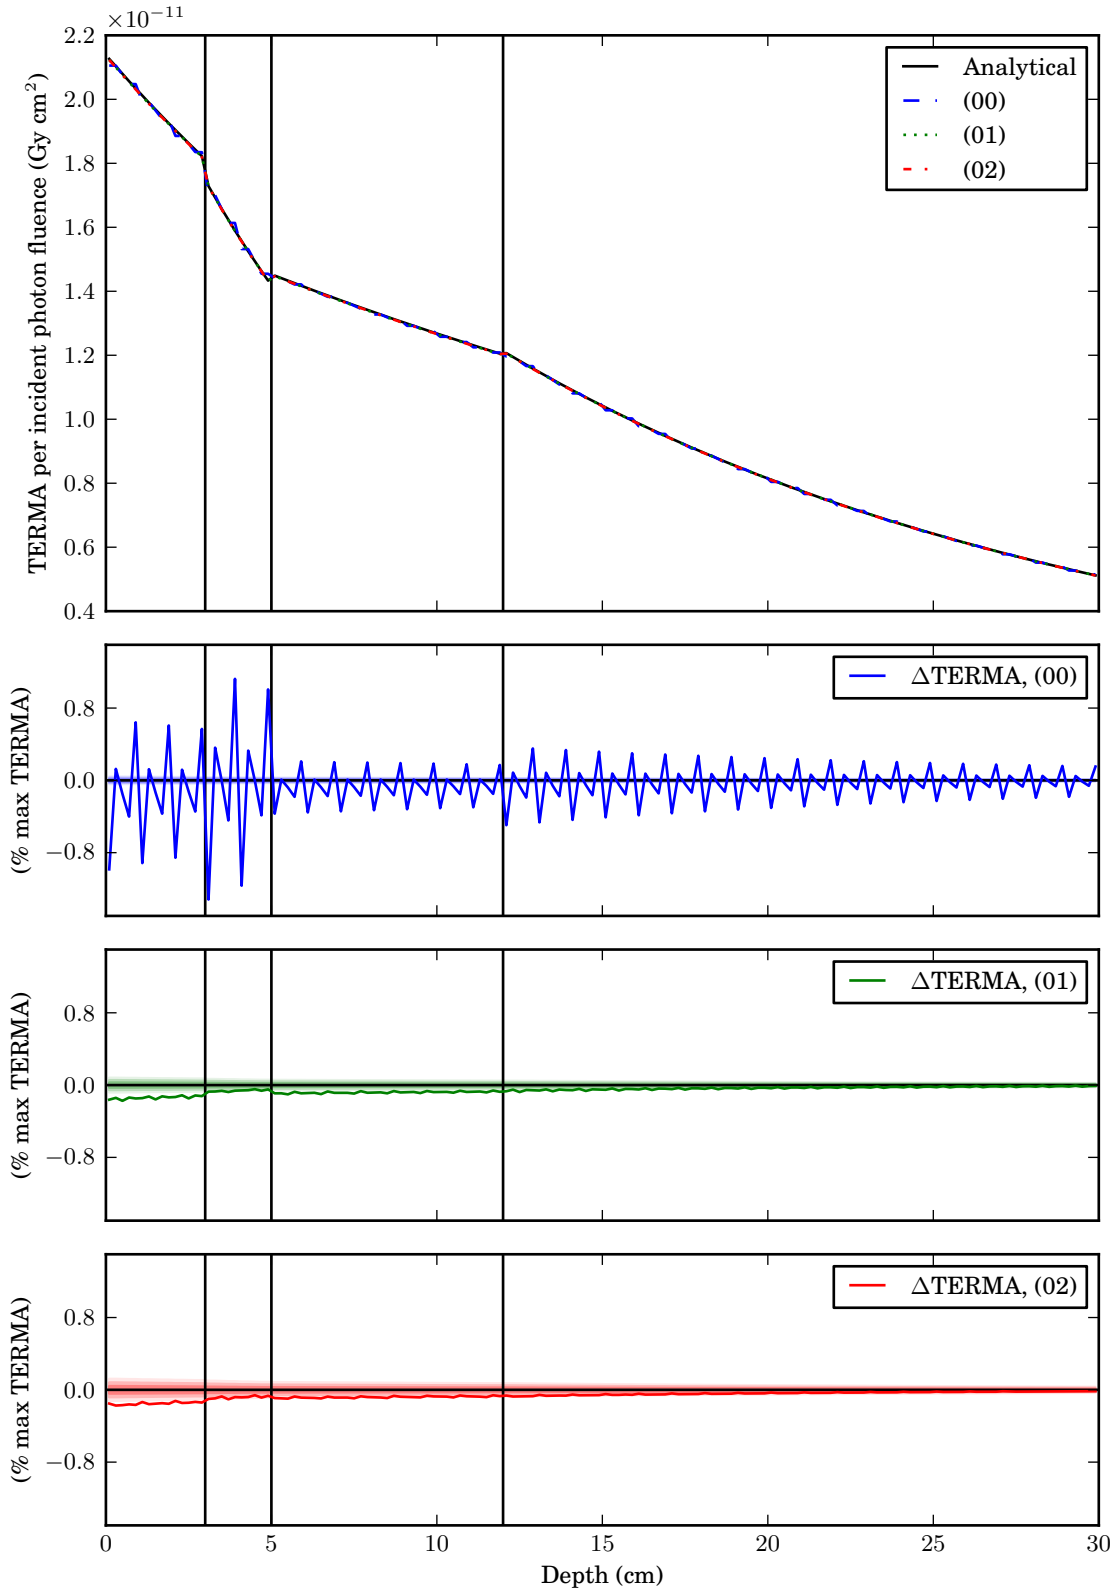

Figure 17: COMET-PE compared with analytical TERMA using 5 mm mesh size as the uncollided spatial expansion order is varied. Shading indicates uncertainty within 1, 2, and 3 estimated standard deviations, and vertical divisions indicate material interfaces.

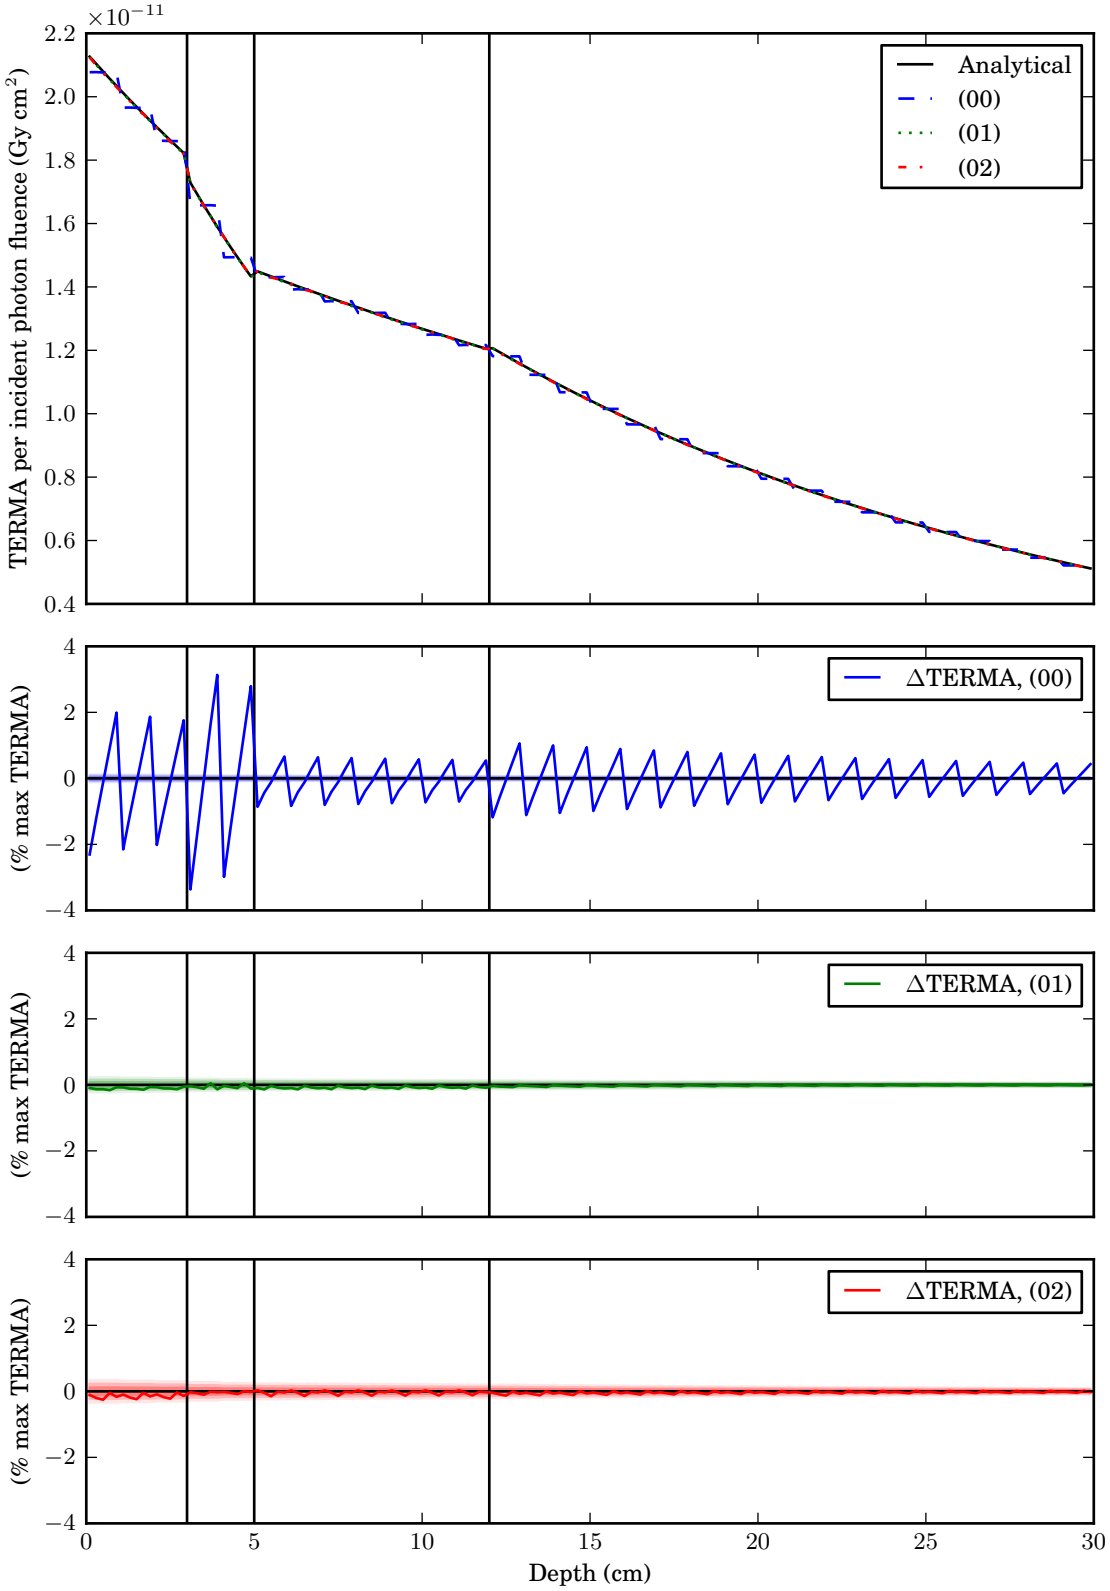

Figure 18: COMET-PE compared with analytical TERMA using 10 mm mesh size as the uncollided spatial expansion order is varied. Shading indicates uncertainty within 1, 2, and 3 estimated standard deviations, and vertical divisions indicate material interfaces.

Table 5: Comparison of TERMA calculations relative to an analytical solution. Both mesh size and uncollided spatial expansion order are varied for the COMET-PE solutions. The values following ± give estimated uncertainty at the level of one standard deviation.

	Error		Relative Error	
Method	Mean (%)	Max (%)	Mean (%)	Max (%)
Monte Carlo	0.11 ±0.01	0.60 ±0.23	0.23 ±0.02	0.98 ±0.38
5 mm, (00)	0.21 ±0.00	1.32 ±0.01	0.40 ±0.00	1.62 ±0.01
5 mm, (01)	0.06 ±0.00	0.17 ±0.03	0.10 ±0.00	0.18 ±0.03
5 mm, (02)	0.06 ±0.00	0.17 ±0.04	0.11 ±0.00	0.18 ±0.04
10 mm, (00)	0.61 ±0.00	3.36 ±0.03	1.15 ±0.00	4.20 ±0.04
10 mm, (01)	0.05 ±0.00	0.16 ±0.07	0.08 ±0.01	0.21 ±0.08
10 mm, (02)	0.05 ±0.01	0.25 ±0.11	0.10 ±0.01	0.26 ±0.11

Table 6: Uncertainty in the TERMA comparison.

	Relative Uncertainty	
Method	Average (%)	Max (%)
Monte Carlo	0.28	0.39
5 mm, (00)	0.01	0.01
5 mm, (01)	0.03	0.03
5 mm, (02)	0.04	0.04
10 mm, (00)	0.04	0.04
10 mm, (01)	0.08	0.08
10 mm, (02)	0.12	0.13

The results for varying the photon angular expansion order for the 5 mm mesh size are given in table 7 on the next page and figure 19 on page 55. The 5 mm COMET-PE solutions each have average relative uncertainty of ~0.15% and maximum relative uncertainty of ~0.16%. Because the COMET-PE uncertainty estimates appear to be relatively insensitive to expansion order, the uncertainties for 5 mm cases (9271), (9281), (9290), and (9291) are not calculated explicitly; instead, these are assumed to be the same as those of the (9261) expansion. From the table, one can see that accuracy improves with increasing angular order up to about the (9261) expansion, after which the differences are on the order of the statistical uncertainty. From figure 19, one can see that the difference between COMET-PE

and Monte Carlo is largely explained by statistical uncertainty for the (9281) expansion.

Table 7: Sensitivity of KERMA calculation to photon angular expansion order for the 5 mm mesh size.

	Error		Relative Error	
Expansion	Mean (%)	Max (%)	Mean (%)	Max (%)
(9201)	1.40 ±0.01	4.59 ±0.13	2.44 ±0.01	5.53 ±0.15
(9211)	0.95 ±0.01	1.68 ±0.13	1.93 ±0.01	2.95 ±0.17
(9221)	0.23 ±0.01	0.39 ±0.07	0.52 ±0.01	1.25 ±0.18
(9231)	0.27 ±0.01	0.56 ±0.11	0.60 ±0.01	1.58 ±0.18
(9241)	0.17 ±0.01	0.38 ±0.12	0.38 ±0.01	0.89 ±0.17
(9251)	0.17 ±0.01	0.38 ±0.12	0.39 ±0.01	1.00 ±0.18
(9261)	0.10 ±0.01	0.33 ±0.13	0.22 ±0.01	0.68 ±0.18
(9271)	0.13 ±0.01	0.31 ±0.15	0.29 ±0.01	0.76 ±0.18
(9281)	0.09 ±0.01	0.33 ±0.13	0.19 ±0.01	0.55 ±0.18
(9291)	0.10 ±0.01	0.31 ±0.15	0.21 ±0.01	0.57 ±0.18

The results for varying the photon spatial expansion order for the 5 mm mesh size are given in table 8 and figure 20 on page 56. The first order spatial expansion provides greater accuracy than the zeroth order expansion. From figure 20, it appears that much of the difference between COMET-PE and Monte Carlo can be explained by statistical uncertainty, but there may be a systematic error of ~0.1% of the maximum KERMA that occurs in the second water slab, beyond a depth of 12 cm. Regardless, the accuracy is excellent, with all relative errors substantially below 1%.

Table 8: Sensitivity of KERMA calculation to photon spatial expansion order for the 5 mm mesh size.

	Error		Relative Error	
Expansion	Mean (%)	Max (%)	Mean (%)	Max (%)
(9290)	0.15 ±0.01	0.40 ±0.14	0.30 ±0.01	0.65 ±0.17
(9291)	0.10 ±0.01	0.31 ±0.15	0.21 ±0.01	0.57 ±0.18

The results for varying the photon angular expansion order for the 10 mm mesh size

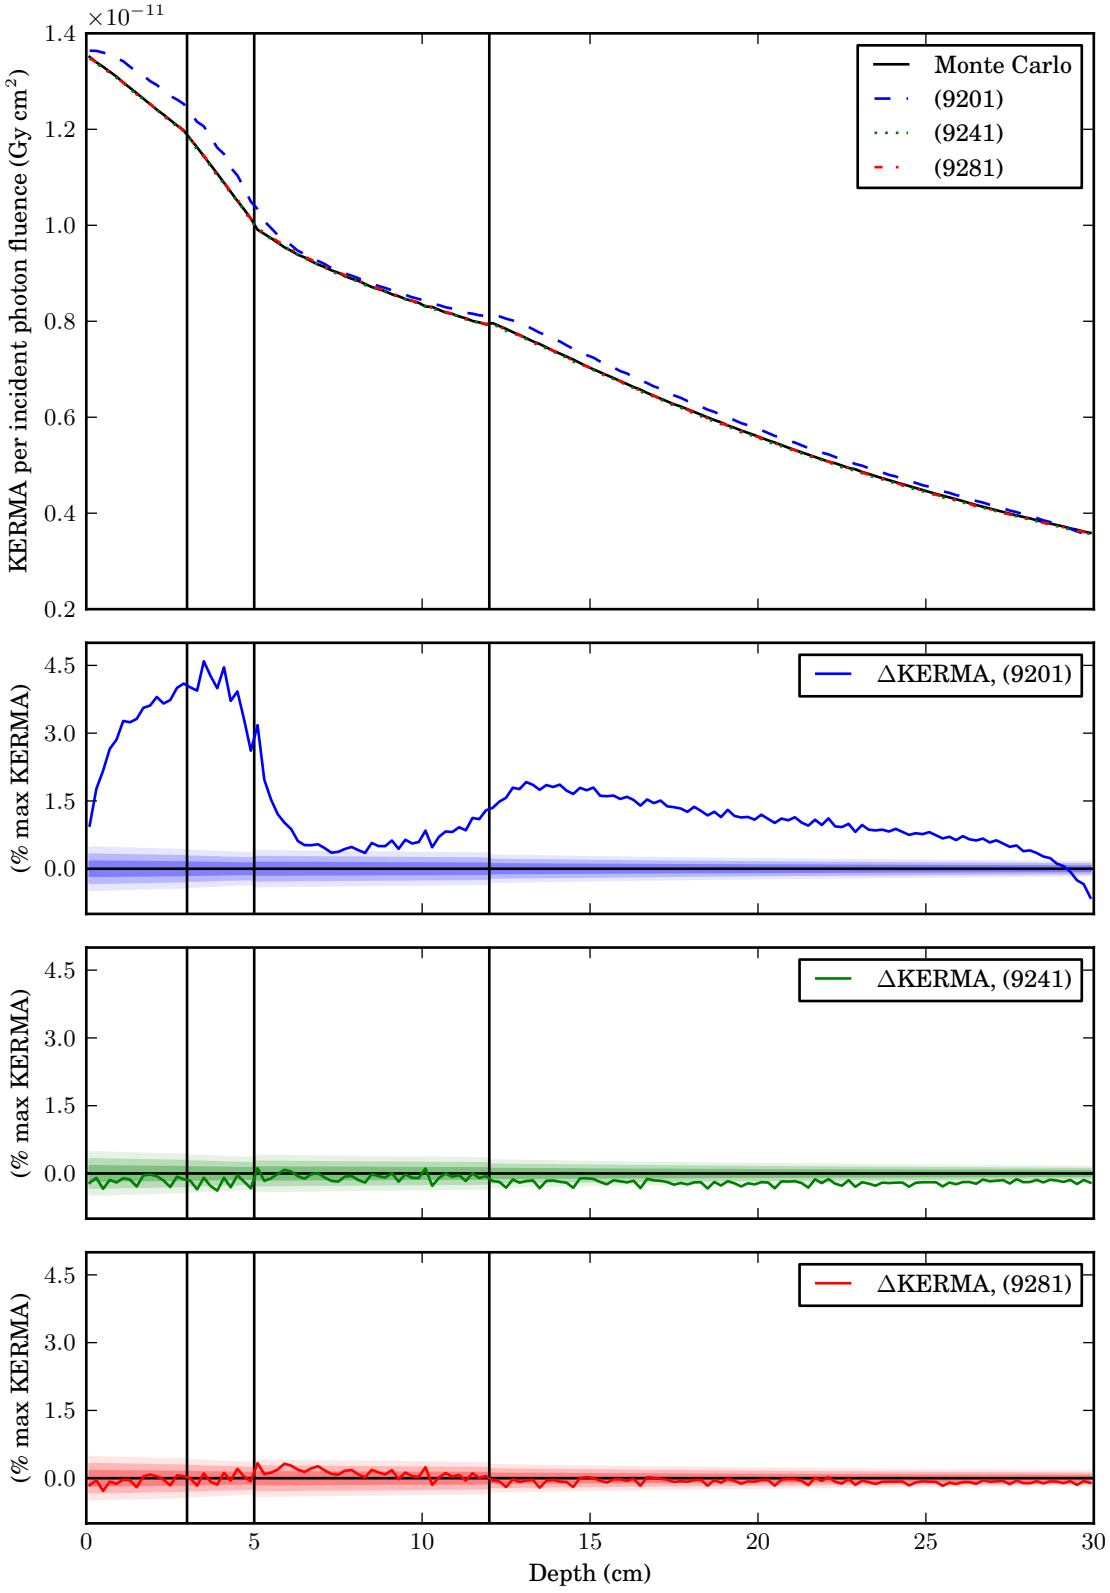

Figure 19: Impact of photon angular order on KERMA for the 5 mm mesh size. Shading indicates uncertainty within 1, 2, and 3 estimated standard deviations, and vertical divisions indicate material interfaces.

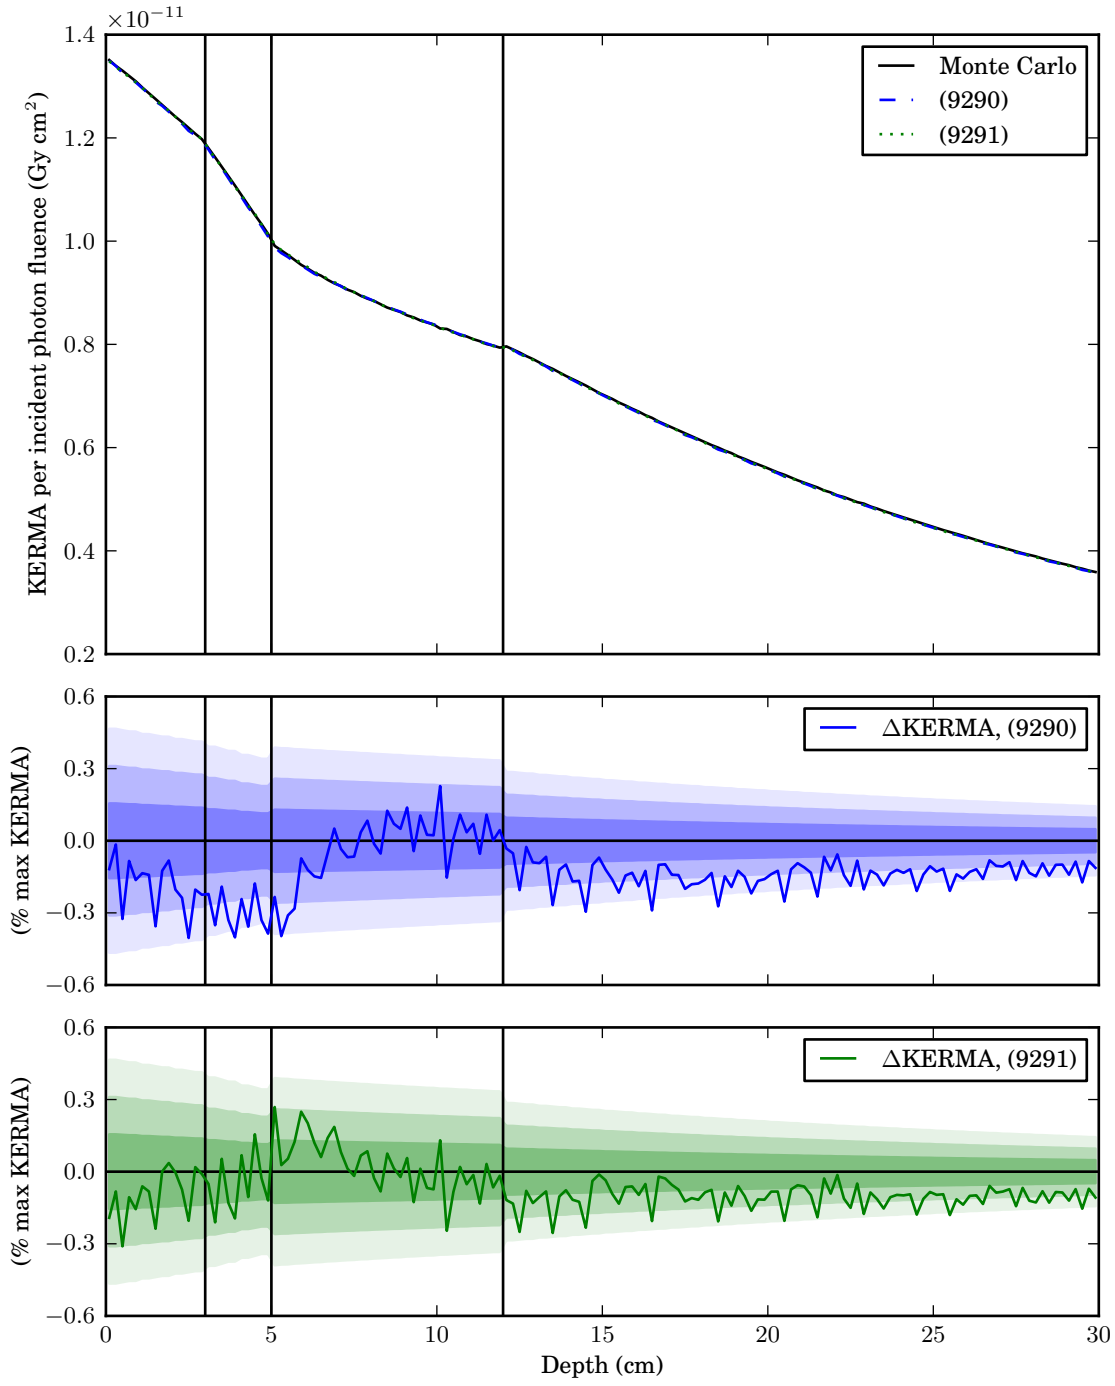

Figure 20: Impact of photon spatial order on KERMA for the 5 mm mesh size. Shading indicates uncertainty within 1, 2, and 3 estimated standard deviations, and vertical divisions indicate material interfaces

are given in table 9 on page 57 and figure 21 on page 58. The 10 mm COMET-PE solutions each have average relative uncertainty of ~0.43% and maximum relative uncertainty of ~0.46%. For the 10 mm cases, all uncertainties are estimated directly. One can see improvement in accuracy with increasing order until (9241); after this, the results are dominated by statistical uncertainty. The impact of statistical uncertainty is greater for the 10 mm mesh size than it was for the 5 mm mesh. This is because the same number of particle histories is used for generating each set of response functions, yet the 10 mm meshes contain 8 times the number of 1 mm^3 regions for scoring energy deposition (TERMA, KERMA, and dose). Thus, it is reasonable to expect uncertainty will be higher, especially for KERMA since it uses a collision estimator rather than a track-length estimator. It is interesting to note the manifestation of this uncertainty in the COMET-PE solution is of a different character than what one would expect to see from uncertainty caused by a purely stochastic method. With COMET-PE, the same pattern of statistical noise repeats in each mesh. In fact, each different material has its own repeating pattern of noise. For example, the pattern seen in the first slab is duplicated in the last slab as well since both are composed of water.

Table 9: Sensitivity of KERMA calculation to photon angular expansion order for the 10 mm mesh size.

	Error		Relative Error	
Expansion	Mean (%)	Max (%)	Mean (%)	Max (%)
(9201)	0.49 ±0.02	2.23 ±0.38	0.82 ±0.04	3.12 ±0.46
(9211)	0.62 ±0.02	1.51 ±0.30	1.32 ±0.04	2.74 ±0.45
(9221)	0.34 ±0.02	1.09 ±0.30	0.71 ±0.04	1.75 ±0.45
(9231)	0.30 ±0.02	0.74 ±0.42	0.65 ±0.04	1.67 ±0.44
(9241)	0.26 ±0.02	0.73 ±0.42	0.52 ±0.04	1.39 ±0.46
(9251)	0.25 ±0.02	0.70 ±0.42	0.50 ±0.04	1.46 ±0.45
(9261)	0.23 ±0.02	0.73 ±0.42	0.45 ±0.04	1.29 ±0.45
(9271)	0.24 ±0.02	0.73 ±0.42	0.46 ±0.04	1.31 ±0.45
(9281)	0.22 ±0.02	0.72 ±0.42	0.43 ±0.04	1.21 ±0.45
(9291)	0.23 ±0.02	0.73 ±0.42	0.44 ±0.04	1.22 ±0.45

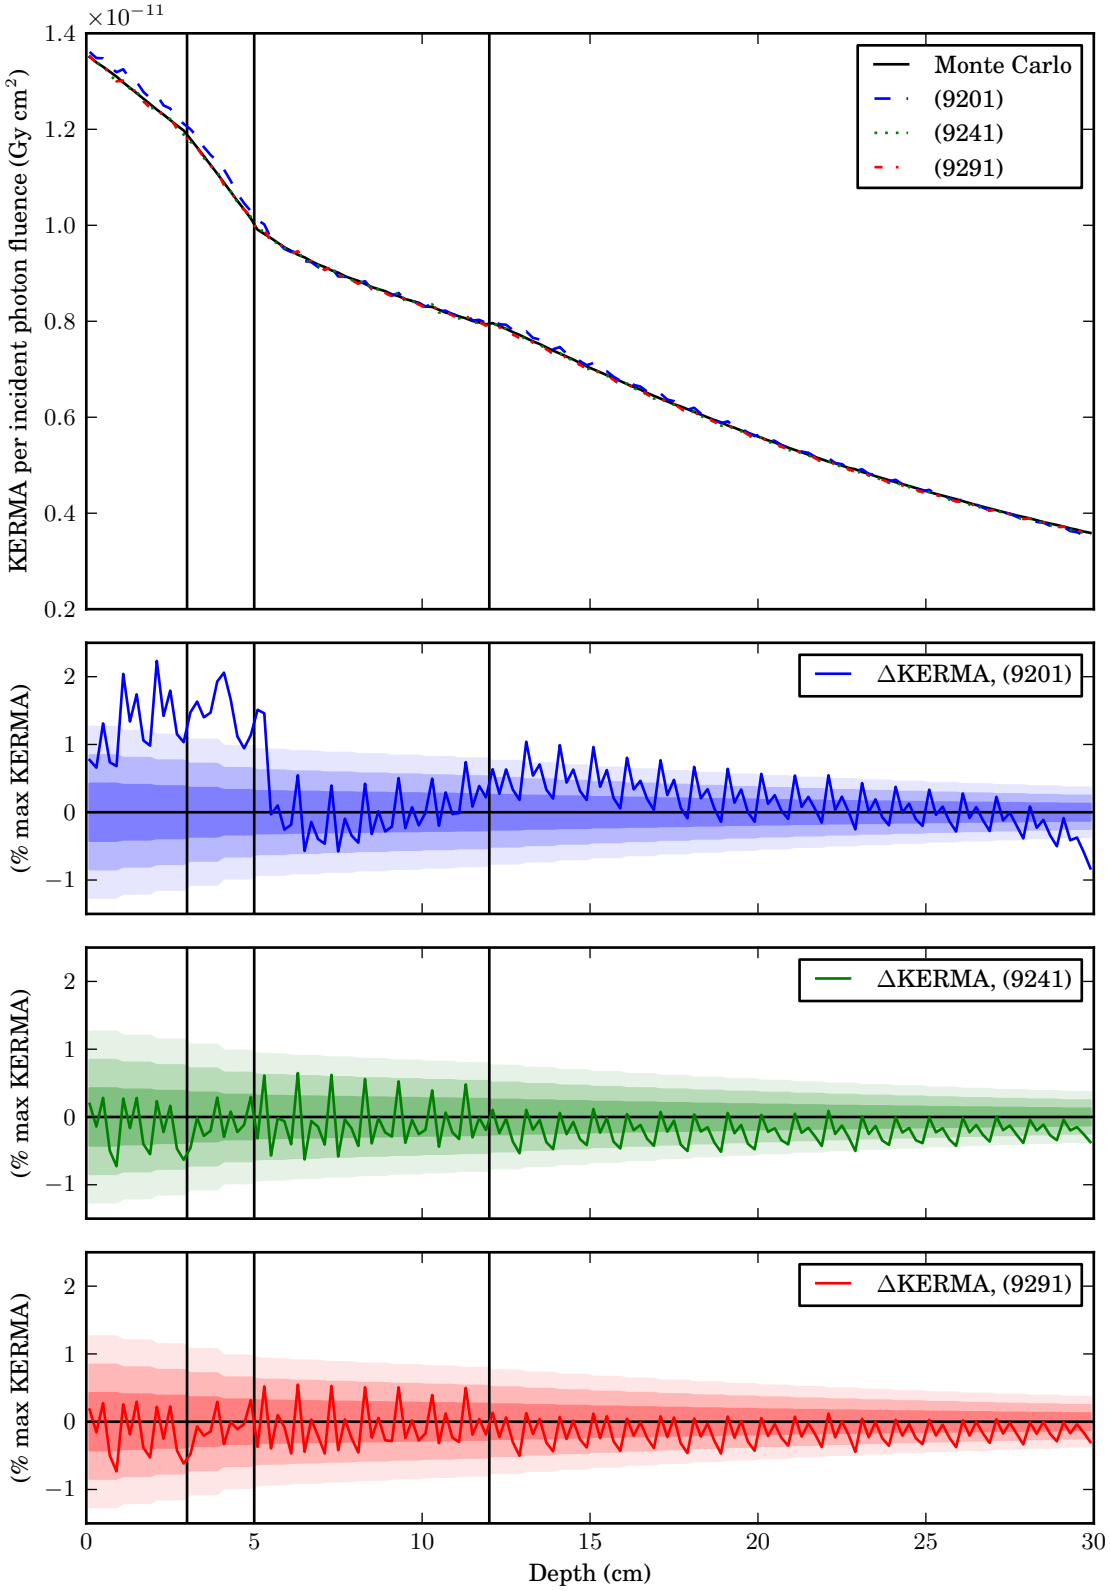

Figure 21: Impact of photon angular order on KERMA for the 10 mm mesh size. Shading indicates uncertainty within 1, 2, and 3 estimated standard deviations, and vertical divisions indicate material interfaces.

The results for varying the photon spatial expansion order for the 10 mm mesh size are given in table 10 and figure 22 on the following page. These results do not clearly favor either (9290) or (9291). This is somewhat surprising, especially when considering the advantage of the first order spatial expansion in the 5 mm mesh size. It is likely that greater statistical precision in the pre-computation would reveal more of a difference here.

Table 10: Sensitivity of KERMA calculation to photon spatial expansion order for the 10 mm mesh size.

	Error		Relative Error	
Expansion	Mean (%)	Max (%)	Mean (%)	Max (%)
(9290)	0.21 ±0.02	0.86 ±0.29	0.37 ±0.04	1.32 ±0.44
(9291)	0.23 ±0.02	0.73 ±0.42	0.44 ±0.04	1.22 ±0.45

6.5.4 Accuracy of Electron Transport (Dose)

The dose calculation involves all parts of the method: analytical integration of the uncollided fluence, photon transport, and electron transport. Since all of these parts except electron transport are tested by other calculations (TERMA and KERMA), the dose calculation is used to test the method's sensitivity to parameters affecting electron transport. The COMET-PE solutions are compared to a reference solution computed with DOSXYZnrc. The Monte Carlo reference solution has average relative uncertainty of 0.07% and maximum relative uncertainty of 0.12%.

The results for varying the electron angular expansion order for the 5 mm mesh size are given in table 11 on page 61 and figure 23 on page 62. The 5 mm COMET-PE solutions each have average relative uncertainty of ~0.03% and maximum relative uncertainty between 0.13% and 0.16%. Uncertainties for the cases (9281i2) for $2 \leq i \leq 7$ were estimated using a calculation of order (924132). The uncertainty for the (928171) and (928170) cases was estimated using order (924131) and (924130) calculations respectively. One can see that accuracy improves with electron angular expansion order up to 7th order, the maximum tested. It is possible that further improvement might be seen by raising the electron

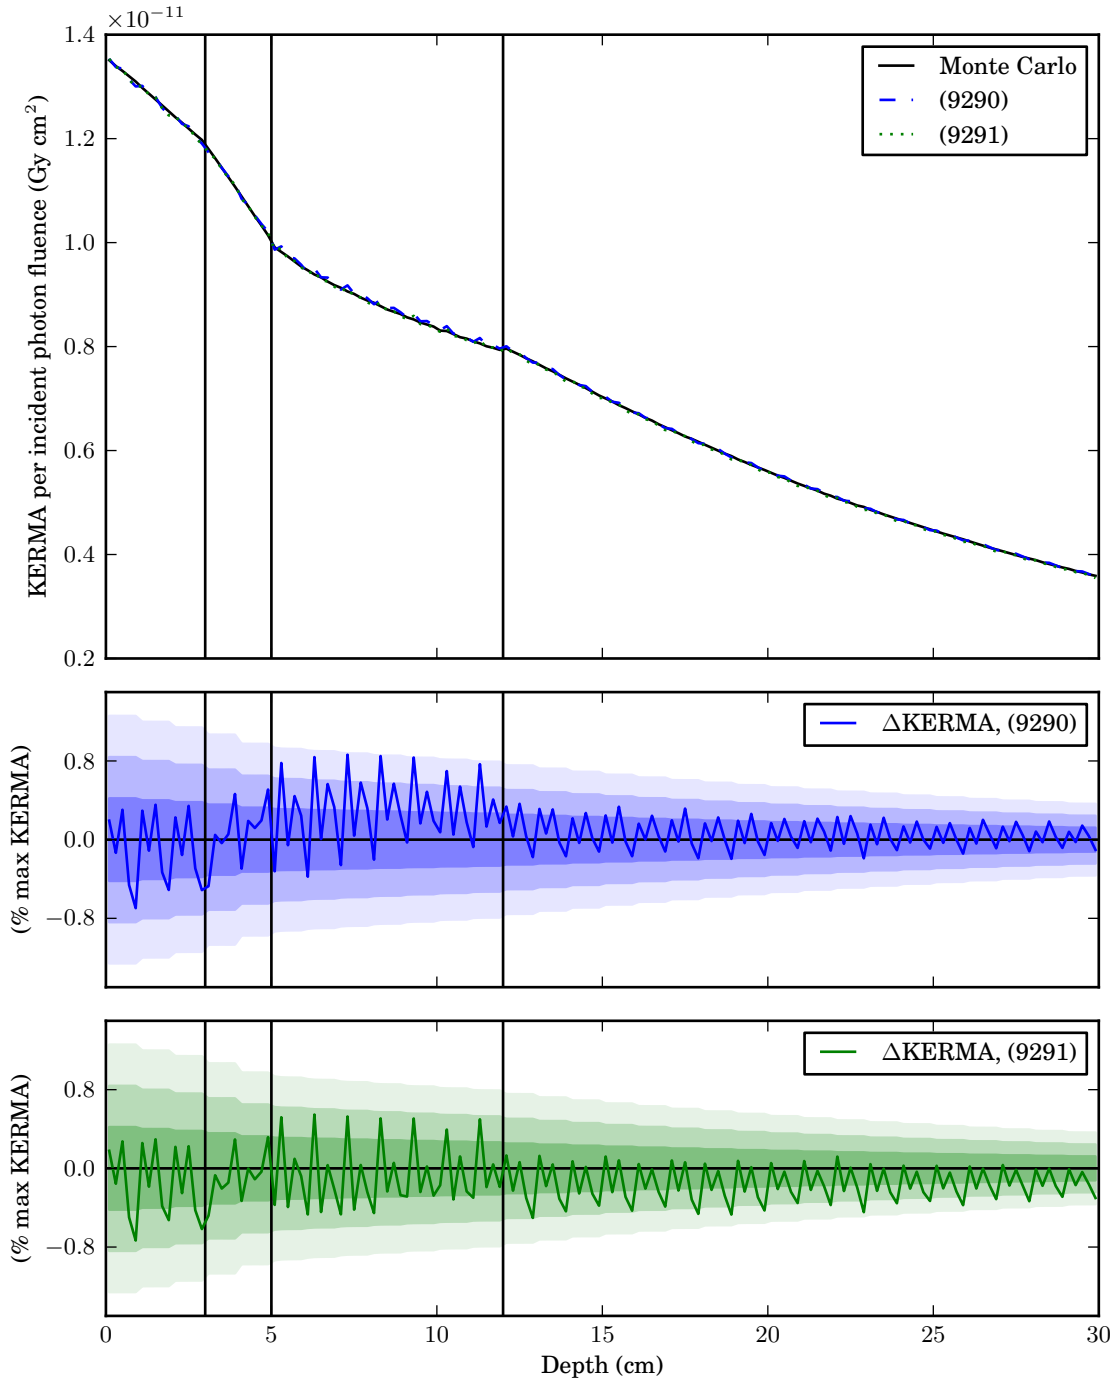

Figure 22: Impact of photon spatial order on KERMA for the 10 mm mesh size. Shading indicates uncertainty within 1, 2, and 3 estimated standard deviations, and vertical divisions indicate material interfaces.

angular expansion order even higher. The electrons set in motion by Compton scattering of primary photons have a highly forwardly-peaked angular distribution. When the electron angular order is set too low, the surface expansion fails to capture this peaked distribution and lateral electron leakage is overestimated.

Table 11: Sensitivity of dose calculation to electron angular expansion order for the 5 mm mesh size.

	Error		Relative Error	
Expansion	Mean (%)	Max (%)	Mean (%)	Max (%)
(928122)	2.86 ±0.00	5.62 ±0.04	6.32 ±0.01	14.72 ±0.09
(928132)	1.78 ±0.00	4.18 ±0.04	3.97 ±0.01	10.87 ±0.09
(928142)	0.99 ±0.00	2.30 ±0.05	2.22 ±0.01	5.47 ±0.10
(928152)	0.72 ±0.00	1.66 ±0.05	1.60 ±0.01	3.66 ±0.10
(928162)	0.56 ±0.00	1.41 ±0.05	1.22 ±0.01	2.65 ±0.09
(928172)	0.49 ±0.00	1.41 ±0.05	1.06 ±0.01	2.76 ±0.09

The results for varying the electron spatial expansion order for the 5 mm mesh size are given in table 12 and figure 24 on page 63. A first order expansion in space results in the highest accuracy of those tested. When the spatial order is set to 0, dose is underestimated in a pattern similar to when the angular expansion order is too low. It appears that the spatial expansion, not just the angular expansion, is important for preventing artificial lateral leakage. This may be caused by the lack of lateral charged particle equilibrium associated with the small 1.5 x 1.5 cm^2 field size.

Table 12: Sensitivity of dose calculation to electron spatial expansion order for the 5 mm mesh size.

	Error		Relative Error	
Expansion	Mean (%)	Max (%)	Mean (%)	Max (%)
(928170)	3.74 ±0.00	8.61 ±0.06	7.54 ±0.01	20.30 ±0.22
(928171)	0.46 ±0.00	1.13 ±0.05	1.01 ±0.01	2.23 ±0.09
(928172)	0.49 ±0.00	1.41 ±0.05	1.06 ±0.01	2.76 ±0.09

The results for varying the electron angular expansion order for the 10 mm mesh size are given in table 13 on page 64 and figure 25 on page 65. The 10 mm COMET-PE solutions

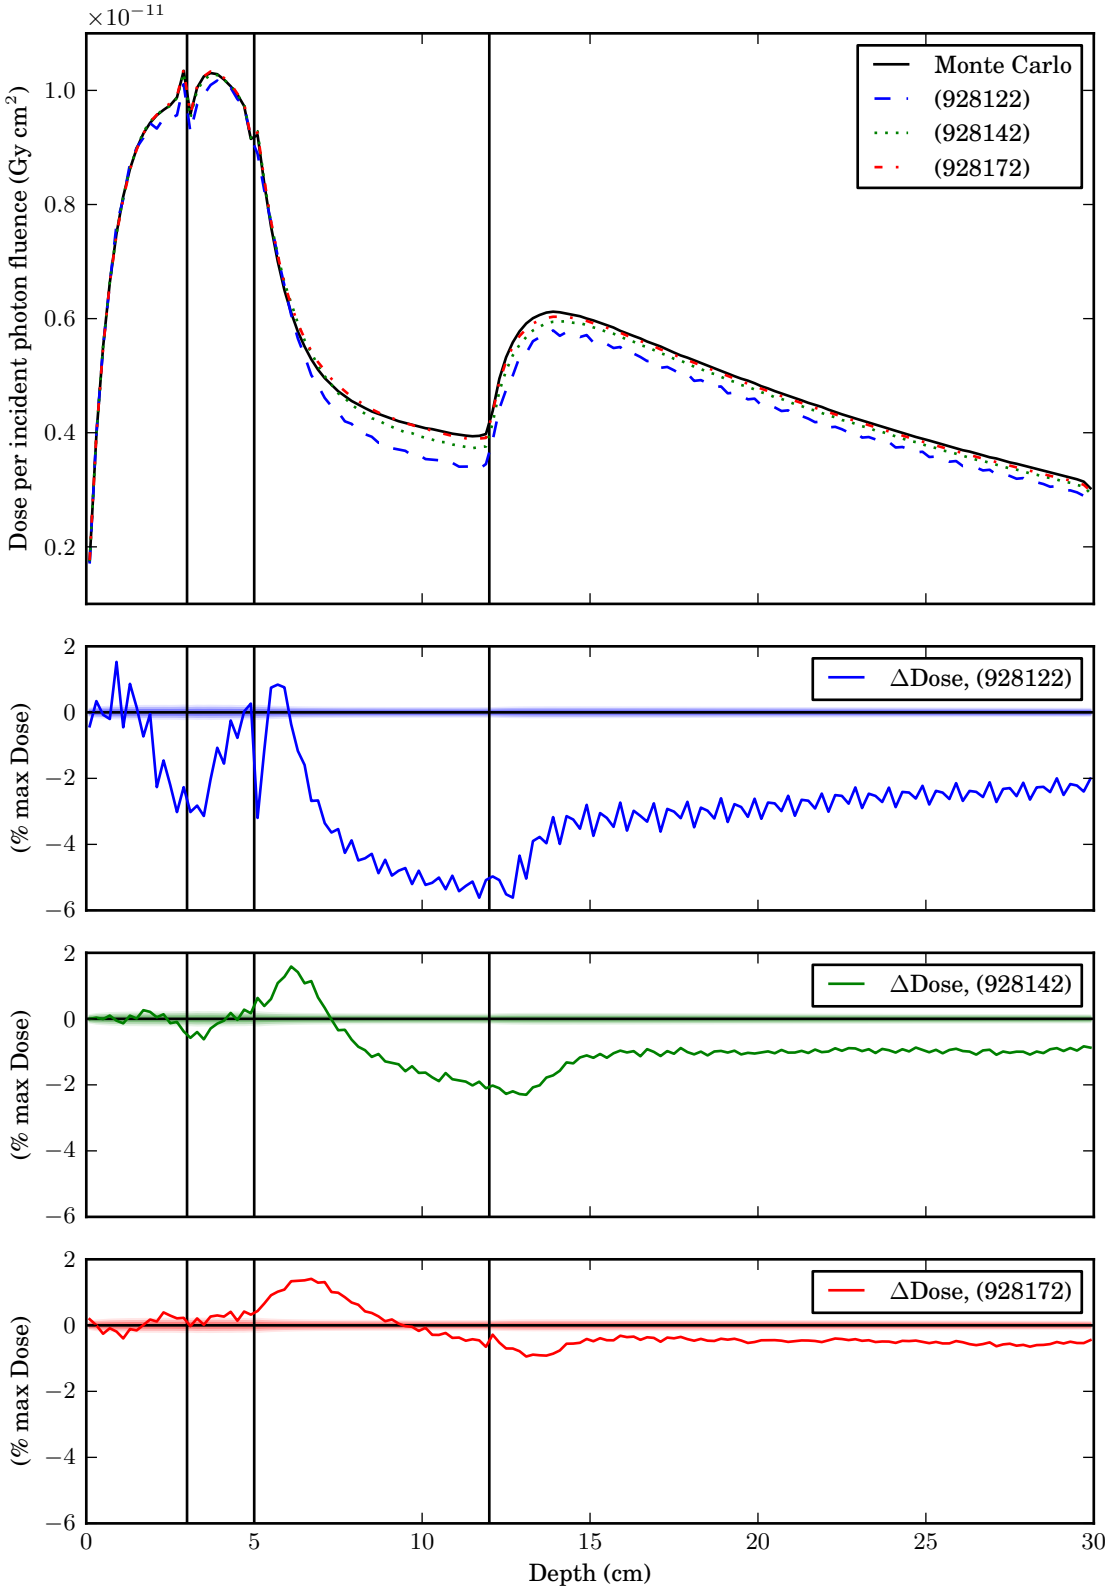

Figure 23: Impact of electron angular order on dose for the 5 mm mesh size. Shading indicates uncertainty within 1, 2, and 3 estimated standard deviations, and vertical divisions indicate material interfaces.

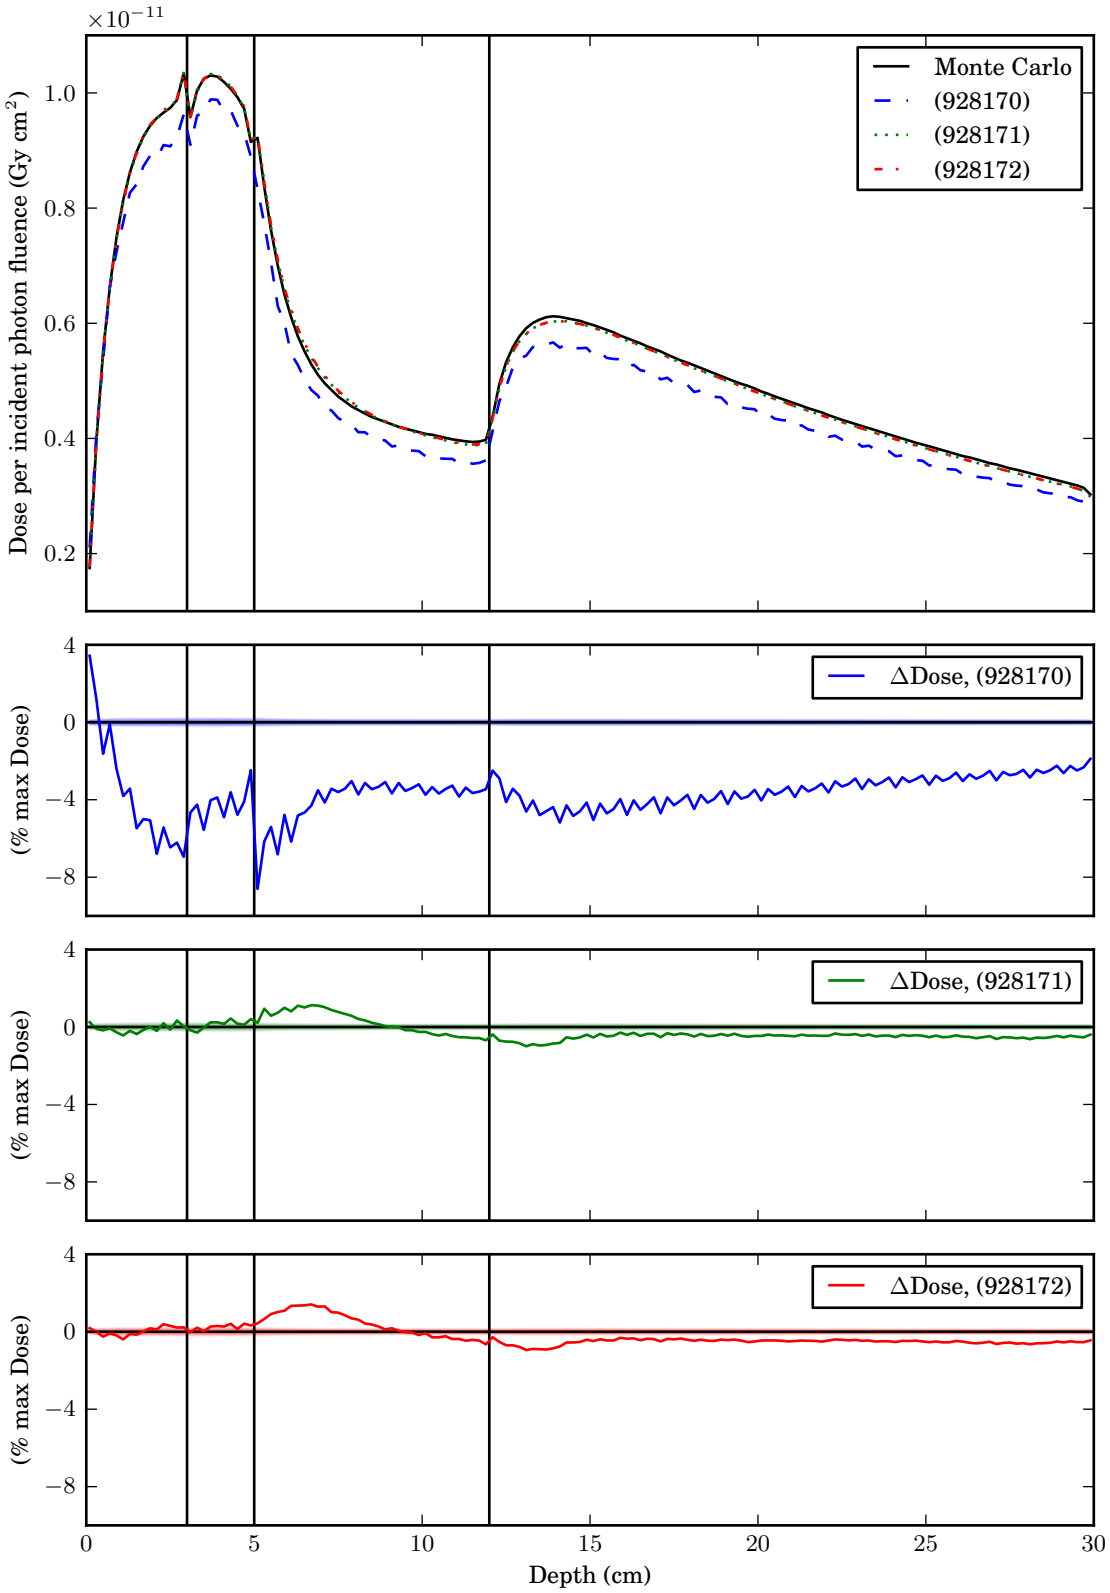

Figure 24: Impact of electron spatial order on dose for the 5 mm mesh size. Shading indicates uncertainty within 1, 2, and 3 estimated standard deviations, and vertical divisions indicate material interfaces.

each have average relative uncertainty of ~0.09% and maximum relative uncertainty of ~0.55%. For the 10 mm cases, uncertainty is computed for all cases except (924172), which is assumed to have the same uncertainty distribution as (924162). Accuracy improves with increasing electron angular order up to (924162). Overall, the behavior with angular order is similar to that seen with the 5mm mesh size. When the angular order is too low, the dose along the beam's central axis is underestimated, especially in the lung region where electrons travel the farthest.

Table 13: Sensitivity of dose calculation to electron angular expansion order for the 10 mm mesh size.

	Error		Relative Error	
Expansion	Mean (%)	Max (%)	Mean (%)	Max (%)
(924122)	1.76 ±0.01	5.01 ±0.04	3.82 ±0.01	13.06 ±0.10
(924132)	0.94 ±0.01	3.20 ±0.05	2.00 ±0.01	8.39 ±0.11
(924142)	0.53 ±0.01	2.17 ±0.05	1.04 ±0.01	3.90 ±0.10
(924152)	0.38 ±0.01	1.86 ±0.06	0.69 ±0.01	3.33 ±0.10
(924162)	0.34 ±0.01	1.43 ±0.05	0.61 ±0.01	2.83 ±0.10
(924172)	0.35 ±0.01	1.53 ±0.05	0.60 ±0.01	3.09 ±0.11

The results for varying the electron spatial expansion order for the 10 mm mesh size are given in table 14 on page 67 and figure 26 on page 66. In this case, accuracy improves with increasing electron spatial order up to 2nd order, the maximum tested. The pattern seen in the error for the (924170) case is particularly interesting. When the electron spatial order is too low, dose is underestimated, and it is underestimated most on the shallowest face of each mesh. This indicates that there is a failure of–what should be–transient charged particle equilibrium; the electrons liberated within each mesh as a result of the collisions of primary photons are not being adequately replaced by electrons flowing downstream across mesh interfaces. Thus, there is a small build-up region within each mesh.

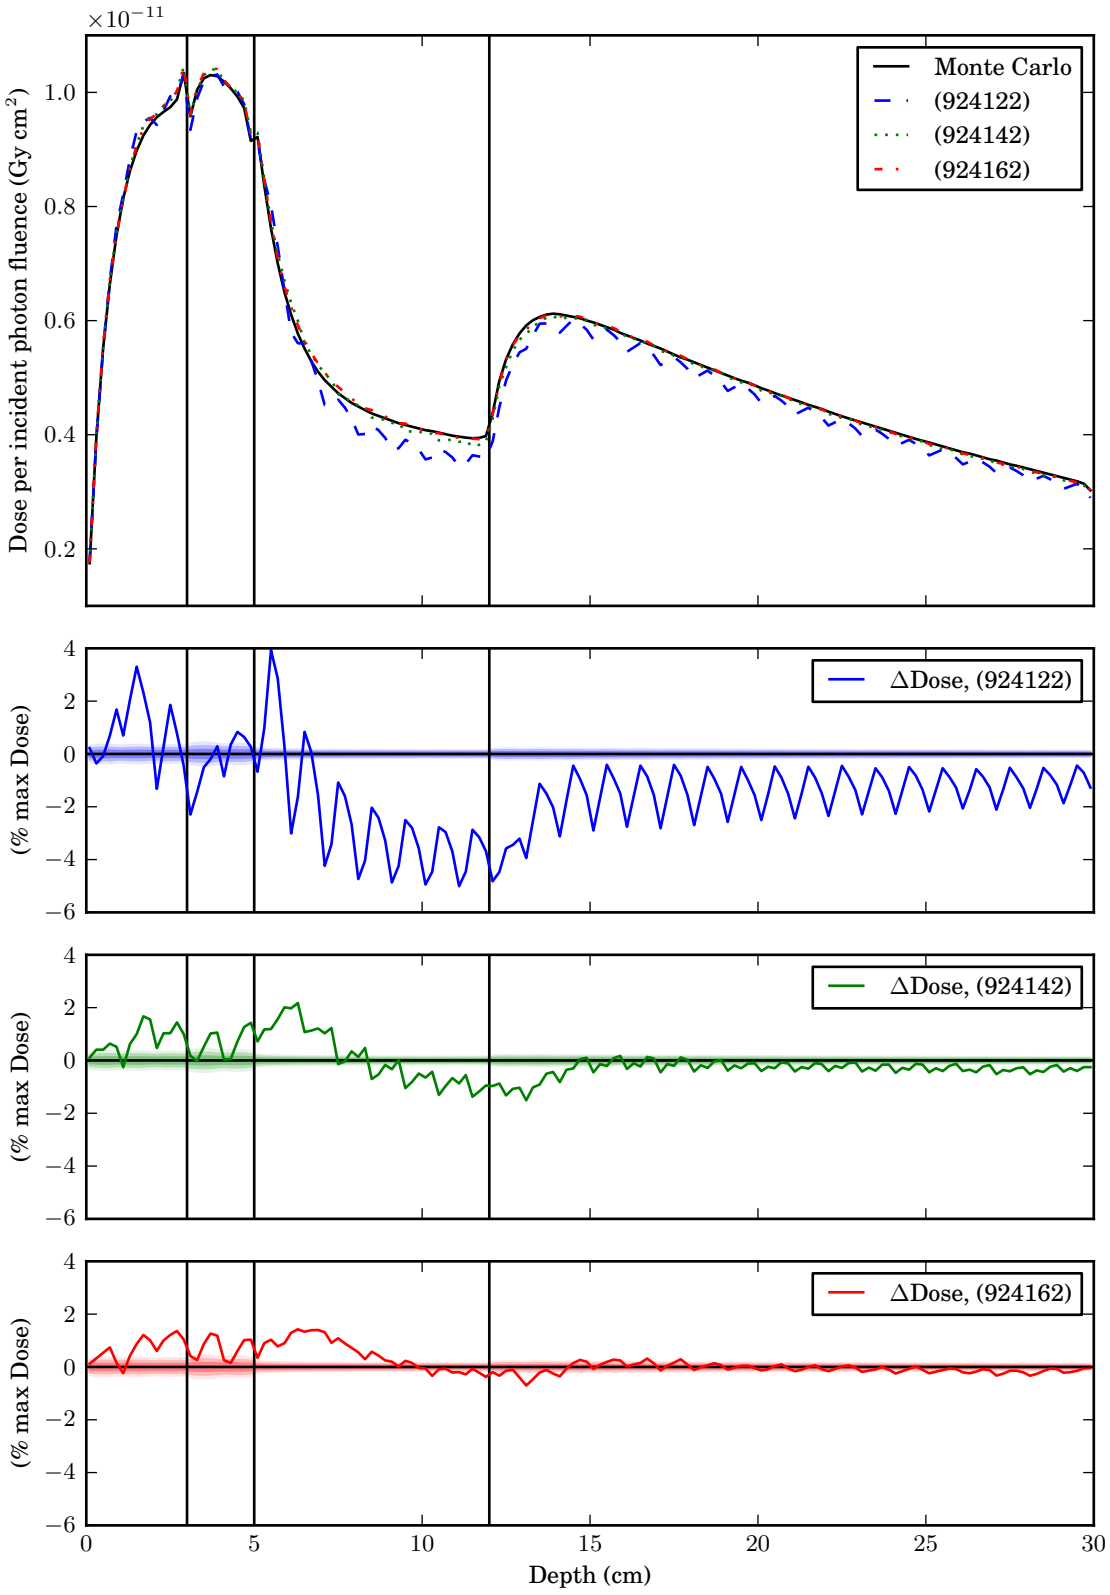

Figure 25: Impact of electron angular order on dose for the 10 mm mesh size. Shading indicates uncertainty within 1, 2, and 3 estimated standard deviations, and vertical divisions indicate material interfaces.

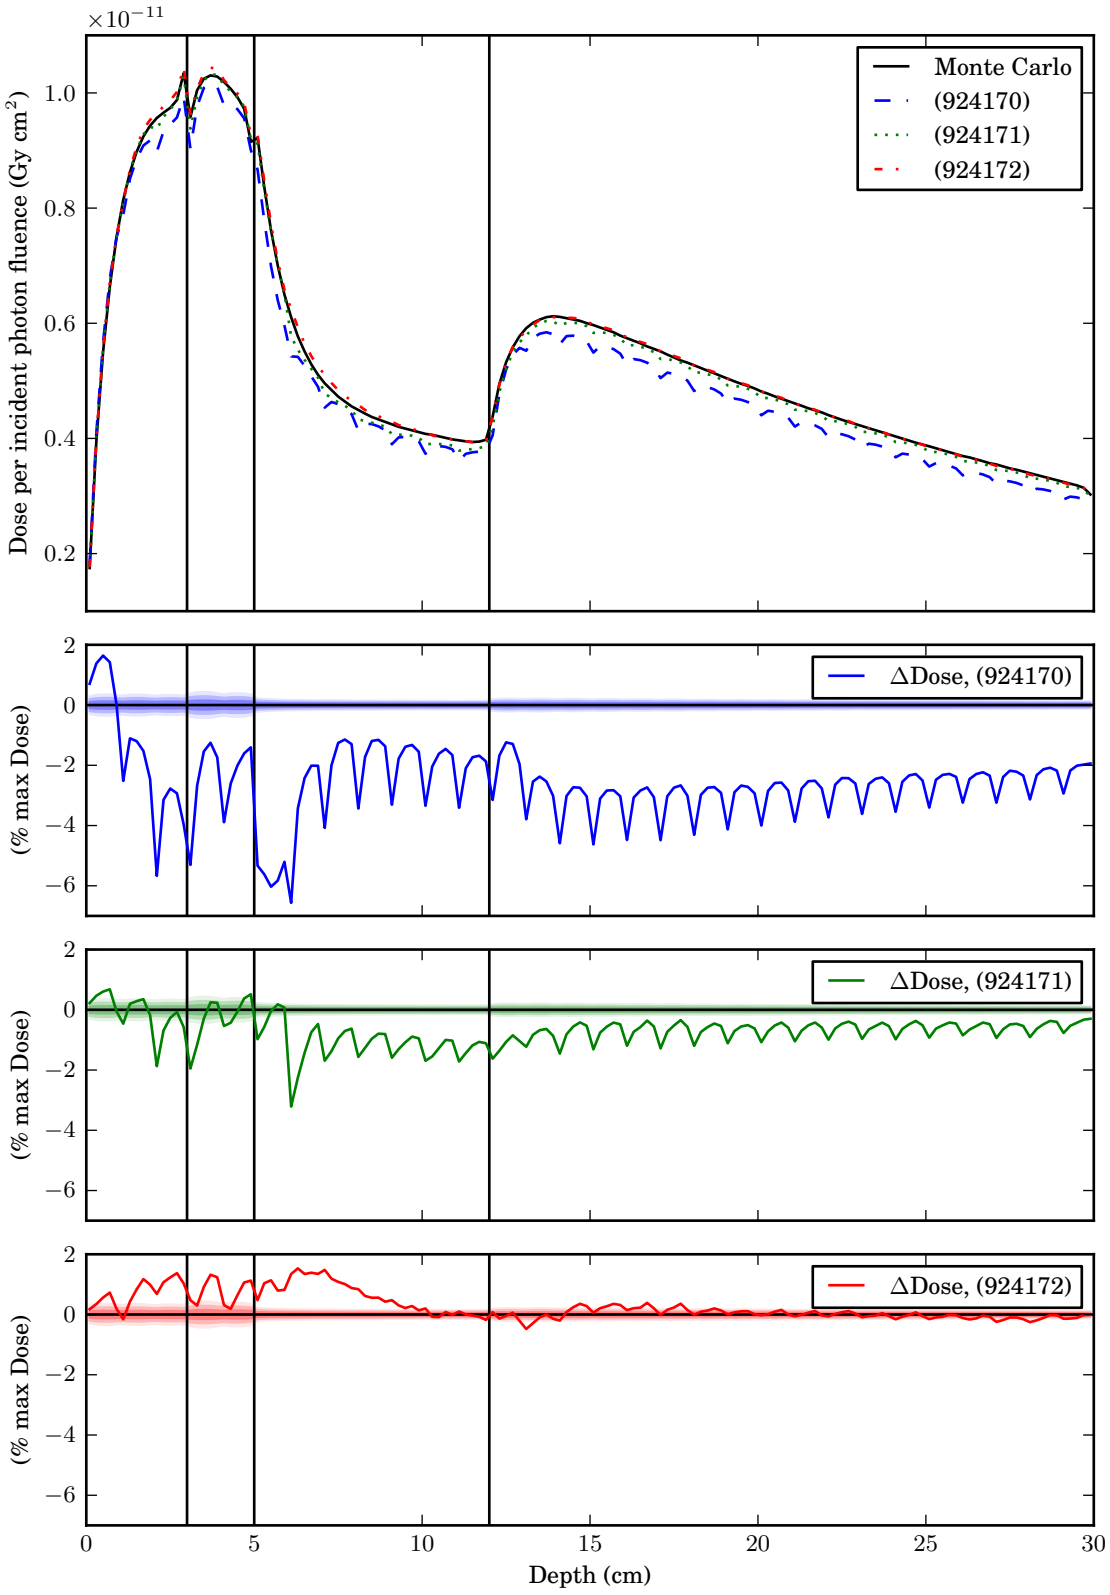

Figure 26: Impact of electron spatial order on dose for the 10 mm mesh size. Shading indicates uncertainty within 1, 2, and 3 estimated standard deviations, and vertical divisions indicate material interfaces.

Table 14: Sensitivity of dose calculation to electron spatial expansion order for the 10 mm mesh size.

	Error		Relative Error	
Expansion	Mean (%)	Max (%)	Mean (%)	Max (%)
(924170)	2.67 ±0.01	6.57 ±0.06	5.50 ±0.01	11.13 ±0.09
(924171)	0.76 ±0.01	3.21 ±0.06	1.62 ±0.01	5.44 ±0.09
(924172)	0.35 ±0.01	1.53 ±0.05	0.60 ±0.01	3.09 ±0.11

6.5.5 Timing

Timing information is given for the most economical calculations that result in maximum relative error of less than 3%. For the 10 mm mesh size, the (924162) case requires just over 23 minutes on one compute node. For the 5 mm mesh size, the (928171) case requires 6.0 hours on one compute node. Compared to the Monte Carlo reference (if run on a single node) the 10 mm case is 31 times faster and the 5 mm case is twice as fast.

6.6 Discussion

Overall, COMET-PE is highly accurate for both KERMA and dose calculation in this benchmark problem. Maximum error of less than 1.5% is achieved for both 5 mm and 10 mm mesh sizes. The dose results are comparable to those found in a comparison between the GATE Monte Carlo method and EGSnrc [22]. The small differences between DOSXYZnrc and COMET-PE arise from statistical uncertainty and from truncation error. The statistical uncertainty has two sources: the response coefficients and the reference solution, both of which were computed using stochastic methods. The statistical uncertainty is most pronounced in the KERMA distributions. The uncertainty is lower for the dose distributions because the dose was tallied with a track-length estimator instead of the relatively inefficient collision estimator that was used for KERMA. For both the dose and KERMA solutions, the 10 mm case has larger uncertainty than the 5 mm case because the 10 mm meshes have 8 times the number of 1 mm^3 scoring voxels. It is reasonable to expect that better agreement would be obtained if uncertainties were smaller.

Truncation error is caused by failure of the truncated expansion to fully represent the true fluence distribution. The impact of truncation error is most obvious in the lung build-down region (5 cm < depth < 7 cm), where electrons born upstream in the aluminum slow down and deposit dose. The physics of this process necessitate a well-resolved energy spectrum because electron track length is most sensitive to energy in the low-density lung. A 1% difference in the energy of a 6 MeV electron is enough to change its range by 1 mm. Because the electron bin structure is relatively unrefined, and because the 1 mm shift is almost identical in both the 5 mm and 10 mm cases, it is likely that this difference is caused by a deficiency of electron energy resolution. In general, the truncation errors are small, and could likely be improved by optimizing both the bin structure and expansion orders.

The timing information provided in this chapter is not sufficient to compare the relative efficiency of DOSXYZnrc and COMET-PE for several reasons. First, DOSXYZnrc scored energy deposition on a 5 x 5 x 2 mm^3 grid compared with COMET-PE, which computed energy deposition on a much finer 1 x 1 x 1 mm^3 grid. With COMET- PE, this high level of detail comes without penalty thanks to the use of pre-computation. It is also unfair to compare these calculation times because COMET-PE modeled a phantom that was shortened to 11 x 11 x 30 cm^3 while DOSXYZnrc modeled the entire 30.5 x 39.5 x 30 cm^3. phantom. Finally, DOSXYZnrc was run on 12 nodes of the cluster while COMET-PE was only run on a single node. The next two chapters provide timing results that are more suitable for comparison.

CHAPTER VII

WATER PHANTOM

Measurements in water phantoms form the basis for the commissioning of linear accelerators and treatment planning systems. It is therefore important to calculate dose accurately in water. In this chapter a standard water phantom is used to assess the accuracy and efficiency of COMET-PE when compared with Monte Carlo.

A composite analysis, combining relative dose difference and distance to agreement, is used to measure the accuracy of the method [13, 28]. A voxel is considered "passing" if it has either relative error within the prescribed tolerance (%) or distance to agreement within the specified tolerance (mm). When these methods are used alone, relative error is too sensitive in regions with high dose gradients and distance to agreement is too sensitive within regions of low dose gradients. Combining the two methods provides a more meaningful analysis. For the sake of simplicity, the algorithm used to compute distance to agreement searches only along the three coordinate axes away from each voxel. This means that the true distance to agreement can be conservatively underestimated, resulting in failures for some voxels that should actually pass.

7.1 Benchmark Specifications

The phantom is a 30 x 30 x 30 cm^3 cube of water. A 6MV point source illuminates the phantom from source to surface distance of 100 cm. The beam is collimated to 10 x 10 cm^2 field on the phantom surface. The energy spectrum is pictured in figure 8 on page 39.

7.2 Monte Carlo Reference Solution

The reference solution is computed using DOSXYZnrc. The physics settings are given by table 18 on page 89. Dose is tallied in 4 x 4 x 4 mm^3 voxels throughout the phantom.

For this calculation, 1.2×10^9 source histories were followed, consuming 9.8 hours on 8 compute nodes or a total of 64 CPU cores. The resulting solution has average relative uncertainty of 0.17%. The maximum relative uncertainty in voxels receiving at least 10% of the maximum dose is 0.46%. The 8-fold symmetry of this problem is used to reduce the Monte Carlo uncertainty by averaging each dose value with its image under each of the 7 non-trivial symmetric transformations. After symmetrization, the solution has average relative uncertainty of 0.06%, and the maximum relative uncertainty in voxels receiving at least 10% of the maximum dose is 0.16%. These uncertainty values are underestimated, however, because covariances between symmetric voxels are not accounted for.

7.3 COMET-PE Solutions

COMET-PE solutions are computed using the same response library that is used for the slab phantom. These computations are performed with the 10 mm mesh size, and truncation order is set to (914162). The convergence criterion is set to the relatively conservative value of $\varepsilon = 10^{-5}$. COMET-PE outputs dose for 1 x 1 x 1 mm^3 voxels throughout the phantom. For comparison with Monte Carlo, these values are then averaged over 4 x 4 x 4 mm^3 voxels. The resulting solution has average relative uncertainty of 0.17%. The maximum relative uncertainty in voxels receiving at least 10% of the maximum dose is 0.43 %. The uncertainty was computed in a calculation that took just under 45 minutes on 8 compute nodes. To test the method's ability to scale with distributed-memory parallelism, the calculation is re-run with different numbers of compute nodes, from 1 to 16. In these scaling tests, uncertainty is not estimated.

7.4 Results: Accuracy

A comparison of the COMET-PE and and Monte Carlo dose distributions is given in figure 27 on the next page. The phantom is composed of 421 875 voxels. The 66 715 ($\sim$ 16%) of these voxels that receive dose greater than 10% of the maximum dose in the reference

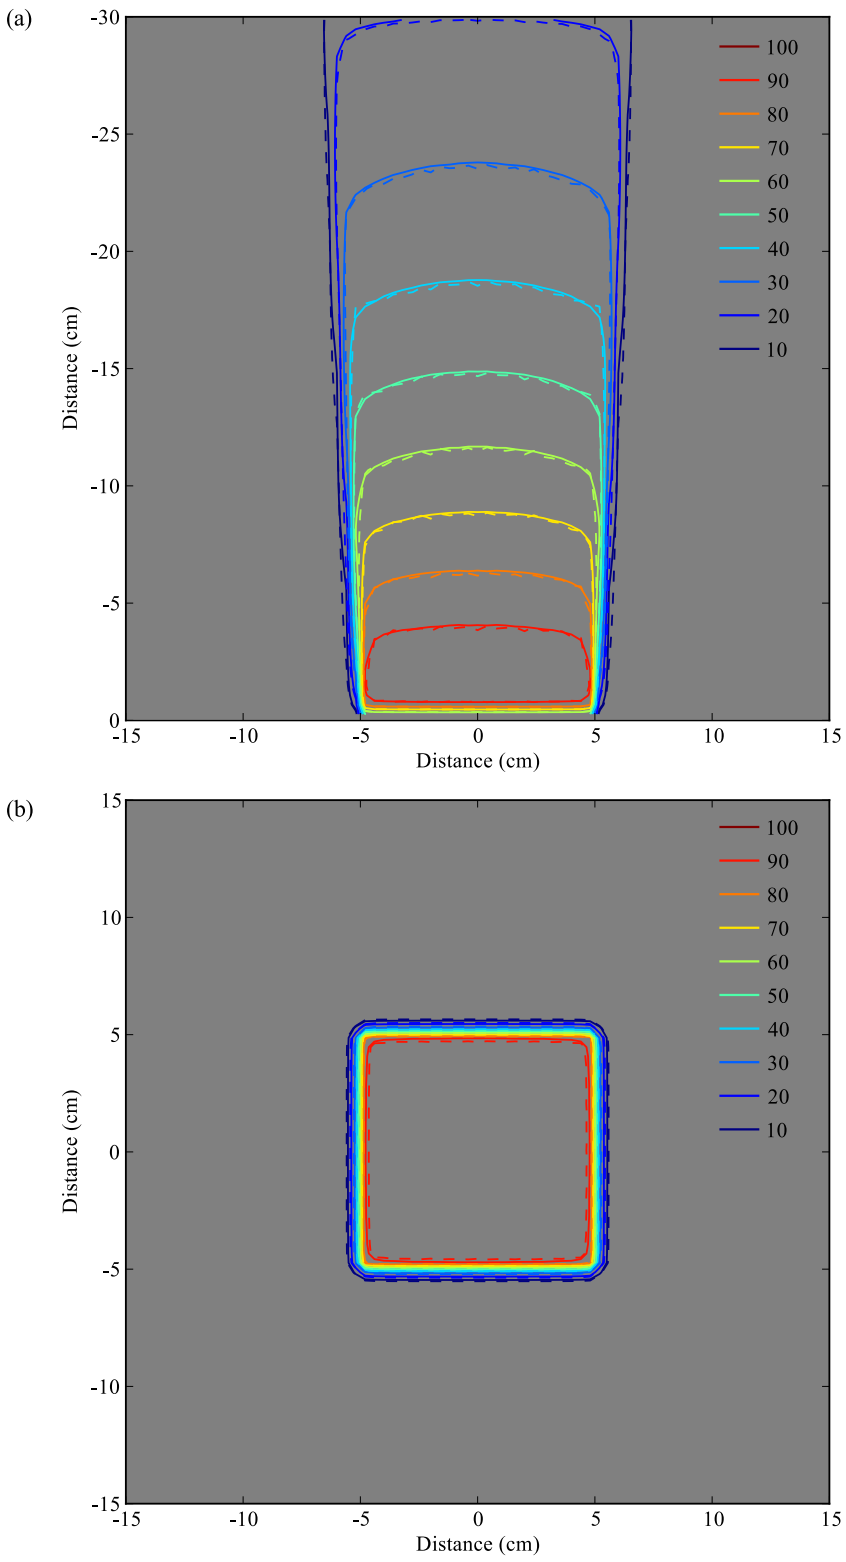

Figure 27: Comparison of isodose contours for the water phantom. Dose is given as a percentage of the maximum reference dose. Solid lines represent the Monte Carlo solution, and dashed lines represent the COMET-PE solution. Part (a) shows a slice through the central axis, and (b) shows a slice perpendicular to the central axis at depth = 2.2 cm.

solution are evaluated for agreement. The results are given in table 15. The 168 voxels that

Table 15: Agreement of voxels receiving greater than 10% of the maximum reference dose.

Criterion	# Passing voxels	% Passing voxels
1% / 1 mm	40 601	60.857
2% / 2 mm	64 921	97.311
3% / 3 mm	66 547	99.748
4% / 4 mm	66 650	99.902
5% / 5 mm	66 715	100.00

failed the 3% / 3 mm test are shown in figure 28 on the next page. All of the failures occur at the corners of the beam. Dose was underestimated at each of the 168 points. Before symmetrization, there were 6 additional failing voxels closer to the beam's central axis in the deepest 2 cm.

7.5 Results: Performance

To evaluate the distributed-memory performance, the calculation was run on 1, 2, 4, 8, 12, and 16 compute nodes. Since each compute node has 8 CPU cores, this represents scaling from 8 up to 96 CPU cores. Figure 29 on page 74 shows speedup as a function of the number of compute nodes. For this performance analysis, the code is considered to consist of three parts:

1. Source integration: the time required for the adaptive cubature to evaluate integrals of the uncollided flux.

2. Response calculation: the time spent applying response functions, most of which is consumed by the fluence iterations.

3. Input/output: the time spent reading (mostly response functions) and writing (dose) to disk.

A breakdown of the time spent on each task is given by figure 30 on page 74. The single-

(a)

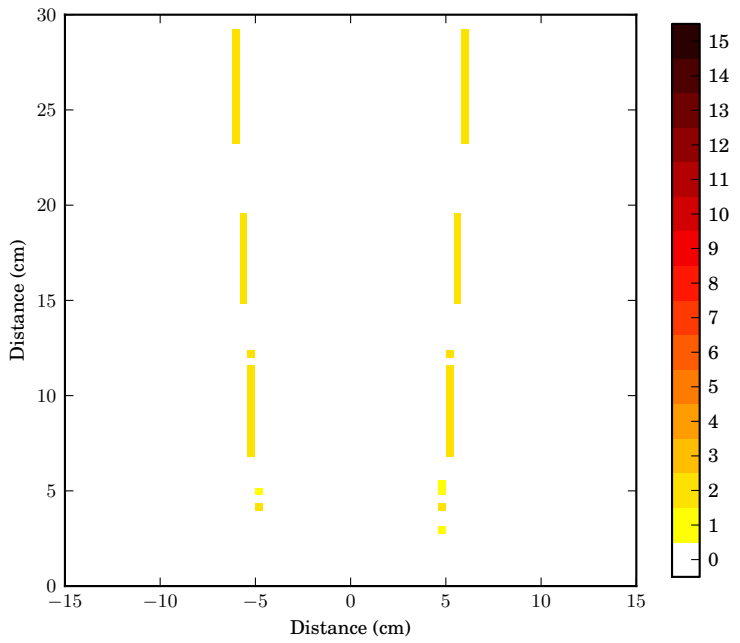

(b)

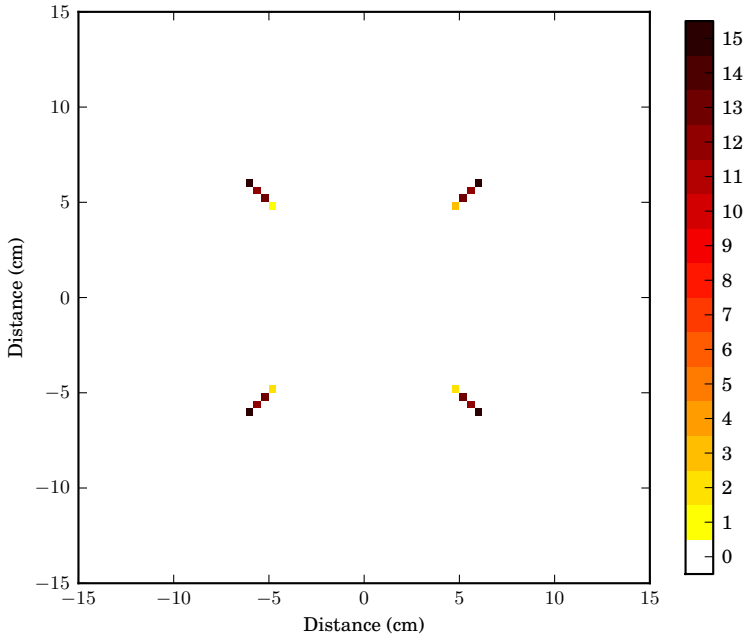

Figure 28: Voxels failing the 3% / 3 mm test. Each pixel represents the total number of failing voxels summed along the direction perpendicular to the image plane. Part (a) shows a plane through the central axis, and (b) shows a plane perpendicular to the central axis.

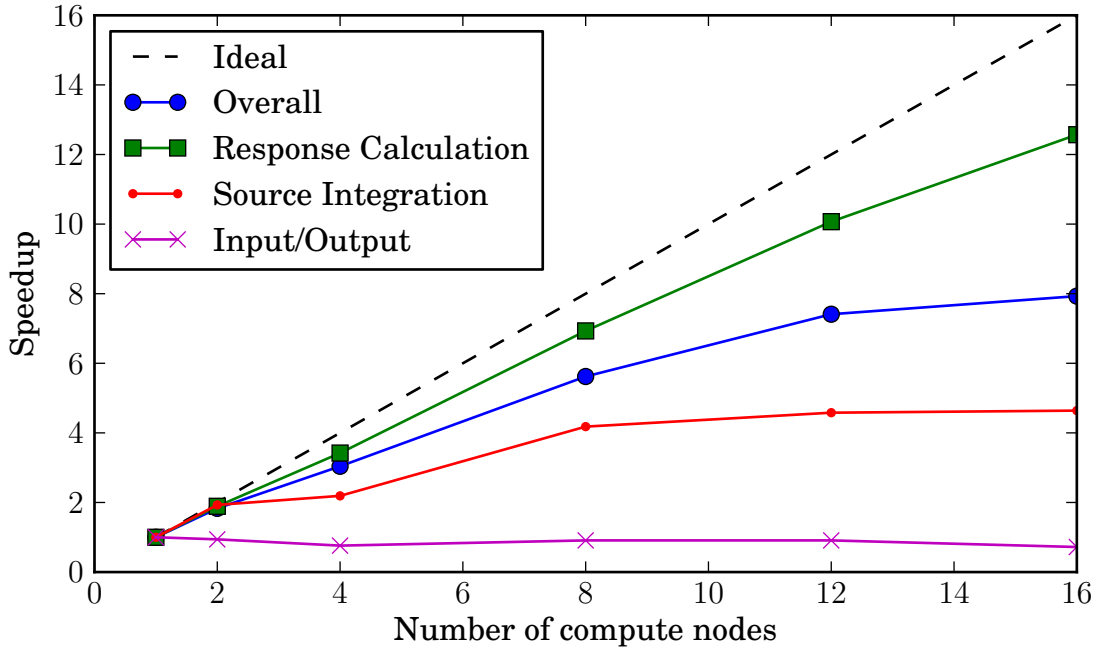

Figure 29: Parallel scaling for distributed memory computation. Speedup is the time required for a many-node calculation divided by the time required for a single-node calculation.

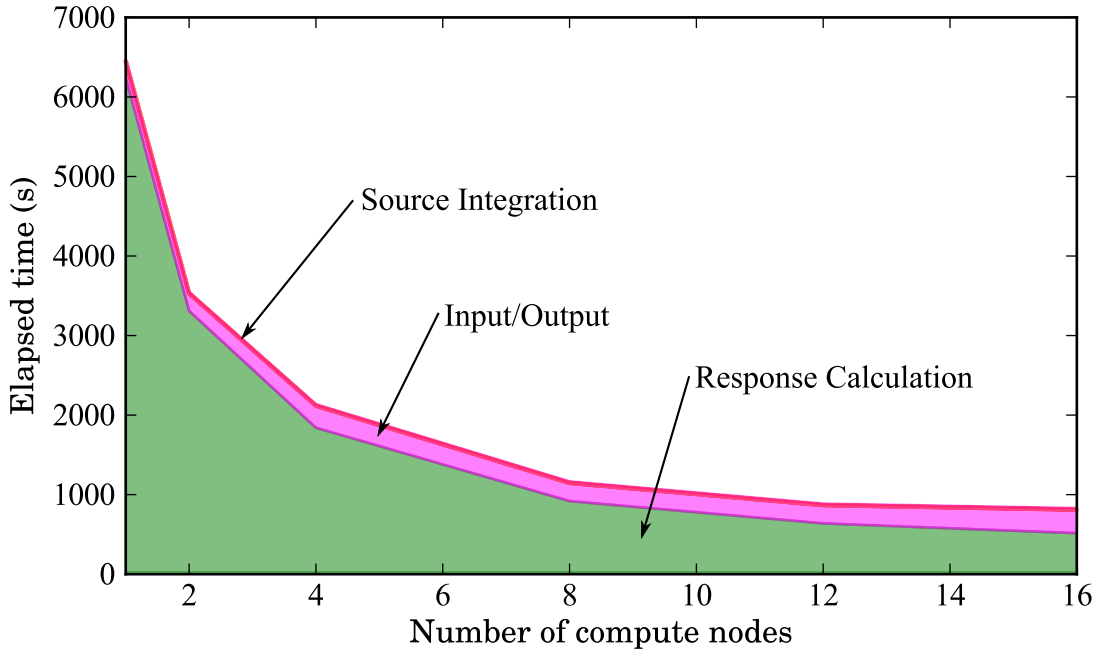

Figure 30: Time spent on each phase of computation.

node computation took 108 minutes while the 16-node computation takes less than 14 minutes, representing a speedup of nearly 8 times. The source integration uses less than 0.16% of the total computation time in all cases. Response computations dominate the total effort, taking 61% of the time on 16 nodes and 96% of the time on a single node. When running on 8 nodes, the same number used for the reference solution, the COMET-PE solution takes less than 20 minutes.

7.6 Discussion

COMET-PE is shown to be highly accurate and efficient for this water phantom problem. The greater than 97% pass rate for the 2% / 2 mm criterion is excellent, and the small discrepancies that do exist may be overestimated. Most of the voxels that fail the 3% / 3 mm test are located on the corners of the beam. Due to the limitations of the method used to calculate distance to agreement, it is likely that many of these failing voxels are actually within 3 mm in the diagonal direction.

The performance results demonstrate excellent scaling for the response computations. The input/output portions of the code are a clear limit to performance when scaling to more than ~8 compute nodes, but this is expected since there has been no effort to parallelize disk access. The scaling of the source integration is related to the parallel domain decomposition. Most of the source integration effort is spent on meshes that straddle the beam edge, so parallel efficiency is affected by the way these edges are distributed among compute nodes. Overall, the source integration is highly efficient, using less than 0.16% of the total computational effort. This represents an advantage for COMET-PE when it comes to problems with complicated source distributions such as those encountered in rotational arc therapy. Superposition methods scale in complexity in proportion to the number of beam angles. By contrast, the COMET-PE calculation time is only weakly dependent on the source complexity; thus it is ideally suited for these applications.

Besides scaling well COMET-PE is efficient. When both methods are run on 8 compute

nodes, COMET-PE is over 13 times faster than Monte Carlo with uncertainty estimation enabled and over 30 times faster without uncertainty estimation. Testing shows that the uncertainty estimated with COMET-PE is insensitive to expansion order. When uncertainty calculation is enabled, however, the same expansion order is used for the fluence calculation as is used for the uncertainty calculation. This means that about 2/3 of the floating point operations are used to estimate uncertainty, and only 1/3 go toward dose calculation. For this reason, an input option will be added to select a response truncation order for the uncertainty calculation independent of the setting used for the fluence. It is expected that these lower order uncertainty calculations will consume an order of magnitude less computational effort without sacrificing the quality of the estimate.

CHAPTER VIII

LUNG CT PHANTOM

In the clinical setting, patient-specific dose calculations use computed tomography (CT) data to provide information on patient geometry and material composition. To further test the capabilities of COMET-PE, a CT-based lung benchmark problem is used.

8.1 Creation of the Benchmark Problem

Given a CT scan, one must somehow derive data on the density and material composition to use for dose calculation. In general, this problem is difficult because it is underdetermined; there is not enough information in (single-energy) CT data to uniquely determine both density and material composition for each voxel. Although this problem is certainly important, it is shared by all transport methods, including Monte Carlo [14]. This issue, along with the use of interpolation to derive response functions for a continuous range of densities, is left to future work. For this work, material conversions are performed as specified by table 16. This conversion is similar to the default conversion performed by the CTCREATE

Table 16: CT to material conversion.

Range of Hounsfield units	Material	Density (g/cm^3)
-1000 $\leq$ HU $<$ -950	Air	1.2048×10^{-3}
-950 $\leq$ HU $<$ -700	Lung	0.26
-700 $\leq$ HU $<$ 125	Soft Tissue	1.0
125 $\leq$ HU $<$ 2000	Bone	1.85

tool [51], but density is mapped to a set of discrete values. To create a phantom compatible with a 10 mm mesh, the CT data was first re-sampled to 10 mm voxels before material conversion. This process is depicted by figure 31. A GUI application, called CTgui, was

(a)

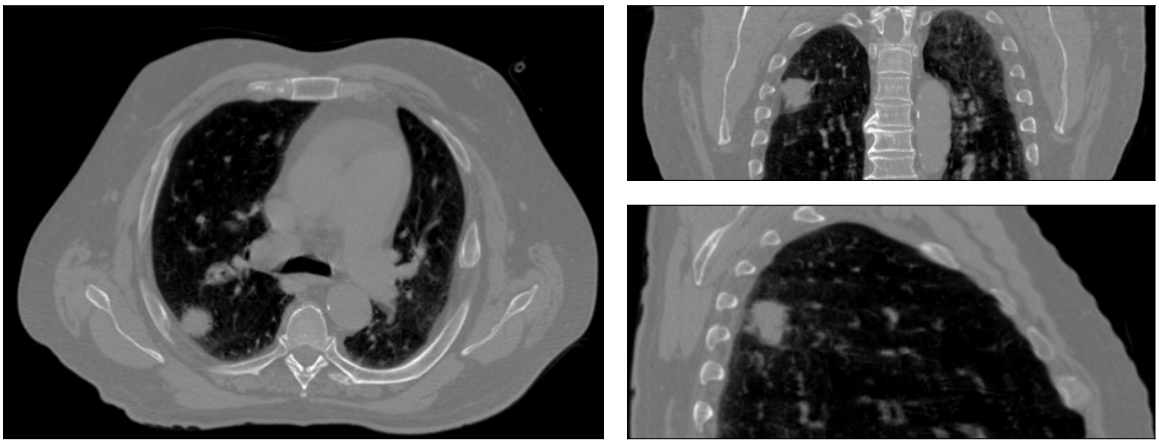

(b)

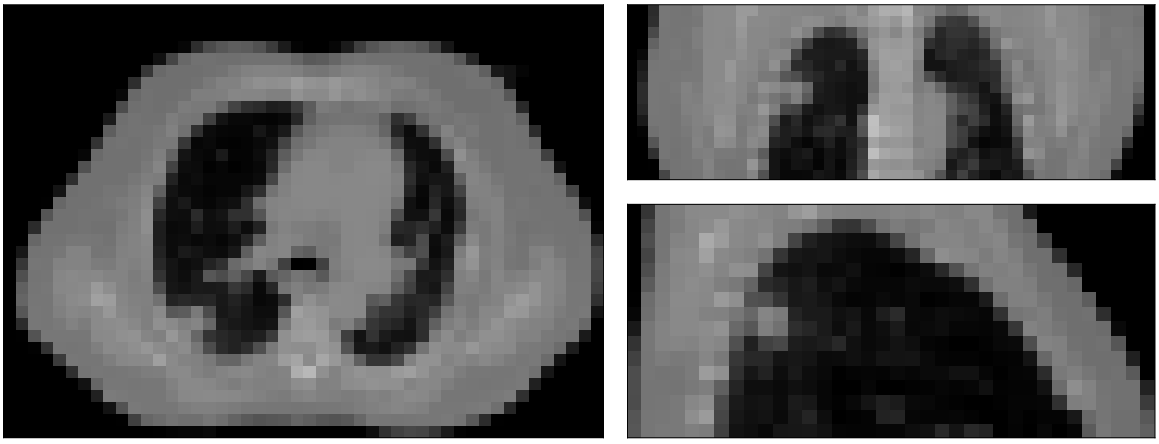

(c)

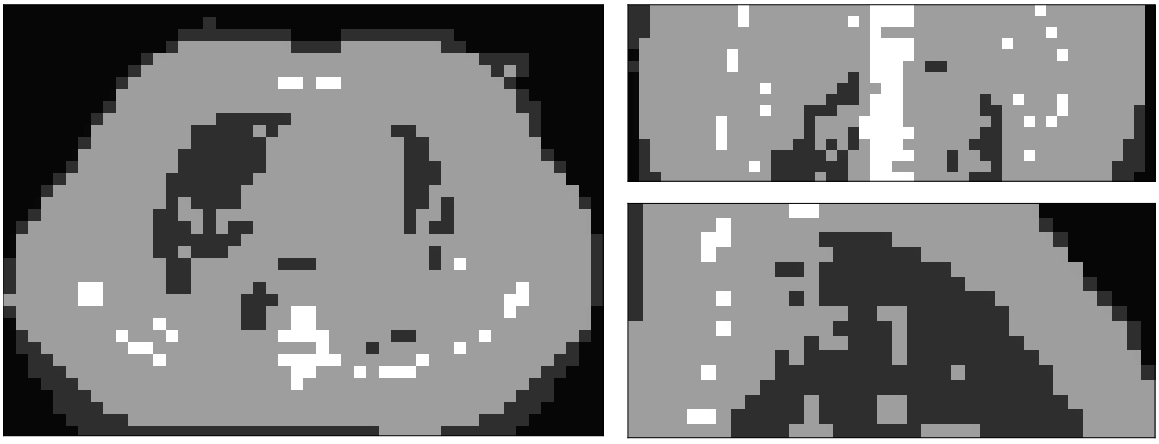

Figure 31: Creation of the lung phantom from CT data. The original CT scan data is shown in part (a). Part (b) shows the scan data resampled to 10 x 10 x 10 mm^3 voxels. Part (c) shows the discretized phantom.

developed to help with this process. The program can import and export both DICOM images and `EGSPHANT` files, which are required for DOSXYZnrc. The same tool can also be used to visualize dose distributions imported from `.3ddose` files. All of the contour plots in this chapter and in the previous chapter were generated this way. Figure 32 on the next page shows screenshots of CTgui.

Using CTgui, a 48 x 36 x 16 cm^3 region surrounding a lung tumor is chosen to provide the benchmark geometry. The coordinate system is the set so that the tumor is located at the treatment isocenter. A posterior-anterior field is chosen to provide the source. The field size is 2 x 2 cm^2 at the isocenter. (Since the tumor is located at a depth of 9.9 cm, and because the source-to-isocenter distance is 100 cm, the field has a source-to-surface distance of 90.1 cm, and the field size is 1.802 x 1.802 cm^2 at the posterior surface. The source energy distribution is given by the same 6MV spectrum (see figure 8 on page 39) that is used with the water phantom in the previous chapter.

8.2 Monte Carlo Reference Solution

The reference solution is computed using DOSXYZnrc. The physics settings are given by table 18 on page 89. Dose is tallied in 5 x 5 x 5 mm^3 voxels throughout the phantom. For this calculation, 10^{10} source histories were followed, consuming 5.5 hours on 8 compute nodes or a total of 64 CPU cores. The resulting solution has average relative uncertainty of 0.06%. The maximum relative uncertainty in voxels receiving at least 10% of the maximum dose is 0.91%.

8.3 Pre-computation

Response functions are pre-computed for 10 mm meshes with expansion up to order (929172). The energy structure is the same as for the slab phantom benchmark (see chapter 6). For each source basis function, 10 million source particles are simulated. Dose is tallied in 1 x 1 x 1 mm^3 voxels within each mesh. The pre-computation uses 9.8 hours on 12 compute

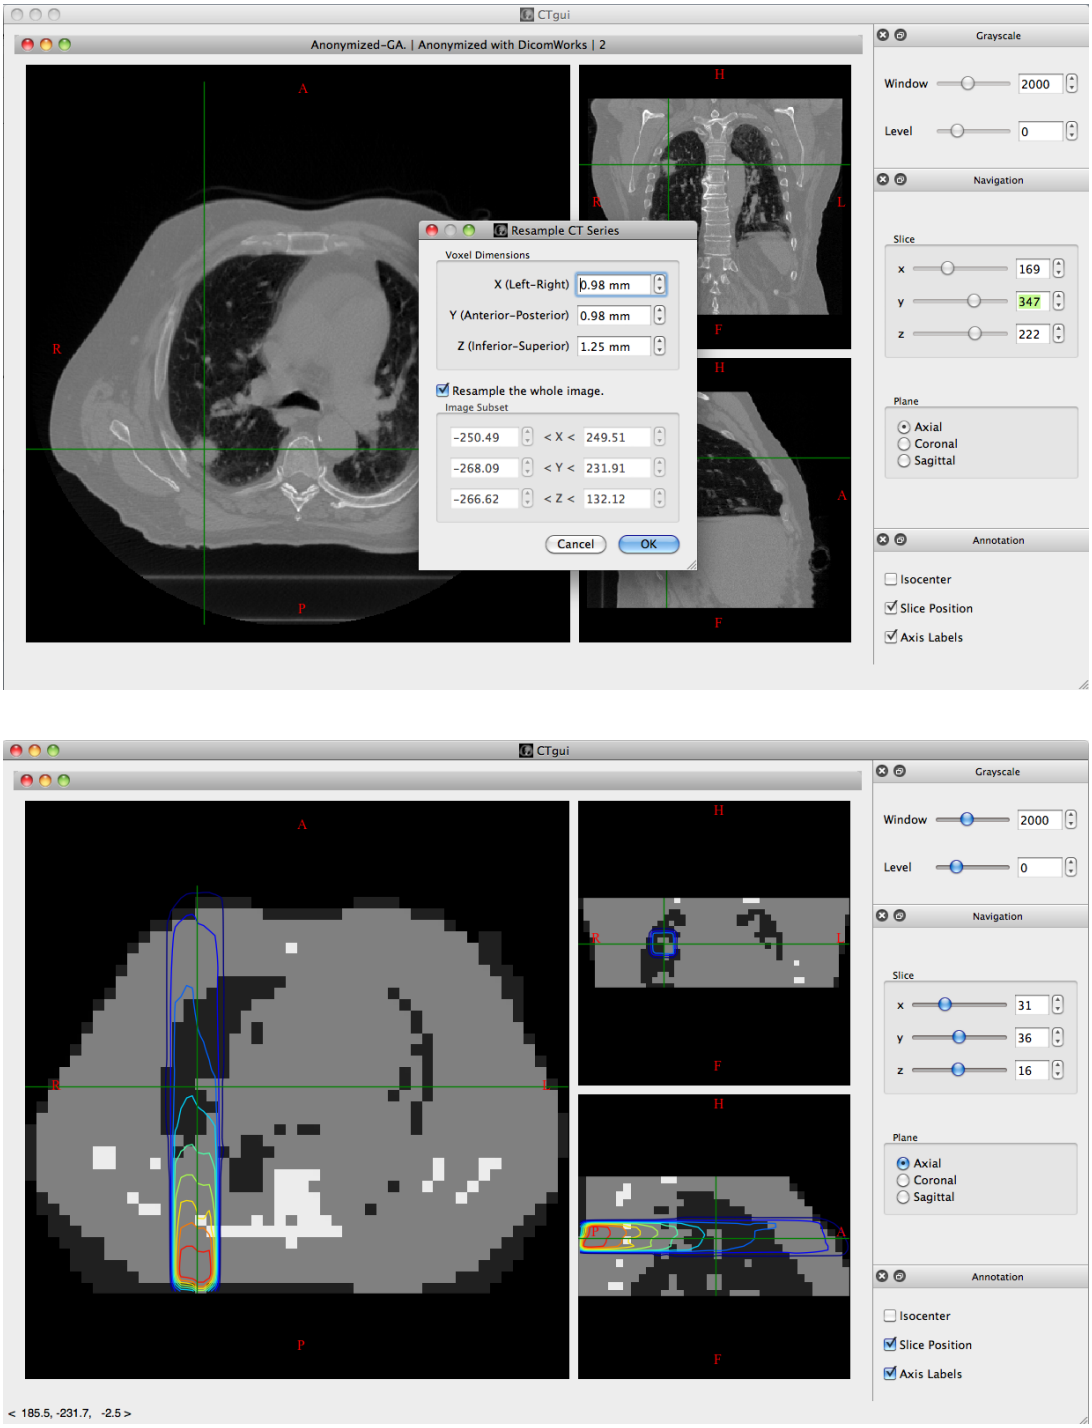

Figure 32: CTgui: a tool for interaction with volume data.

nodes (96 cores).

8.4 COMET-PE Solution

The COMET-PE solution is computed with an expansion or order (924162) using a convergence criterion of $\varepsilon = 10^{-3}$. COMET-PE outputs dose for 1 x 1 x 1 mm^3 voxels throughout the phantom. For comparison with Monte Carlo, these values are then averaged over 5 x 5 x 5 mm^3 voxels. The resulting solution has average relative uncertainty of 0.14%. The maximum relative uncertainty in voxels receiving at least 10% of the maximum dose is 0.69 %. Including uncertainty calculation, the solution uses 81 minutes on 8 compute nodes. Without uncertainty calculation, the time is 33 minutes.

8.5 Results

A comparison of the COMET-PE and Monte Carlo dose distributions is given in figures 33 to 35 on pages 82–84. As with the water phantom, dose is compared for voxels that received greater than 10% of the maximum dose. In this case there are 6007 comparison voxels. The 37 voxels that fail the 3% / 3 mm test are shown in figure 36 on page 85. As

Table 17: Agreement of voxels receiving greater than 10% of the maximum reference dose.

Criterion	# Passing voxels	% Passing voxels
1% / 1 mm	5 106	85.00
2% / 2 mm	5 931	98.73
3% / 3 mm	5 970	99.38
4% / 4 mm	5 987	99.67
5% / 5 mm	6 000	99.88

with the water phantom, all of these voxels are located either on the beam corner or near the beam's exit. In each of the 37 voxels, COMET-PE predicts a higher dose than Monte Carlo.

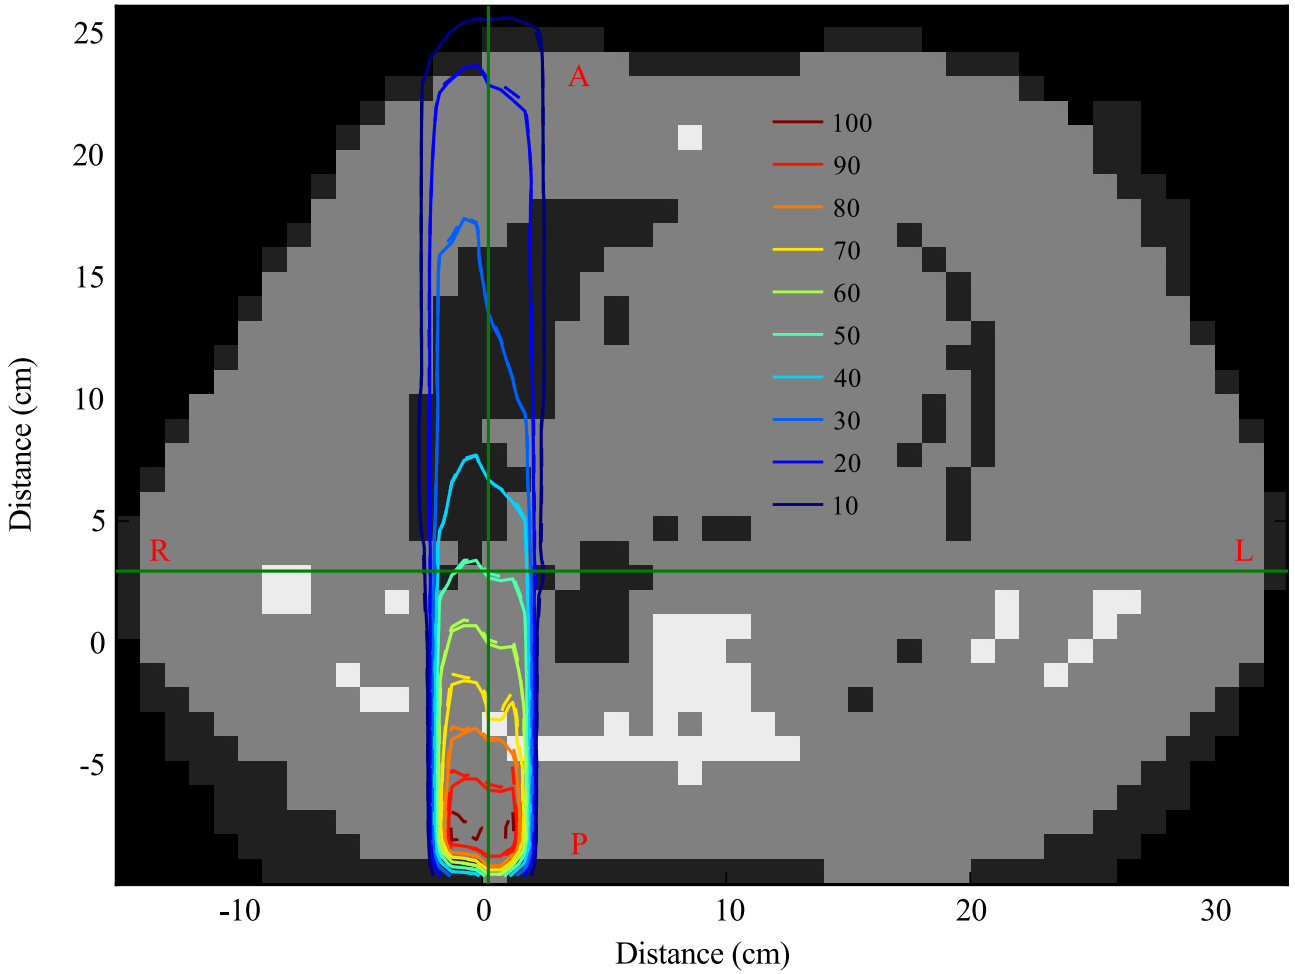

Figure 33: Lung phantom isodose contours in the axial plane. The contours give dose as a percentage of the maximum reference dose. Solid contours represent the Monte Carlo solution and dashed contours represent the COMET-PE solution. The letters A, P, R and L indicate anterior, posterior, left, and right directions. The horizontal and vertical green lines correspond to the coronal plane of figure 35 and the sagittal plane of figure 34.

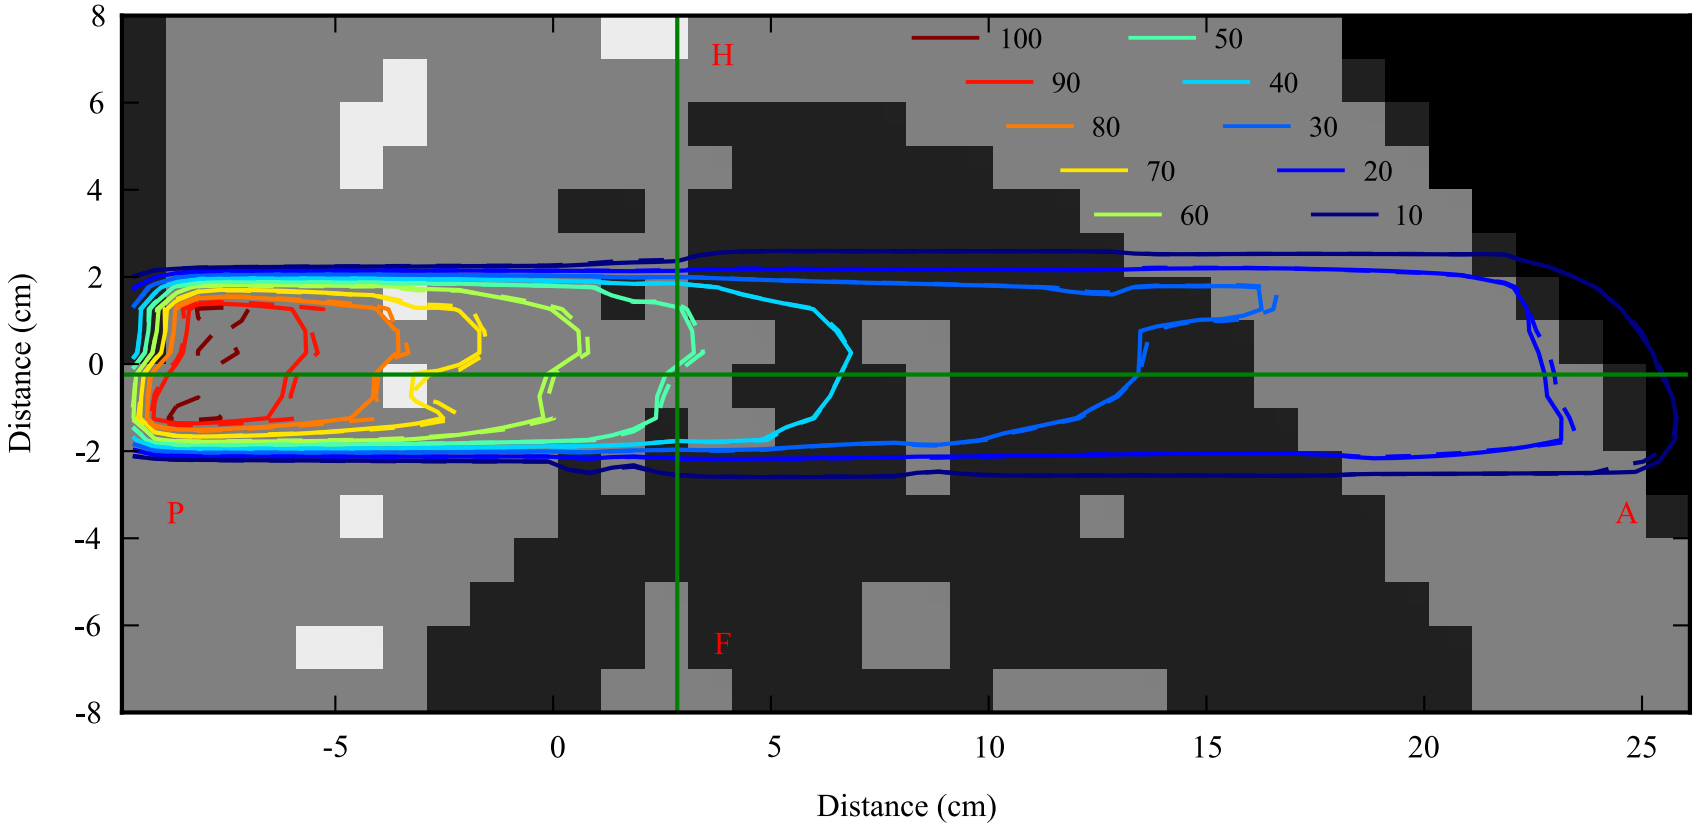

Figure 34: Lung phantom isodose contours in the sagittal plane. Solid contours represent the Monte Carlo solution, and dashed contours represent the COMET-PE solution. The letters A, P, H and F indicate anterior, posterior, head, and foot directions. The horizontal green and vertical green lines correspond to the axial plane of figure 33 and the coronal plane of figure 35.

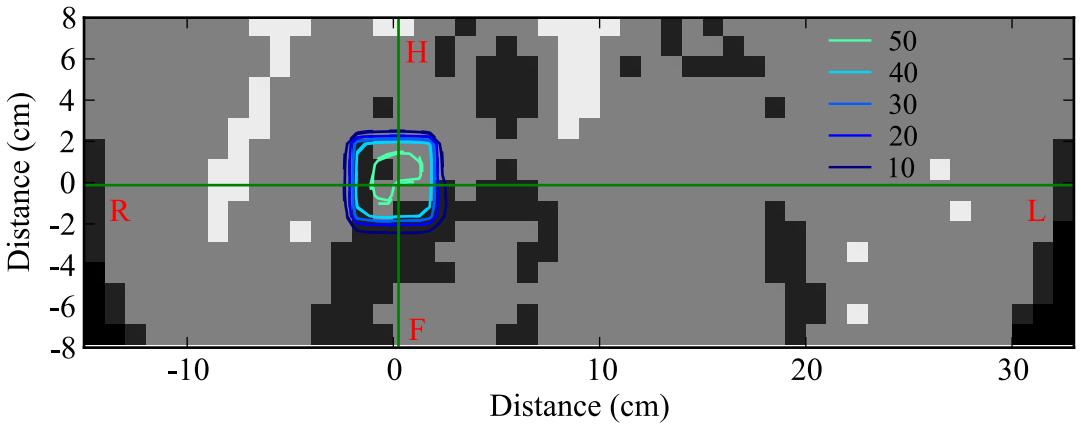

Figure 35: Lung phantom isodose contours in the coronal plane. Solid contours represent the Monte Carlo solution, and dashed contours represent the COMET-PE solution. The letters R, L, H and F indicate right, left, head, and foot directions. The horizontal and vertical green lines correspond to the axial plane of figure 33 and the sagittal plane of figure 34.

8.6 Discussion

COMET-PE is shown to be highly accurate and efficient for this CT-based problem. Over 98% of voxels with dose greater than 10% of maximum dose pass the 2% / 2 mm criterion. Furthermore, the number of passing voxels, especially on the beam corners, may be underestimated. This is due to limitations in the way that distance to agreement is evaluated. Most of the voxels that fail the 3 % / 3 mm test are located within 3 cm of the anterior surface, where the dose is lowest, and where Monte Carlo uncertainties are highest. A similar failure pattern was observed when using the raw Monte Carlo solution as a reference for the water phantom. In that case, the failures disappeared when symmetry was used to reduce the uncertainty of the Monte Carlo reference solution in post-processing. For this reason, it seams likely that many of the failures near the phantom's anterior surface are caused by the limited precision of the Monte Carlo solution.

Furthermore, performance is excellent. With uncertainty calculation enabled, COMET-PE is 4 times as fast as the Monte Carlo reference solution. Without uncertainty estimation,

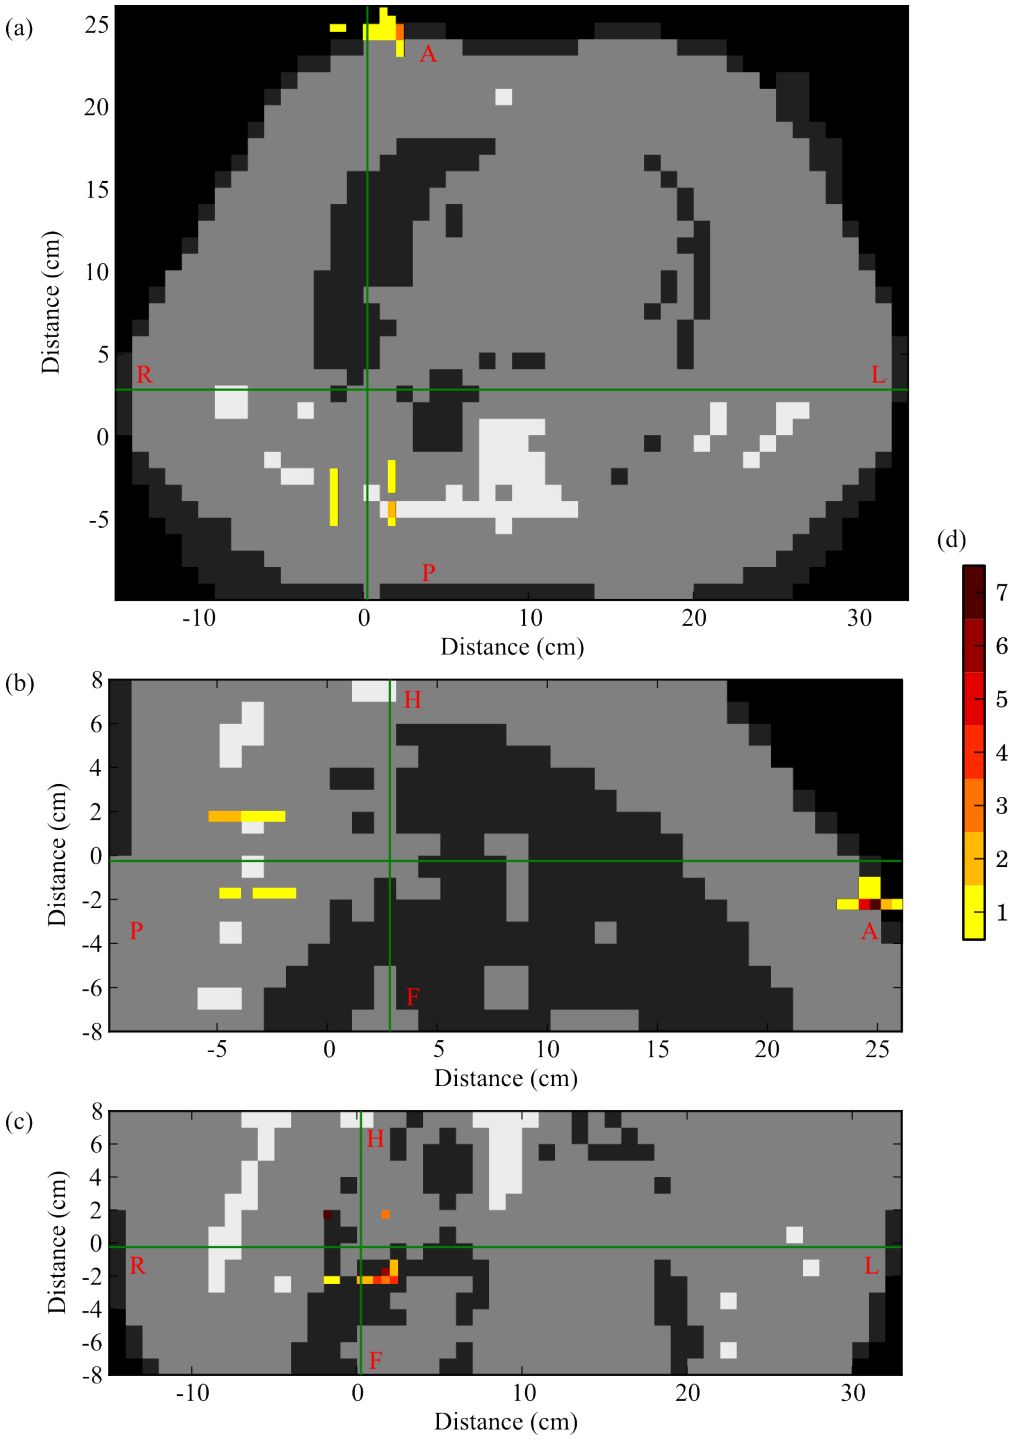

Figure 36: Voxels failing the 3% / 3 mm test. Projections are shown on the axial (a), sagittal (b), and coronal (c) planes. Each colored pixel represents the total number of failing voxels summed along the direction perpendicular to the image plane according to the scale in part (d).

the COMET-PE calculation is 10 times as fast as Monte Carlo. As mentioned in the previous chapter, enhancements to the uncertainty calculation will likely decrease the cost of COMET-PE uncertainty estimation in the near future.

CHAPTER IX

CONCLUSION AND FUTURE WORK

9.1 Summary

A new method for radiotherapy calculations that is comparable in accuracy to Monte Carlo has been developed. The new method is an extension of the COMET method for coupled photon-electron transport in 3-D. Because it takes advantage of stochastic pre-computation, the method is capable of achieving high precision without long calculation times. To adapt the method for radiotherapy, both a new set of angular basis functions and a new analytical treatment of highly anisotropic sources are used. The efficiency of the iterative solver is improved, and it can now take full advantage of both shared-memory and distributed-memory parallel architectures. The method's accuracy is demonstrated by comparison with the DOSXYZnrc Monte Carlo code. For a highly heterogeneous slab phantom, COMET-PE calculates dose with maximum error of less than 1.5% and with maximum relative error of less than 3% for both 5 mm and 10 mm mesh sizes. When a COMET-PE solution is compared with Monte Carlo for a homogeneous water phantom 97.31% of voxels receiving more than 10% of the maximum dose pass a 2%/2 mm distance to agreement test and 99.75% pass a 3%/3 mm test. When COMET-PE is compared with Monte Carlo for a CT-based dose calculation benchmark, 98.73% of the voxels receiving more than 10% of the maximum dose pass a 2%/2 mm test, and 99.38% pass a 3%/3 mm test. These results shows that COMET-PE is capable of high fidelity physical simulation in the presence of strong heterogeneity and in the absence of charged particle equilibrium as well as in complicated CT-based geometries. Timing comparisons show that COMET-PE is roughly 10-30 times faster than Monte Carlo for calculations using a 10 mm mesh size. Thus, COMET-PE is not only highly accurate, but fast.

9.2 The Path Forward

To be clinically useful, the method must be able to model a realistic linear accelerator and it must be able to represent the continuous variation in material properties associated with CT-based dose calculation. Future work should investigate response function interpolation as a means of simulating continuous density variation. Additionally, a more accurate source model that includes accelerator head scatter, collimator leakage, and electron contamination should be investigated.

There are many opportunities to further improve the performance of the iterative solver. Foremost among these is the Low Order Acceleration (LOA) method, which has already been applied to COMET calculations in nuclear reactor physics. Finally, the search for basis functions that provide a better representation of the fluence using fewer terms should continue.

APPENDIX A

EGSNRC TRANSPORT PARAMETERS

Table 18: EGSnrc transport parameters.

Parameter	Value
Photon Energy Cutoff	10 keV
Electron Energy Cutoff	700 keV
Photon Cross Sections	XCOM
Electron Step Algorithm	PRESTA-II
Boundary Crossing Algorithm	EXACT
Spin Effects	On
Bound Compton Scattering	norej
Photoelectron Angular Sampling	On
Rayleigh Scattering	On
Atomic Relaxations	On
Bremsstrahlung Cross Sections	NIST
Bremsstrahlung Angular Sampling	KM
Pair Production Angular Sampling	KM

APPENDIX B

QUANTITIES USED TO DESCRIBE ACCURACY AND PRECISION

Table 19: Quantities used to evaluate accuracy and precision. $\bar{x}$ denotes the reference solution; $\tilde{x}$, the approximate solution; y, a generic vector; and δ_y its estimated standard deviation.

Quantity	Definition
Average Relative Uncertainty [39]	$\text{RMS}\left(\left\{ \frac{\delta_{yi}}{y_i} \;\middle|\; y_i > \frac{\max(y)}{2} \right\}\right)$
Maximum Relative Uncertainty	$\max\left(\delta_y \middle/ y\right)$
Mean Relative Error	$\text{mean}\left\lvert(\tilde{x} - \bar{x})\,/\bar{x}\right\rvert$
RMS Relative Error	$\text{RMS}\left((\tilde{x} - \bar{x})/\bar{x}\right)$
Maximum Relative Error	$\max\left\lvert(\tilde{x} - \bar{x})\,/\bar{x}\right\rvert$
Mean Error	$\text{mean}\left\lvert\tilde{x} - \bar{x}\right\rvert / \max\left\lvert\bar{x}\right\rvert$
RMS Error	$\text{RMS}\,(\tilde{x} - \bar{x})/\max\left\lvert\bar{x}\right\rvert$
Maximum Error	$\max\left\lvert\tilde{x} - \bar{x}\right\rvert/\max\left\lvert\bar{x}\right\rvert$
Root Mean Square (RMS)	$\sqrt{\frac{1}{N}\sum_{i=1}^{N} {y_i}^2}$
Mean	$\frac{1}{N}\sum_{i=1}^{N} y_i$

Given a reference distribution $\bar{f}(\mathbf{r})$ and an approximate distribution $\tilde{f}(\mathbf{r})$, the distance-to-agreement at a point $\mathbf{r}$ is defined as

$$\text{DTA}(\mathbf{r}) = \min\left\{ h \mid \bar{f}(\mathbf{r}) = \tilde{f}(\mathbf{r} + h\hat{\Omega}) \quad \forall\, h \geq 0, \hat{\Omega} \in 4\pi \right\}.$$

In the case of values defined on a set of voxels, tri-linear interpolation is used to find

continuous distributions. In this work, distance-to-agreement is estimated by searching only along the three coordinate axes ($\pm\hat{\mathrm{x}}, \pm\hat{\mathrm{y}}, \pm\hat{\mathrm{z}}$) rather than the full space of possible directions (4π). The approximation is then given by

$$\mathrm{DTA}(\mathbf{r}) \approx \min\left\{h \mid \bar{f}(\mathbf{r}) = \tilde{f}(\mathbf{r} + h\hat{\Omega}) \quad \forall\, h \geq 0, \hat{\Omega} \in \{\pm\hat{\mathrm{x}}, \pm\hat{\mathrm{y}}, \pm\hat{\mathrm{z}}\}\right\}.$$

This approximation is an upper bound for the true distance-to-agreement as defined above.

REFERENCES

[1] Abramowitz, M. and Stegun, I. A., *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. New York: Dover Publications, 1965.

[2] Agostinelli, S., Allison, J., Amako, K., Apostolakis, J., Araujo, H., Arce, P., Asai, M., Axen, D., Banerjee, S., Barrand, G., Behner, F., Bellagamba, L., Boudreau, J., Broglia, L., Brunengo, A., Burkhardt, H., Chauvie, S., Chuma, J., Chytracek, R., Cooperman, G., Cosmo, G., Degtyarenko, P., Dell'Acqua, A., Depaola, G., Dietrich, D., Enami, R., Feliciello, A., Ferguson, C., Fesefeldt, H., Folger, G., Foppiano, F., Forti, A., Garelli, S., Giani, S., Giannitrapani, R., Gibin, D., G?mez Cadenas, J. J., Gonz?lez, I., Gracia Abril, G., Greeniaus, G., Greiner, W., Grichine, V., Grossheim, A., Guatelli, S., Gumplinger, P., Hamatsu, R., Hashimoto, K., Hasui, H., Heikkinen, A., Howard, A., Ivanchenko, V., Johnson, A., Jones, F. W., Kallenbach, J., Kanaya, N., Kawabata, M., Kawabata, Y., Kawaguti, M., Kelner, S., Kent, P., Kimura, A., Kodama, T., Kokoulin, R., Kossov, M., Kurashige, H., Lamanna, E., Lamp?n, T., Lara, V., Lefebure, V., Lei, F., Liendl, M., Lockman, W., Longo, F., Magni, S., Maire, M., Medernach, E., Minamimoto, K., Mora de Freitas, P., Morita, Y., Murakami, K., Nagamatu, M., Nartallo, R., Nieminen, P., Nishimura, T., Ohtsubo, K., Okamura, M., O'Neale, S., Oohata, Y., Paech, K., Perl, J., Pfeiffer, A., Pia, M. G., Ranjard, F., Rybin, A., Sadilov, S., Di Salvo, E., Santin, G., Sasaki, T., Savvas, N., Sawada, Y., Scherer, S., Sei, S., Sirotenko, V., Smith, D., Starkov, N., Stoecker, H., Sulkimo, J., Takahata, M., Tanaka, S., Tcherniaev, E., Safai Tehrani, E., Tropeano, M., Truscott, P., Uno, H., Urban, L., Urban, P., Verderi, M., Walkden, A., Wander, W., Weber, H., Wellisch, J. P., Wenaus, T., Williams, D. C., Wright, D., Yamada, T., Yoshida, H., and Zschiesche, D., "GEANT4–a simulation toolkit," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 506, pp. 250–303, July 2003.

[3] Ahnesjö, A., "Collapsed cone convolution of radiant energy for photon dose calculation in heterogeneous media," *Medical Physics*, vol. 16, no. 4, pp. 577–592, 1989.

[4] Ahnesjö, A. and Aspradakis, M. M., "Dose calculations for external photon beams in radiotherapy," *Physics in Medicine and Biology*, vol. 44, p. R99, 1999.

[5] Ahnesjo, A., Saxner, M., and Trepp, A., "A pencil beam model for photon dose calculation," *Medical Physics*, vol. 19, no. 2, pp. 263–273, 1992.

[6] Andreo, P., "Monte carlo techniques in medical radiation physics," *Physics in medicine and biology*, vol. 36, p. 861, 1991.

[7] Barrett, R., Berry, M., Chan, T. F., Demmel, J., Donato, J., Dongarra, J., Eijkhout, V., Pozo, R., Romine, C., and Van der Vorst, H., *Templates for the Solution of Linear*

Systems: Building Blocks for Iterative Methods. Philadelphia, PA: SIAM, 2nd ed., 2006.

[8] Berntsen, J., Espelid, T. O., and Genz, A., "An adaptive algorithm for the approximate calculation of multiple integrals," *ACM Trans. Math. Softw.*, vol. 17, p. 437–451, Dec. 1991.

[9] Blackburn, M. S., "Numerical benchmarking of a coarse-mesh transport (COMET) method for medical physics applications," 2009.

[10] Boman, E., Tervo, J., and Vauhkonen, M., "Modelling the transport of ionizing radiation using the finite element method," *Physics in Medicine and Biology*, vol. 50, pp. 265–280, Jan. 2005.

[11] Boyer, A. and Mok, E., "A photon dose distribution model employing convolution calculations," *Medical physics*, vol. 12, p. 169, 1985.

[12] Bush, K., Gagne, I. M., Zavgorodni, S., Ansbacher, W., and Beckham, W., "Dosimetric validation of acuros® XB with monte carlo methods for photon dose calculations," *Medical Physics*, vol. 38, p. 2208, 2011.

[13] Cheng, A., Sr, W. B. H., Gerber, R. L., Wong, J. W., and Purdy, J. A., "Systematic verification of a three-dimensional electron beam dose calculation algorithm," *Medical Physics*, vol. 23, no. 5, pp. 685–693, 1996.

[14] Chetty, I. J., Curran, B., Cygler, J. E., DeMarco, J. J., Ezzell, G., Faddegon, B. A., Kawrakow, I., Keall, P. J., Liu, H., Ma, C. M., Rogers, D. W. O., Seuntjens, J., Sheikh-Bagheri, D., and Siebers, J. V., "Report of the AAPM task group no. 105: Issues associated with clinical implementation of monte carlo-based photon and electron external beam treatment planning," tech. rep., 2007.

[15] Daskalov, G. M., Baker, R. S., Rogers, D. W. O., and Williamson, J. F., "Multigroup discrete ordinates modeling of i 6702 seed dose distributions using a broad energy-group cross section representation," *Medical Physics*, vol. 29, p. 113, 2002.

[16] Daskalov, G. M., Baker, R. S., Rogers, D. W. O., and Williamson, J. F., "Dosimetric modeling of the microselectron high-dose rate [sup 192]Ir source by the multigroup discrete ordinates method," *Medical Physics*, vol. 27, no. 10, pp. 2307–2319, 2000.

[17] Duderstadt, J. J. and Martin, W. R., *Transport Theory*. John Wiley & Sons Inc, June 1979.

[18] Gautron, P., Krivanek, J., Pattanaik, S., and Bouatouch, K., "A novel hemispherical basis for accurate and efficient rendering," in *Rendering Techniques 2004, Eurographics Symposium on Rendering*, p. 321–330, Citeseer, 2004.

[19] Gifford, K. A., Horton, J. L., Wareing, T. A., Failla, G., and Mourtada, F., "Comparison of a finite-element multigroup discrete-ordinates code with monte carlo for radiotherapy calculations," *Physics in Medicine and Biology*, vol. 51, pp. 2253–2265, May 2006.

[20] Hartmann, S. C., Walling, R. S., Daly, T. P., Faddegon, B., Albright, N., Bergstrom, P., Bielajew, A. F., Chuang, C., Garrett, D., and House, R. K., "Description and dosimetric verification of the PEREGRINE monte carlo dose calculation system for photon beams incident on a water phantom.," *Medical physics*, vol. 28, no. 7, p. 1322, 2001.

[21] Hendricks, J. S., McKinney, G. W., Waters, L. S., Roberts, T. L., Egdorf, H. W., Finch, J. P., Trellue, H. R., Pitcher, E. J., Mayo, D. R., and Swinhoe, M. T., "MCNPX extensions: version 2.5. 0," *Los Alamos National Laboratory*, 2005.

[22] Jan, S., Benoit, D., Becheva, E., Carlier, T., Cassol, F., Descourt, P., Frisson, T., Grevillot, L., Guigues, L., and Maigne, L., "GATE v6: a major enhancement of the GATE simulation platform enabling modelling of CT and radiotherapy," *Physics in medicine and biology*, vol. 56, no. 4, p. 881, 2011.

[23] Kawrakow, I., "VMC++, electron and photon monte carlo calculations optimized for radiation treatment planning," in *Advanced Monte Carlo for Radiation Physics, Particle Transport Simulation and Applications: Proc. Monte Carlo 2000 Meeting*, pp. 229–236, Berlin: Springer, 2001.

[24] Kawrakow, I., Mainegra-Hing, E., Rogers, D. W. O., Tessier, F., and Walters, B. R. B., "The EGSnrc code system," Tech. Rep. PIRS-701, NRCC, May 2011.

[25] Larsen, E. W., Miften, M. M., Fraass, B. A., and Bruinvis, I. A. D., "Electron dose calculations using the method of moments," *Medical Physics*, vol. 24, no. 1, pp. 111–125, 1997.

[26] Lewis, E. E. and Miller, W. F., *Computational Methods of Neutron Transport*. La Grange Park, IL, USA: American Nuclear Society, Inc., 1993.

[27] Lorence Jr, L. J., Morel, J. E., and Valdez, G. D., "Physics guide to CEPXS: a multigroup coupled electron-photon cross-section generating code," *SAND89-1685, Sandia National Laboratory*, 1989.

[28] Low, D. A., Harms, W. B., Mutic, S., and Purdy, J. A., "A technique for the quantitative evaluation of dose distributions," *Medical physics*, vol. 25, p. 656, 1998.

[29] Mackie, T. R., Bielajew, A. F., Rogers, D. W. O., and Battista, J. J., "Generation of photon energy deposition kernels using the EGS monte carlo code," *Physics in medicine and biology*, vol. 33, p. 1, 1988.

[30] Mackie, T. R., Scrimger, J. W., and Battista, J. J., "A convolution method of calculating dose for 15-MV x rays," *Medical physics*, vol. 12, p. 188, 1985.

[31] Makhotkin, O. A., "Analysis of radiative transfer between surfaces by hemispherical harmonics," *Journal of Quantitative Spectroscopy and Radiative Transfer*, vol. 56, no. 6, pp. 869–879, 1996.

[32] Metropolis, N., "The beginning of the monte carlo method," *Los Alamos Science*, vol. 15, no. 584, p. 125–130, 1987.

[33] Naqvi, S. A., Earl, M. A., and Shepard, D. M., "Convolution/superposition using the monte carlo method," *Physics in Medicine and Biology*, vol. 48, pp. 2101–2121, July 2003.

[34] Nelson, W. R., Hirayama, H., and Rogers, D. W., "EGS4 code system," tech. rep., Stanford Linear Accelerator Center, Menlo Park, CA (USA), 1985.

[35] O'Connor, J. E., "The variation of scattered x-rays with density in an irradiated body," *Physics in Medicine and Biology*, vol. 1, pp. 352–369, Apr. 1957.

[36] Papanikolaou, N., Battista, J. J., Boyer, A. L., Kappas, C., Klein, E., Mackie, T. R., Sharpe, M., and Dyk, J. V., "Tissue inhomogeneity corrections for megavoltage photon beams," Tech. Rep. 85, 2004.

[37] Raeside, D. E., "Monte carlo principles and applications," *Physics in Medicine and Biology*, vol. 21, pp. 181–197, Mar. 1976.

[38] Rogers, D. W. O., "Fifty years of monte carlo simulations for medical physics," *Physics in medicine and biology*, vol. 51, p. R287, 2006.

[39] Rogers, D. W. O. and Mohan, R., "Questions for comparison of clinical monte carlo codes," in *The Use of Computers in Radiation Therapy: XIIIth International Conference* (Schlegel, W. and Bortfeld, T., eds.), (Heidelberg, Germany), p. 120–122, Springer, May 2000.

[40] Salvat, F., Fernández-Varea, J. M., and Sempau, J., "PENELOPE-2006: a code system for monte carlo simulation of electron and photon transport," in *Workshop Proceedings Barcelona, Spain*, vol. 4, Citeseer, 2006.

[41] Sempau, J., Wilderman, S. J., and Bielajew, A. F., "DPM, a fast, accurate monte carlo code optimized for photon and electron radiotherapy treatment planning dose calculations," *Physics in Medicine and Biology*, vol. 45, p. 2263, 2000.

[42] Siddon, R. L., "Fast calculation of the exact radiological path for a three-dimensional CT array," *Medical Physics*, vol. 12, p. 252, 1985.

[43] Smelov, V. V., "On completeness of hemispherical harmonics system," *Sibirskii Zhurnal Vychislitel'noi Matematiki*, vol. 1, no. 4, pp. 391–395, 1998.

[44] Spanier, J. and Gelbard, E. M., *Monte Carlo principles and neutron transport problems*. Mineola, N.Y: Dover Publications, dover ed ed., 2008.

[45] Szegö, G., *Orthogonal polynomials*. American Mathematical Society, 3rd ed., 1967.

[46] Tervo, J., Vauhkonen, M., and Boman, E., "Optimal control model for radiation therapy inverse planning applying the boltzmann transport equation," *Linear Algebra and its Applications*, vol. 428, pp. 1230–1249, Mar. 2008.

[47] Thomas E. Booth, Forrest B. Brown, Jeffrey S. Bull, R. Arthur Forster, John T. Goorley, H. Grady Hughes, Russell D. Mosteller, Richard E. Prael, Avneet Sood, Jeremy E. Sweezy, Anthony Zukaitis, Marsha Boggs, and Roger Martz, "MCNP—A general monte carlo n-particle transport code, version 5," Tech. Rep. LA-UR-03-1987, Los Alamos National Laboratory, 2003.

[48] Van Esch, A., Tillikainen, L., Pyykkonen, J., Tenhunen, M., Helminen, H., Siljamäki, S., Alakuijala, J., Paiusco, M., Iori, M., and Huyskens, D. P., "Testing of the analytical anisotropic algorithm for photon dose calculation," *Medical physics*, vol. 33, p. 4130, 2006.

[49] Vassiliev, O. N., Wareing, T. A., Davis, I. M., McGhee, J., Barnett, D., Horton, J. L., Gifford, K., Failla, G., Titt, U., and Mourtada, F., "Feasibility of a multigroup deterministic solution method for three-dimensional radiotherapy dose calculations," *International Journal of Radiation Oncology* Biology* Physics*, vol. 72, no. 1, pp. 220–227, 2008.

[50] Vassiliev, O. N., Wareing, T. A., McGhee, J., Failla, G., Salehpour, M. R., and Mourtada, F., "Validation of a new grid-based boltzmann equation solver for dose calculation in radiotherapy with photon beams," *Physics in Medicine and Biology*, vol. 55, p. 581, 2010.

[51] Walters, B., Kawrakow, I., and Rogers, D. W. O., "DOSXYZnrc users manual," Tech. Rep. PIRS-794revB, NRCC, May 2011.

[52] Williams, M. L., Ilas, D., Sajo, E., Jones, D. B., and Watkins, K. E., "Deterministic photon transport calculations in general geometry for external beam radiation therapy," *Medical Physics*, vol. 30, p. 3183, 2003.

[53] Yuan, J., Jette, D., and Chen, W., "Deterministic photon kerma distribution based on the boltzmann equation for external beam radiation therapy," *Medical physics*, vol. 35, p. 4079, 2008.

[54] Zhang, D. and Rahnema, F., "A heterogeneous coarse mesh method for coupled photon electron transport problems," *Transport Theory and Statistical Physics*, vol. 40, no. 3, pp. 127–152, 2011.

[55] Zhang, D. and Rahnema, F., "An efficient hybrid Stochastic/Deterministic coarse mesh neutron transport method," *Annals of Nuclear Energy*, vol. 41, pp. 1–11, Mar. 2012.