AN ABSTRACT OF THE DISSERTATION OF

Elanchezhian Somasundaram for the degree of Doctor of Philosophy in Nuclear Engineering presented on September 15, 2015.

Title: Application of Advanced Hybrid Transport Methods for Nonproliferation Test Problems

Abstract approved: _____

Todd S. Palmer

The implementation of advanced hybrid (Monte Carlo/Deterministic) transport methods for realistic test problems has been a challenge due to the overhead efforts associated with interfacing a solution generated by a deterministic solver with a Monte Carlo based radiation transport code. In this work, with the help of Transpire, Inc., a hybrid methods workbench that consists of a multigroup Monte Carlo code has been developed within the framework of Attila^{TM}, a commercial deterministic tool for radiation transport problems. This workbench is used to test two advanced variance reduction techniques, the LIFT method and the VVR method on source-detector type test problems that are of interest for nonproliferation research. Further, the more common weight windows and source biasing based variance reduction techniques are also implemented based on the CADIS methodology. The performance of the hy-

brid methods has been analyzed on two photon test problems and one neutron test problem, both individually and when applied together. The behavior of the methods for different configurations of the problems have been studied and the limitations of the methods when applied to complex problem has been analyzed. From these observations, the possible research directions to improve the test bench as well as the methods itself are presented.

Application of Advanced Hybrid Transport Methods for Nonproliferation Test Problems

by

Elanchezhian Somasundaram

A DISSERTATION

submitted to

Oregon State University

in partial fulfillment of
the requirements for the
degree of

Doctor of Philosophy

Presented September 15, 2015
Commencement June 2016

Doctor of Philosophy dissertation of Elanchezhian Somasundaram presented on September 15, 2015.

APPROVED:

Major Professor, representing Nuclear Engineering

Head of the School of Nuclear Science and Engineering

Dean of the Graduate School

I understand that my dissertation will become part of the permanent collection of Oregon State University libraries. My signature below authorizes release of my dissertation to any reader upon request.

Elanchezhian Somasundaram, Author

ACKNOWLEDGEMENTS

I like to express my deep appreciation and thanks to my advisor, Professor Dr. Todd Palmer, you have been an wonderful mentor for me. I would like to thank you for encouraging my research, for providing me with all the assistance needed to complete my research successfully and also for championing me through all the highs and lows of my graduate life. Your advice on both research as well as on my career have been priceless.

I would also like to thank professor Dr. Andrew C. Klein, professor Dr. Brian Woods, professor Dr. Leah Minc and my minor professor, Dr. Lisa J. Madsen for serving as my committee members. I would like to thank you for letting my defense be an enjoyable moment, and for your brilliant comments and suggestions.

I also want to thank Dr. Ed Larsen, professor at the University of Michigan, Dr. Jeffery Densmore and Dr. Kendra P. Keady, for your excellent suggestions and ideas to enhance my research. I would also like to express my thanks and gratitude to Transpire, Inc. for providing the software tools required to conduct my research, especially to Ian Davis and Dr. Todd Wareing.

I am also indebted to the School of Nuclear Science and Engineering, all the professors and colleagues for their support and warmth. I am also grateful to Oregon State University and the Office of Nonproliferation Research, Department of Energy, for providing the funding to support my graduate studies and my PhD research.

A special thanks to my family, friends and relatives for all of the sacrifices that you have made on my behalf. Your prayers for me is what sustained me thus far.

TABLE OF CONTENTS

TABLE OF CONTENTS (Continued)

TABLE OF CONTENTS (Continued)

LIST OF FIGURES

LIST OF TABLES

Chapter 1: Introduction

In this dissertation work, the implementation and performance analysis of advanced automated "hybrid" (Deterministic/Monte Carlo) variance reduction methods for realistic radiation shielding problems that are of interest in non-proliferation applications has been studied. This chapter gives a brief introduction to the type of radiation transport problems that are of interest in this work. A discussion of the major solution techniques for such problems and the motivation for selecting the specific techniques (hybrid methods) that are the research focus in this work are also presented.

1.1 Radiation transport problems for non-proliferation applications

The transport of nuclear particles in various media has been a major area of research in nuclear reactor engineering, radiation shielding, health physics, medical physics, remote sensing etc. There are numerous physical phenomena that influence the transport of a nuclear particle and are modeled to various degrees of accuracy depending on the application. Two major classes of problems that are of interest in nuclear science are eigenvalue problems and Source-Detector problems. The eigenvalue problems find application primarily in reactor design and engineering while the source-detector (including source-region) problems find application in radiation shielding, dosimetry,

non-proliferation monitoring and radiation therapy. In this work, our focus is on source-detector problems for non-proliferation related applications, but the methods developed can be easily extended to other similar applications.

Nuclear non-proliferation technologies often utilize radiation detection and spectral analysis to identify and locate sources of radiation. In many instances, it is not feasible to experimentally evaluate the performance of these technologies in the field prior to their deployment because of the safety issues and/or financial costs involved. Computer simulations can be used to develop the proposed technologies, and to understand the complexity of various scenarios of interest in non-proliferation science, provided the simulations generate high-fidelity information in "reasonable" periods of time.

Radiation transport problems are computationally challenging in that the quantities to be determined (particle density functions or *fluxes*) can be extremely complicated functions of seven independent variables: $\boldsymbol{x}$ - three position, $\boldsymbol{\Omega}$ - two direction, E - one energy, and t - one time.

Two example problems of interest to NA-22 (Office of Non-proliferation Research, Department of Energy) are discussed here to illustrate the requirements for radiation transport software packages. The first, illustrated in Figure 1.1, is a radiation portal monitor with four separate detectors exposed to radiation from (i) environmental sources (soil, concrete barriers, etc.), (ii) transported naturally-occurring radioactive material (NORM), (iii) medical sources, and/or (iv) true special nuclear material

sources. This scenario involves volumetric sources, potentially point-like sources (depending on the size of the problem), and large thin planar detectors. The substantial attenuation provided by the walls of the vehicle and its cargo likely qualifies this problem as "deep-penetration". Possible objectives in this analysis include design and evaluation of candidate detection systems and the calculation of detector response for the development of alarming algorithms. This is an optimization problem, which will require many radiation transport simulations to achieve a "best" result. Each simulation must be accurate and computed very quickly if the optimization problem is to be solved in a reasonable time period.

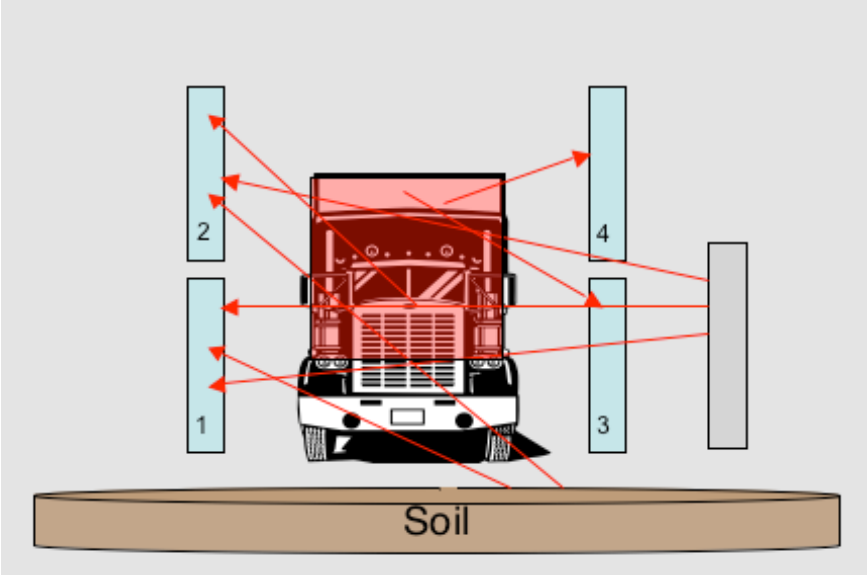

Figure 1.1: Example radiation transport calculation - portal monitoring

A second example problem is the urban search scenario, shown in Figure 1.2. Here the goals are to (i) locate sources of radiation in a relatively large urban area, (ii) discriminate between NORM (trees, concrete), medical sources (person crossing street),

and threat sources (van), and (iii) provide this detection information to local authorities for decision making and response. Simulations are used to understand minimum detectable quantities of different materials, to evaluate candidate detection systems, and to guide the improvement of decision-making processes. The physical size of this problem, relative to the physical size of the detectors and threat sources, and the presence of large thick attenuators (buildings) makes it extremely challenging to simulate, especially from the perspective of spatial and angular resolution.

Figure 1.2: Example radiation transport calculation: urban source search scenario

Traditionally, the transport problems of the kinds described above are simulated using deterministic methods or Monte-Carlo methods. There is also a newer breed of methods known as hybrid methods which combine both deterministic and Monte-Carlo techniques. A brief description of the three solution techniques will be presented in the following sections.

1.2 Deterministic methods

Deterministic techniques, as the name suggests solve the transport problem numerically by employing various discretization schemes and algorithms. The most accurate representation of neutral radiation transport is given by the Boltzmann Transport Equation (BTE) as below,

$$\left[\frac{1}{\nu(E)}\frac{d}{dt} + \Omega \cdot \nabla + \sigma_t(r,E)\right]\psi\left(r,E,\Omega,t\right) = \int_{4\pi} d\Omega' \int_0^\infty dE' \sigma_s(r,E' \rightarrow E, \Omega' \cdot \Omega)\psi\left(r,E',\Omega',t\right) + Q\left(r,E,\Omega,t\right), \tag{1.1}$$

$r-$ position (cm),

$\Omega-$ direction (sr),

$E-$ energy of the particle (MeV),

$t-$ time (sec)

$\nu-$ neutron speed at energy E (cm/sec),

$\psi-$ angular flux $(particles/cm^2/MeV/sr/sec)$,

σ_t- total macroscopic cross section about energy E and position r (cm^{-1}/MeV),

σ_s- scattering cross section for energy E' to E and direction Ω' to Ω $(cm^{-1}/MeV/sr)$,

$Q-$ source $(particles/cm^3/MeV/sr/sec)$.

In the above equation, the terms in the left describe the particle loss terms, while the terms in the right describe the production terms. The total cross section (σ_t) determines the total interaction probability of the particles with the materials they are traveling in. The scattering cross section (σ_s) determines the probability of the particles undergoing a scattering interaction with the material. The BTE in Equation (1.1) is known as the integro-differential form of the transport equation. It describes the forward transport of particles from the source region. There are various other forms that are derived from this general transport equation such as the integral form, parity form etc.

Most of the problems that are of interest in non-proliferation research are fixed source problems as described in section 1.1. Hence the time dependent term in the BTE is usually neglected and the steady state problem is solved. The steady state problem itself has six different variables, two in the angular domain, three in the spatial domain and one in the energy domain. The energy domain is usually discretized using the multigroup scheme in which the energy domain is subdivided into a number of energy groups with upper and lower boundaries. The angular domain is usually discretized using what is known as the discrete ordinates approach (where the angular domain is divided into a finite number of directions with weights to represent the entire domain). There are also schemes that use a finite element approach to discretize the angular domain. The scattering kernel (The first term in the right hand side of the BTE) which models the scattering of particles from one direction to another is represented using a Legendre polynomial expansion. Finally, the spatial co-ordinates

are discretized using a wide variety of finite element and finite difference schemes [1].

Deterministic methods produce detailed, system-wide solutions and are computationally efficient. However, deterministic methods contain uncertainties associated with the discretization of the independent variables (space, energy and angle) of the transport equation and can admit solutions that exhibit non-physical features (e.g., ray effects and negative fluxes). This is especially the case in shielding applications. Hence, a significant degree of insight and expertise is required to mitigate these undesirable characteristics and ultimately produce usable results.

1.3 Monte-Carlo methods

Monte Carlo methods are a family of stochastic simulation techniques that get their name from the famous casinos in Monaco as coined by Von Neumann [2]. Since the transport and interaction of radiation particles in a medium is a random process with fixed probabilities, Monte-Carlo techniques lend themselves naturally for simulating particle transport. They are more accurate than deterministic methods, especially for deep penetration problems with several orders of magnitude attenuation and are devoid of any truncation error (ideally) that are characteristic of all deterministic methods. However, as it is true with any statistical method, the Monte Carlo methods suffer from statistical errors and large simulation times are required to obtain a reasonable estimate of the target values (which is usually detector flux or reaction rate for source-detector problems). This has led to the development of biased

Monte-Carlo techniques which we will describe after a brief introduction to unbiased (analog) Monte-Carlo simulations as applied to transport problems.

1.3.1 Analog Monte Carlo simulations

To model the particle transport using Monte Carlo methods, all the processes the particle undergoes during its life, from being born (or emitted) from the source region to the various interactions it undergoes in the medium to finally being absorbed or escape the problem boundaries are represented using probability density functions. The probabilistic functions are determined based on the material properties (material type, densities), the nature of the source particle (neutron or photon) and the particle's velocity (or energy). Once the probability density functions (PDFs) are defined, the simulation of the particle transport can proceed by using pseudo-random numbers to sample parameters that determine the flight of the particle and its interaction with the medium. The following section will introduce the sampling methods for a one-dimensional slab problem with isotropic scattering.

1.3.1.1 Sampling methods for particle transport

Let us consider a slab of length L in the x-direction. The particle simulation begins by first sampling the source parameters such as the initial location, energy (energy

group in multigroup problems) and direction. Let us consider a uniform unit source distribution for a multigroup problem given by,

$$\begin{aligned} Q_g(x,\mu) &= 1 \text{ for } x \in (0, x_s), \mu \in (-1, 1) \text{ and } g \in 1 \text{ to } G. \\ &= 0, \text{ elsewhere} \end{aligned} \tag{1.2}$$

where G is the total number of energy groups, x_s be some point in the slab of length L and μ is the direction cosine of the angle θ which varies from 0 to π. The density function for sampling a source location between 0 and x_s is given by,

$$p(x) = \frac{Q_g(x)}{\int_0^L Q_g(x)dx}, \tag{1.3}$$

Since we have a uniform source distribution, the above equation can be written as,

$$p(x) = \frac{1}{\int_0^{x_s} dx} \tag{1.4}$$

To sample an initial location x_i, we convert the PDF into a CDF (Cumulative Distribution Function $P(x_i)$) as shown below,

$$P(x_i) = \frac{\int_0^{x_i} dx}{\int_0^{x_s} dx} = \frac{x_i}{x_s}. \tag{1.5}$$

To find x_i, a random number ξ is sampled between 0 and 1 (the range of values a CDF can take) such that,

$$\xi = \frac{x_i}{x_s}, \tag{1.6}$$

and this equation is solved for x_i to yield,

$$x_i = \xi x_s. \tag{1.7}$$

This process of determining x_i directly by sampling an uniform random number between 0 and 1 is known as sampling through inversion of the CDF. To sample an initial energy group g_i for the particle the PDF is given by,

$$p(g) = \frac{Q_g}{\sum_{g=1}^{G} Q_g} = \frac{1}{\sum_{g=1}^{G} Q_g} \tag{1.8}$$

and the corresponding equation for sampling an energy group after inverting the CDF is

$$g_i = \xi G. \tag{1.9}$$

Next step would be to sample an initial direction μ_i from an isotropic distribution given by,

$$p(\mu) = \frac{Q(\mu)}{\int_{-1}^{1} d\mu Q(\mu)} = \frac{1}{\int_{-1}^{1} d\mu} \tag{1.10}$$

and the equation for sampling a new direction is obtained after inverting the CDF to yield,

$$\mu_i = 2\xi - 1 \tag{1.11}$$

where ξ is a random number.

Once the source particle's initial location, energy and direction are sampled, the

sampling of the next distance to collision and the type of collision follow.

$$p(s) = C\frac{\psi(s)}{\psi(x_i)} \tag{1.12}$$

where C is the normalization constant. $\psi(s)$ is the flux determined by the equation that describes transport through a constant medium along the trajectory of a particle given by the reduced form of the transport equation,

$$\frac{d\psi(s)}{ds} + \sigma_t\psi(s) = 0 \tag{1.13}$$

with initial condition,

$$\psi(0) = \psi(x_i) \tag{1.14}$$

The solution to this equation is,

$$\psi(s) = \psi(x_i)e^{-\sigma_t s} \tag{1.15}$$

and the corresponding PDF is given by,

$$p(s) = \sigma_t e^{-\sigma_t s} \tag{1.16}$$

Calculating the CDF and inverting would yield the following expression for the new distance to collision s,

$$s = \frac{-log(\xi)}{\sigma_t} \tag{1.17}$$

Now, the particle should be moved to its new location x_{new} as calculated below

$$x_{new} = x_i + s\mu_i. \tag{1.18}$$

Once the particle reaches its new location, the type of collision it undergoes should also be sampled based on the probability of occurrence of a particular interaction (in other words the cross section of the reaction).

In shielding applications, when a particle collides with an atom (or molecule) of the medium, it either scatters or is captured according to some probability. The probability that a particle collision results in a scattering event is given by

$$p_{scat} = \frac{\sigma_s}{\sigma_t} \tag{1.19}$$

If $p_{scat} < \xi$ (the sampled random number), the collision results in a scattering event; otherwise, the particle is captured. If the particle is absorbed, it is terminated and a new source particle will be sampled, whereas if the particle is scattered, a new direction and a new energy group based on the scattering distribution is sampled. For isotropic scattering, the new direction is sampled using Equation (1.18). The new energy group g_{new} is then sampled using the PDF given by

$$p(g_{new}) = \frac{\sigma_{s,g,g_{new}}}{\sum_{g_{new}=1}^{G} \sigma_{s,g,g_{new}}}. \tag{1.20}$$

A new distance to collision is then sampled and the transport process continues until

the particle is absorbed or leaks out of the problem boundaries. The sampling distribution we looked at are all directly invertible, but in problems where continuous energy representation is used or in highly anisotropic scattering problems, the distribution may not easily be invertible and in such cases other computationally heavy sampling mechanisms such as rejection sampling are employed. The process continues until all particle histories are simulated.

1.3.1.2 Estimating target entities

During the particle transport, a score (also known as contribution) of the distance traveled by each particle in each region of interest is tallied. The tallies are binned in energy group and spatial region and in some rare applications in directions. Here we describe the most common tally that is employed in radiation transport codes, the path-length estimator. A path-length estimator keeps track of the number of times a particle enters the detector region and the corresponding distance it traverses within the detector volume. The scalar flux ϕ (which has units particles/cm^2/sec) in the tally is then estimated using the formula,

$$\phi = \frac{1}{V} \sum_{i \in T} l_i, \tag{1.21}$$

where T is the set of all the particle's trajectories within the desired volume and l_i is the length of the i-th trajectory and V is the volume of the spatial region that is

tallied. The average flux ϕ_{est} is then estimated using the law of large numbers after simulating a relatively large number of the particles N and the corresponding scores ($l_{i,n}$) are recorded,

$$\phi_{est} = \frac{1}{N}\sum_{n=1}^{N}\frac{1}{V}\sum_{n,i\in T} l_{i,n}. \tag{1.22}$$

R is the estimated relative error defined to be one estimated standard deviation of the mean divided by the estimated mean

$$R = \frac{S_{\phi_{est}}}{\phi_{est}}, \tag{1.23}$$

where ϕ_{est} is the average score in the tally and an unbiased estimate of the variance of ϕ_{est} is given by,

$$S^2_{\phi_{est}} = \frac{1}{N-1}\left(\frac{1}{N}\sum_{n=1}^{N}\phi_n^2 - \phi_{est}^2\right). \tag{1.24}$$

The statistical error in an analog Monte Carlo simulation is inversely proportional to the square root of the particle history N. Hence, for deep penetration problems, to get an accurate estimate of the detector response, significantly large amount of particles need to be simulated which in turn result in large computation times. In applications like radiation monitoring for non-proliferation purposes it is essential to have an accurate estimate of the response variable in reasonable time. This requirement has led to the development of variance reduction techniques that would result in lesser computation times and increase the efficiency of the simulations.

1.3.2 Variance reduction techniques

Variance reduction techniques have been researched right from the emergence of Monte Carlo methods for radiation transport applications. The basic concept is to modify the sampling distributions to selectively transport the particles towards regions of interest (non analog) so that the variance in the measurement tally goes down as noted in Equation (1.24). To compensate for the bias introduced in the simulations, the concept of weight is introduced. Each particle is assigned an initial weight (w) and the weights are adjusted according to Equation 1.25 such that the contributions made to the tally are closer to the expected value and result in a lower variance than the analog (unbiased) simulation. This technique is also known as importance sampling [3].

$$w_{non-analog} \cdot PDF_{non-analog} = w_{analog} \cdot PDF_{analog}. \tag{1.25}$$

$PDF_{non-analog}$ is the modified distribution that transports particles towards regions of interest while PDF_{analog} is the natural distribution with which the particles travel. $w_{non-analog}$ and w_{analog} are the weights for non analog and analog Monte Carlo transport, respectively. In analog Monte Carlo, the weights of the particle would remain constant throughout their lifetime (usually $w_{analog} = 1$ for an unit volume source.) The major difficulty in employing variance reduction techniques is the determination of problem dependent variance reduction parameters present in the biased terms. There are simple techniques like forced collisions and implicit capture that are based

on the material properties of the problem and can be easily implemented [3]. But, more sophisticated techniques like weight windows, source biasing and angular biasing (described in chapter 2) require more information based on the tally's position, energy group and the material properties of all regions in the problem [4]. In the initial days, the parameters used by the biasing techniques were determined by the users based on their experience and using trial and error methods. Later developments lead to the usage of a crude solution of the problem as the basis to estimate the biasing parameters for the variance reduction methods. This crude solution is usually generated by performing an analog Monte Carlo calculation over a small particle history. But, the biasing distributions generated from an approximate Monte-Carlo simulation would not be very accurate, while using a more refined solution would in turn affect the overall efficiency. This led to the idea of using an approximate deterministic solution of the problem to generate the parameters required for variance reduction techniques in a Monte Carlo simulation. Deterministic methods are faster to solve than Monte Carlo simulations to get a reasonable estimate of the solution and are therefore ideal for generating the information needed in biasing methods. Thus, a new class of solution techniques that use a deterministic solution to reduce the variance of Monte Carlo simulations known as the hybrid methods (discussed in the next section) came into existence.

1.3.3 Hybrid methods

Hybrid methods for solving radiation transport problems have gained significant development in the past decade. They strive to bring together the advantages of the two traditional solution techniques for radiation transport problems, Monte Carlo simulations and deterministic techniques. The resulting methods are computationally more efficient than stand alone Monte Carlo simulations while retaining the accuracy benefits of the Monte-Carlo simulations over deterministic solutions [5]. Therefore, the hybrid methods are potentially capable of playing a crucial role in various non-proliferation research and monitoring scenarios like development of new radiation detection systems, analysis of source term associated with nuclear materials diversion scenarios, and the prediction of radiation doses from weapons of mass destruction.

Most hybrid variance reduction methods are developed based on the benefits of solving the adjoint problem [6] [4]. The adjoint transport equation is similar to the forward BTE (Equation (1.1)) and is given by,

$$\left[-\frac{1}{\nu(E)}\frac{d}{dt} + \Omega\cdot\nabla + \sigma_t(r,E)\right]\hat{\psi}\left(r,E,\Omega,t\right) = \int_{4\pi} d\Omega' \int_0^\infty dE'\sigma_s(r,E\rightarrow E',\Omega\cdot\Omega')\hat{\psi}\left(r,E',\Omega',t\right) + Q_d\left(r,E,\Omega,t\right), \quad (1.26)$$

where, Q_d is the detector response function. According to adjoint theory, the adjoint solution $\hat{\psi}$ gives the relative importance of the corresponding phase-space region to the detector response. Hence the adjoint solution, also known as the importance

function, can be effectively used to implement variance reduction techniques.

There are numerous hybrid methods that have been developed in the past, but only few simple techniques have found their way into production codes and even those are not straight forward to implement. A classic example is the provision to implement weight windows in MCNP (Monte Carlo N-Particle), the most common Monte Carlo simulation tool [2]. Although this provision has existed from the inception of MCNP, users need to be well versed with the problem in hand to develop weight windows or rely on other deterministic tools to generate them. The associated overhead in implementing these variance reduction features is often times not justifiable for the returns in computational efficiency. There are also other sophisticated hybrid techniques that have been tested only on restricted problems [6] [7]. In this project, one of the main objectives is to develop a platform that uses the results of a deterministic solution in a hybrid Monte-Carlo simulation with minimal user effort and overhead time and to test the performance of the variance reduction techniques on realistic test problems that are of interest to nonproliferation research.

Chapter 2: Literature Review

In this chapter we review some of the major hybrid variance reduction approaches that have been developed in the past and other acceleration schemes for Monte Carlo simulations that are based on adjoint solutions.

2.1 CADIS

The Consistent Adjoint Driven Importance Sampling (CADIS) method is developed at Oak Ridge National Labs (ORNL) and used primarily to optimize calculations using the multigroup Monte Carlo code Monaco, by creating an importance map based on the adjoint solution calculated with the discrete ordinates code Denovo [8], all of which are part of the SCALE 6 package [4].

2.1.1 CADIS based weight windows

In weight windows method, the particles are biased towards regions of higher importances in the phase-space region by applying Russian roulette and splitting techniques (described below) such that the variance in the weights of the particles traveling

through that region is reduced. A set of weight window target values and boundaries (upper and lower limits to perform splitting and rouletting, respectively) are constructed for every spatial cell and energy group. In the CADIS method, the weight window targets are constructed based on the importance function generated by performing an adjoint calculation on the problem using a deterministic tool.

In CADIS, the source $\widehat{q}$ for the adjoint calculation is the detector response function σ_d,

$$\widehat{q}\,(r, E) = \sigma_d\,(r, E) \tag{2.1}$$

where σ_d is usually the detector material absorption cross section, a set of flux to dose conversion factors or detector count rates per unit flux, etc. The target weight for energy group g and spatial cell n is then inversely proportional to the importance function and is given by,

$$w_{targ,g,n} = \frac{C}{\widehat{\phi}_{g,n}} \tag{2.2}$$

where C is a constant and $\widehat{\phi}_{g,n}$ is the adjoint scalar flux at cell n and energy group g. The weight window boundaries (w_{high} and w_{low}) are selected based on an arbitrary ratio of upper limit to lower limit and are usually based on trial and error as the overall efficiency of the method is quite robust [4].

The weight of the particle is checked whenever the particle crosses a cell boundary and whenever the particle scatters into a different energy group. If the particle's weight lies outside the weight window limits, the particle undergoes either splitting or rouletting. Splitting is the process in which if the particle's weight is above the

upper weight window boundary, the parent particle is split into a number of daughter particles such that the weight of each daughter is reduced and is closer to the weight window target and the aggregate of daughter weights would equal the parent's weight. After splitting, each daughter particle is treated as an individual particle and the transport process continues until all the daughters are terminated. Splitting helps to increase the survival rate of particles appearing in regions that are more likely to contribute to the detector tally by splitting them into a number of particles and as a result, more particles reach the detector with lower weights that are closer to the expected value of the tally, which in turn results in a lower variance of the estimate.

Rouletting, on the other hand, occurs when the particle's weight is below the lower weight window boundary. The survival probability of the particles undergoing rouletting is fixed to be an arbitrary value. If the particle survives rouletting, its weight is increased such that it is now within the weight window limits and continues travel whereas if it fails to survive, the particle's history is terminated. Rouletting ensures that computation time is not wasted in tracking particles with lower weights that are most likely not going to contribute to the detector tally.

2.1.2 CADIS based source biasing

Source biasing is a technique to ensure that particles are selectively born in the phase-space regions that are more likely to contribute to the detector tally of interest. The weights of the particle are adjusted to maintain a fair game. Source biasing plays a

crucial role in ensuring the new particles are born within the weight window limits and therefore do not undergo early termination.

The starting weight of the particle is given by the ratio of the original source distribution to the modified source distribution and is defined as:

$$w_0 = \frac{q(r,E)}{q_b(r,E)} \tag{2.3}$$

Using the definition for target weights in Equation (2.2), the biased distribution can be expressed as:

$$q_b(r,E) = \frac{1}{C} q\,(r,E)\,\widehat{\phi}\,(r,E) \tag{2.4}$$

The biased source distribution should be normalized such that it can be converted to a cumulative probability distribution function from which the source particles can be sampled. This is accomplished by setting the constant C equal to the detector response R ($particles/cm^3/sec$). By adjoint theory [4], the detector response can be defined as:

$$R = \int\int q(r,E)\widehat{\phi}\,(r,E)\,d^3r dE \tag{2.5}$$

The target weights can then be calculated using the formula,

$$w_{targ,g,n} = \frac{R}{\widehat{\phi}_{g,n}} \tag{2.6}$$

2.2 LIFT

The LIFT (The Local Importance Function Transform) method is based on the zero variance problem [9] and provided an exact adjoint solution is available a single particle history should be enough to find the exact forward solution. Of course, an exact adjoint solution is never possible and therefore the LIFT method, like the CADIS based techniques uses the approximate adjoint based importance function to bias the particles towards the regions of interest and yield a very low variance solution. The LIFT method is lot more sophisticated than weight windows and bias the entire physics of the particle transport. LIFT uses an adjoint angular flux expression based on the exponential transform functional locally in each spatial cell and energy group based on the importance function (adjoint scalar flux) and material properties. This angular flux reconstruction helps to bias the particle in the angular domain in addition to spatial and energy biasing. Another important idea in the LIFT method is the construction of a biasing parameter (ρ). The biasing parameter is a vector that points in the direction of the flux gradient within the cell and whose magnitude determines the amount of biasing. Some important concepts of the LIFT method are discussed below while a detailed explanation can be found in Turner [9] [6].

2.2.1 The zero variance problem and the LIFT approximation

The Boltzmann transport equation for the forward radiation problem and the corresponding adjoint problem are operated on each other to yield the transport equation for the zero variance problem (also known as the contributon transport equation) as given below.

$$\Omega \cdot \nabla \xi_g (r,\Omega) + \Sigma_{t,g} (r,\Omega)\, \xi_g (r,\Omega) = \frac{1}{4\pi} \sum_{g'=1}^{G} \int_{4\pi} \Sigma_{s,g'\to g} (r,\Omega)\, \xi_{g'} (r,\Omega')\, d\Omega' + \frac{\widehat{\psi}(r,\Omega)}{4\pi} Q(r) \delta_{g,1}, \tag{2.7}$$

where,

$$\xi_g (r,\Omega) = \psi_g (r,\Omega)\, \widehat{\psi}_g (r,\Omega)\,. \tag{2.8}$$

$\psi_g (r,\Omega)$ is the forward angular flux, $\widehat{\psi}_g (r,\Omega)$ is the adjoint angular flux and $\xi_g (r,\Omega)$ is known as the contributon flux. The boundary conditions are given by,

$$\xi_g (r_b,\Omega) = 0, \text{ for } \Omega \cdot n_b < 0, r_b \in \delta D, 1 \le g \le G, \tag{2.9}$$

where n_b is the outward normal vector and δD is the differential volume within the problem boundary. G is the total number of energy groups. The source is defined by,

$$Q(r) = \begin{cases} 1, & r \in D_s \\ 0, & \text{otherwise.} \end{cases} \tag{2.10}$$

where D_s is the source volume.

The zero variance transport equation exhibits certain unique features which make them ideal to optimize a Monte Carlo simulation. In a zero variance transport problem, all particles that are born in the source are absorbed only in the detector region and no particles can leak out of the system. The effective cross sections (Σ_t and Σ_s) for such a problem are given by,

$$\Sigma_{t,g}\left(r,\Omega\right) = \sigma_{t,g}(r) - \frac{1}{\widehat{\psi}_g(r,\Omega)}\Omega\cdot\nabla\widehat{\psi}_g\left(r,\Omega\right) \tag{2.11}$$

$$\Sigma_{s,g'\to g}\left(r,\Omega'\to\Omega\right) = \sigma_{s,g'\to g}(r,\Omega'\cdot\Omega)\frac{\widehat{\psi}_g(r,\Omega)}{\widehat{\psi}_{g'}(r,\Omega')} \tag{2.12}$$

Since we have access only to the adjoint scalar fluxes from the deterministic calculations, an expression to reconstruct the angular fluxes is essential to bias the calculation based on the zero variance problem. The LIFT method defines the adjoint angular flux based on the exponential transform with some modifications to account for linear anisotropic scattering [6].

2.2.2 Expression for adjoint flux reconstruction in the LIFT method

The LIFT approximation of the adjoint angular flux for each cell is given by,

$$\widehat{\psi}_{g,n} = \widehat{\phi}_{g,n}V_n\left[\beta_{g,n}\frac{\sigma_{s0,g\to g,n}b_{g,n}(\Omega)}{\sigma_{t,g,n}-\rho_{g,n}\cdot\Omega}\,e^{\rho_{g,n}\cdot(r-r_n)}\right] \tag{2.13}$$

$\widehat{\phi}_{g,n}$ is the scalar flux in group g and cell n, V_n is the volume of cell n and $\beta_{g,n}$is the normalization factor given by,

$$\beta_{g,n} = \left\{ \left[\int_{V_n} e^{\rho_{g,n} \cdot (r - r_n)} dr \right] \times \left[\int_{4\pi} \frac{\sigma_{s0,g \to g,n} b_{g,n}(\Omega)}{\sigma_{t,g,n} - \rho_{g,n} \cdot \Omega} d\Omega \right] \right\}^{-1} \tag{2.14}$$

and $b_{g,n}$ is the linear anisotropic factor given by,

$$b_{g,n}(\Omega) = 1 + 3\mu_{g \to g,n} \frac{\sigma_{t,g,n} - \sigma_{s0,g \to g,n}}{|\rho_{g,n}|^2} \rho_{g,n} \cdot \Omega \tag{2.15}$$

$\rho_{g,n}$ is known as the biasing parameter and is defined by,

$$\rho_{g,n} = \sigma_{t,g,n} \cdot \lambda_{g,n} \tag{2.16}$$

where λ is the deep penetration eigen value in the material. In the LIFT method however, the biasing parameters are determined locally in each spatial cell and energy group from corresponding adjoint solution. The biasing parameter is essentially a measure of the adjoint flux gradient in the cell and therefore is a vector quantity in three-dimensional space whose magnitude determines the amount of biasing. In Turner's work the biasing parameter in each spatial cell was determined using an approximation of the adjoint flux gradient from the cell edge fluxes calculated in the deterministic calculation. In this work, an improved method to calculate the biasing parameter from the adjoint current and flux ratio is adopted based on Kendra and Larsen's work [10].

2.3 VVR

The Variational Variance Reduction (VVR) method, developed at the University of Michigan has proven to be an effective technique for increasing the efficiency of Monte Carlo simulations for k-eigen value transport problems and can be extended to source-detector problems [11] [12] [13] [14]. This method is based on a zero-order problem that is different to the one considered in the LIFT technique and is based on variational principles. Let us consider the following simple isotropic, mono-energetic forward problem,

$$L\Psi = Q\left(r, \Omega\right) \tag{2.17}$$

where, Q is the source term and L is the forward transport operator defined by,

$$L\Psi = \Omega \cdot \nabla\Psi(r, \Omega) + \sigma_t\Psi(r, \Omega) - \sigma_s \int_{4\pi} \Psi(r, \Omega) d\Omega. \tag{2.18}$$

The corresponding adjoint problem is,

$$\hat{L}\hat{\Psi} = \Sigma\left(x, \Omega\right) \tag{2.19}$$

where, Σ is the detector response (and the source for the adjoint problem) and $\hat{L}$ is the adjoint transport operator defined by,

$$\hat{L}\hat{\Psi} = -\Omega \cdot \nabla\hat{\Psi}(r, \Omega) + \sigma_t\hat{\Psi}(r, \Omega) - \sigma_s \int_{4\pi} \hat{\Psi}(r, \Omega) d\Omega \tag{2.20}$$

Having defined the transport and adjoint operators, we can define variation functionals based on adjoint theory as below,

$$\begin{aligned} F[\hat{\Psi}, \Psi] &= \int_V \int_{4\pi} \Sigma \Psi dV d\Omega - \int_V \int_{4\pi} \hat{\Psi} \left(L\Psi - \frac{Q}{4\pi} \right) dV d\Omega \\ &\approx R \end{aligned} \tag{2.21}$$

This functional has several desirable properties:

- If $\Psi(\boldsymbol{x}, \boldsymbol{\Omega}) = \psi(\boldsymbol{x}, \boldsymbol{\Omega})$, then for any $\hat{\Psi}(\boldsymbol{x}, \boldsymbol{\Omega})$,

$$F[\hat{\Psi}, \psi] = \int_V \int_{4\pi} \frac{\Sigma(\boldsymbol{x})}{4\pi} \psi(\boldsymbol{x}, \boldsymbol{\Omega})\, d\Omega dV = R \tag{2.22a}$$

- If $\hat{\Psi}(\boldsymbol{x}, \boldsymbol{\Omega}) = \hat{\psi}(\boldsymbol{x}, \boldsymbol{\Omega})$, then for any $\Psi(\boldsymbol{x}, \boldsymbol{\Omega})$,

$$\begin{aligned} F[\hat{\psi}, \Psi] &= \int_V \int_{4\pi} \frac{\Sigma(\boldsymbol{x})}{4\pi} \Psi(\boldsymbol{x}, \boldsymbol{\Omega})\, d\Omega dV \\ &\quad - \int_V \int_{4\pi} \left[\left(\hat{L}\hat{\psi}(\boldsymbol{x}, \boldsymbol{\Omega}) \right) \Psi(\boldsymbol{x}, \boldsymbol{\Omega}) - \hat{\psi}(\boldsymbol{x}, \boldsymbol{\Omega}) \frac{Q(\boldsymbol{x})}{4\pi} \right] d\Omega dV \\ &= \int_V \int_{4\pi} \frac{Q(\boldsymbol{x})}{4\pi} \hat{\psi}(\boldsymbol{x}, \boldsymbol{\Omega})\, d\Omega dV \\ &= R \end{aligned} \tag{2.22b}$$

- If $\Psi(\boldsymbol{x},\boldsymbol{\Omega}) = \psi(\boldsymbol{x},\boldsymbol{\Omega}) + \delta\psi(\boldsymbol{x},\boldsymbol{\Omega})$ and $\hat{\Psi}(\boldsymbol{x},\boldsymbol{\Omega}) = \hat{\psi}(\boldsymbol{x},\boldsymbol{\Omega}) + \delta\hat{\psi}(\boldsymbol{x},\boldsymbol{\Omega})$, then:

$$
\begin{aligned}
F[\hat{\Psi},\Psi] &= F[\hat{\psi}+\delta\hat{\psi},\psi+\delta\psi] \\
&= \int_V\int_{4\pi}\frac{\Sigma}{4\pi}\big[\psi+\delta\psi\big]\,d\Omega dV - \int_V\int_{4\pi}\big[\hat{\psi}+\delta\hat{\psi}\big]\Big[L\psi+L\delta\psi-\frac{Q}{4\pi}\Big]\,d\Omega dV \\
&= \int_V\int_{4\pi}\frac{\Sigma}{4\pi}\big[\psi+\delta\psi\big]\,d\Omega dV - \int_V\int_{4\pi}\big[\hat{\psi}+\delta\hat{\psi}\big]\big[L\delta\psi\big]\,d\Omega dV \\
&= \int_V\int_{4\pi}\frac{\Sigma}{4\pi}\big[\psi+\delta\psi\big]\,d\Omega dV - \int_V\int_{4\pi}\big[\hat{L}\hat{\psi}+\hat{L}\delta\hat{\psi}\big]\big[\delta\psi\big]\,d\Omega dV \\
&= \int_V\int_{4\pi}\frac{\Sigma}{4\pi}\big[\psi+\delta\psi\big]\,d\Omega dV - \int_V\int_{4\pi}\big[\frac{\Sigma}{4\pi}+\hat{L}\delta\hat{\psi}\big]\big[\delta\psi\big]\,d\Omega dV \\
&= \int_V\int_{4\pi}\frac{\Sigma}{4\pi}\psi\,d\Omega dV - \int_V\int_{4\pi}\big[\hat{L}\delta\hat{\psi}\big]\big[\delta\psi\big]\,d\Omega dV = R + O(\delta^2)
\end{aligned}
\tag{2.22c}
$$

Therefore, if $\delta\hat{\psi}$ and $\delta\psi$ are small ($\hat{\Psi}$ and Ψ are first-order approximations to ψ and $\hat{\psi}$ respectively), then $F[\hat{\Psi},\Psi]$ is a second-order approximation of R.

In particular, suppose that in Equation (2.21), $\hat{\Psi}(\boldsymbol{x},\boldsymbol{\Omega}) = \hat{\psi}(\boldsymbol{x},\boldsymbol{\Omega})$, and $\Psi(\boldsymbol{x},\boldsymbol{\Omega})$ is an estimate of $\psi(\boldsymbol{x},\boldsymbol{\Omega})$ obtained from a forward Monte Carlo simulation. Then, by Eq. (2.22b), we have $F[\hat{\psi},\Psi] = R$, so we again have a zero variance method. Unfortunately, just as before, this method requires us to have complete knowledge of the adjoint flux $\hat{\psi}(\boldsymbol{x},\boldsymbol{\Omega})$, so this zero variance method is just as impractical as the one in the LIFT method.

However, as before, we can approximate this impractical zero-variance method in ways that render it practical. Suppose that we solve a deterministic adjoint problem

and construct representations of $\hat{\psi}(\boldsymbol{x}, \boldsymbol{\Omega})$ in each spatial cell V_j:

$$\hat{\Psi}(\boldsymbol{x}, \boldsymbol{\Omega}) = \hat{\Psi}_j(\boldsymbol{x}, \boldsymbol{\Omega}) \ \ , \ \ \boldsymbol{x} \in V_j \ \ . \tag{2.23}$$

Then by substituting in Equation (2.21),

$$\begin{aligned} F[\hat{\psi}, \Psi] &= \int_V \int_{4\pi} \frac{\Sigma}{4\pi} \Psi \, d\Omega dV - \sum_{j=1}^{J} \int_{V_j} \int_{4\pi} \hat{\Psi}_j \left[L\Psi - \frac{Q}{4\pi} \right] d\Omega dV, \\ &= \int_V \int_{4\pi} \frac{\Sigma}{4\pi} \Psi \, d\Omega dV \\ &\quad - \sum_{j=1}^{J} \left[\int_{\partial V_j} \int_{4\pi} \boldsymbol{\Omega} \cdot \boldsymbol{n} \hat{\Psi}_j \Psi_j d\Omega dS + \int_{V_j} \int_{4\pi} \left(\hat{L} \hat{\Psi}_j \right) \Psi \right] d\Omega dV \\ &\quad + \sum_{j=1}^{J} \int_{V_j} \int_{4\pi} \hat{\Psi} \frac{Q}{4\pi} \, d\Omega dV \ . \end{aligned} \tag{2.24}$$

The right side of this equation contains volume and surface integrals of Ψ over each cell V_j and its surface ∂V_j. The integrands are weighted by $\hat{\Psi}_j$. If $\hat{\Psi}_j$ are chosen to be "simple" functions, then these integrals can be estimated in the natural course of tracking Monte Carlo particles from the source to the detector. To implement this "Variational Variance Reduction" (VVR) method, it is necessary to tally information about the Monte Carlo particles everywhere in the system, not just in the detector region.

Chapter 3: Tortilla - Hybrid Methods Test Bench

To implement the hybrid methods, a multigroup Monte Carlo (MGMC) code that works within the framework of a commercially available deterministic radiation transport tool, Attila^{TM}was developed at Oregon State University with help from Transpire,Inc (Developers of Attila). In this project, the MGMC code is modified to implement hybrid methods and we call the entire framework to implement hybrid variance reduction methods as "Tortilla."

Tortilla consists of various sub units that interact with each other in a fashion such that the user can implement variance reduction techniques that make use of an approximate deterministic solution in a Monte-Carlo simulation. The problem geometry and the geometric mesh needs to be defined only once to perform both the deterministic and Monte-Carlo calculations. The cross section libraries and source definitions are also shared between the two calculations effortlessly using appropriate tools (Mr mixer and XREP). The basic setup and the specific features are explained below.

3.1 Basic architecture of the test bench

Figure 3.1 shows the basic architecture of the test bench. To perform a transport calculation using the hybrid methods in Tortilla, the user needs to define the problem using Attila's project manager just like they would to run a normal deterministic calculation in Attila. After the initial setup, the user can perform a forward or adjoint deterministic calculation and export the resulting flux moments for future use. Both the adjoint and forward source definitions can be created within the project manager. All the necessary data files (geometric mesh, cross section library, etc.,) from the deterministic calculation are exported to a new directory and the MGMC code is then executed to perform the Monte Carlo simulation.

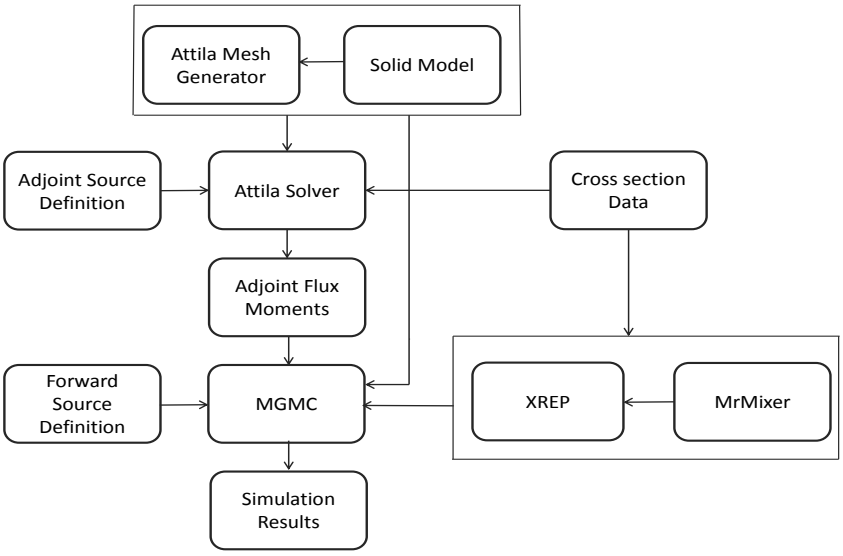

Figure 3.1: Tortilla - Basic architecture

3.2 Computational mesh

Any deterministic solver requires a computational mesh that approximates the spatial domain of the problem using simpler spatial cells in which the physical equations are then solved. The type of mesh used and its refinement play an important role in determining the performance of the deterministic methods. Monte Carlo simulations by nature do not require a spatial mesh and a simpler definition of the different material regions in the problem is suffice. However, for hybrid methods, the solution generated by the deterministic calculations can be applied to the Monte Carlo simulation only if a geometric mesh is available and in previous applications [15] a fine mesh is usually employed for generating the deterministic solution and the quantities are then transferred to a coarser mesh for performing the Monte Carlo simulation. In this project, one of the main objectives is to reduce the overhead associated with generating two different computational meshes for the deterministic and Monte Carlo part of the simulation and exchanging calculation parameters between them. Hence it was decided to use a single mesh for both the adjoint deterministic calculation and the forward Monte Carlo simulation. One limitation in using a single mesh is that, the mesh size needs to be optimal such that the deterministic calculation is fairly accurate but at the same time it does not reduce the efficiency of the Monte Carlo simulation and the variance reduction techniques.

Attila supports the RTT mesh file format developed at Los Alamos National Lab (LANL) [16]. There are few mesh generators like LaGrit (a public domain mesh gener-

ator from LANL), Tetra mesh generator from ICEM CFD Engineering and "'Aetius"' GUI interface (used in Attila) that can generate mesh files in the RTT format. The mesh generated should contain arbitrarily connected tetrahedral elements so that it could be used by the Attila solver. For our test problems the mesh is generated using Attila's GUI interface.

3.3 Cross section libraries

Attila can be used to solve radiation transport problems involving neutral particles like neutrons and photons as well as charged particles like electrons. The cross sections for a particular problem are usually generated from the master data sets using specific cross section processing utilities like NJOY-TRANSX (from LANL) and AMPX (from ORNL) for nuclear engineering problems and CEPXS(from SNL) for charged particle problems. Attila can accept cross section libraries generated from these tools as long as they are in one of the supported formats. Attila solver currently supports DTF, AMPX-77, MENDF and BXSLIB2 formats. However, the Attila project manager supports only the DTF format for now (version 7), but there are codes available to convert from other formats to DTF format [16].

The mixing of material cross sections to calculate the macroscopic cross sections for the different regions of the problem is performed within the Attila solver. The density, atomic weights and atomic or mass fractions of the isotopes in different regions are required during problem setup in Attila's GUI and the rest of the calculations are

performed by algorithms. Further, the Attila solver also produces a material definition file, that contains the densities and isotope fractions for a particular problem, which could be used outside the solver to generate the material cross sections.

3.4 Source definitions

Attila provides for user input of both volumetric and boundary anisotropic, multi-group, radiation sources. Both isotropic and anisotropic point sources can also be described using the point source definition file. It is also possible to build a spectral distribution within the project manager and assign them to specific regions in the problem. A source file (in RTT format) would be generated after the calculation is performed which can then be used by the MGMC code [16].

3.5 Deterministic solver

The deterministic solver in Attila is used to generate the forward or adjoint flux moments for the problem. These flux moments would then be used by variance reduction methods in the Monte Carlo simulation. Attila's deterministic solver can solve three dimensional transport problems using the Linearly Discontinuous (LD) spatial differencing scheme. This is a well damped spatial differencing scheme that is second order accurate locally and third order accurate in global quantities [17]. The

nature of this differencing scheme is that it is not strictly positive, but oscillations are damped and the scheme is stable. Finer meshes result in more accurate solutions, but because of the advanced differencing scheme and the use of arbitrary tetrahedral meshes, the number of cells required in comparison to an equivalent traditional mesh is significantly lower.

The discretization in the angular domain is done using the discrete ordinates (S_N) method [18] and particle scattering is handled by spherical harmonics expansion of the scattering source. Multi-group discretization is employed in the energy domain. The resulting solution method is effective for most transport problems. In cases of optically thick problems, there is a provision to reduce the number of source iterations by employing the simplified diffusion synthetic acceleration (DSA) method.

The flux moments calculated in an Attila calculation are saved in a moments file that would be imported into the MGMC code. This moments file contains all the flux moments data at the vertices of each tetrahedral cell and the cell averaged quantities. There are also other output edits that such as a visualization link file for GMV (General Mesh Viewer) or Tecplot that can be requested [16].

3.6 Preprocessing routines

Before executing the Monte Carlo calculation a preprocessing script needs to be executed. This script would invoke the routines to calculate the macroscopic material

cross sections (MrMixer), the cumulative probability distribution for sampling scattering directions (XREP) and other routines to make sure the format of the data files from the deterministic calculations are compatible with the MGMC code.

3.6.1 MrMixer

In Attila the cross section mixing is done entirely within the program while for external uses a tool for mixing material cross sections known as MrMixer is used [16]. A unique feature of MrMixer is its ability to read a file of standard material definitions and mix these materials by volume fraction into regions. The input for MrMixer is an XML file generated by the Attila calculation that has the material densities and atom fractions for each region in the problem. The output from MrMixer is an HDF5 format file containing both the microscopic and the region wise macroscopic cross sections that can be imported into the MGMC code.

3.6.2 XREP

XREP is a program developed to aid multigroup calculations in MCNP [19]. The main purpose of this program is to eliminate the negative scattering cross sections that may arise due to the truncation of the Legendre polynomials. X-rep uses a maximum entropy approach to the generation of group-ta-group scattering cross sections for

Monte Carlo. It allows for arbitrary number of scattering bins based on the scattering cross section moments that are devoid of any ray effects and provides better accuracy and efficiency for multigroup Monte Carlo calculations.

3.7 Multigroup Monte Carlo (MGMC) code

The multigroup Monte Carlo code is developed in C++ using an object oriented approach and making use of open source libraries and functions for efficient implementation of the methods. Some of the important features of the code and methods for transporting particles in the Monte Carlo simulation are discussed below.

3.7.1 Pre-defined libraries and APIs

The MGMC code is developed by making use of predefined libraries from BOOST and Transpire. The BOOST libraries are a group of open source C++ libraries for general purpose functions that are optimized to provide efficient operation and memory management [20] [21]. Apart from these general purpose functions, Transpire has provided a library (which is also developed using BOOST) that contains various APIs (Application Programmable Interface), classes and functions that take advantage of the polymorphism and inheritance principles in C++. The APIs are used to import the data files needed to implement an independent Monte Carlo simulation on

a problem that has been setup using Attila's project manager. These APIs enable the MGMC code to interact seamlessly with the data from the Attila calculation. The four main data files that are imported using the APIs into the MGMC code are the source definition file, geometric mesh file, cross section library and the flux moments file.

The Transpire libraries also contain predefined functions that facilitate direct access to specific components of the data files imported through the APIs, handle vector operations, multi-dimensional arrays and other functions that are useful for implementing particle transport in three dimensions.

3.7.2 Random number generator

It was determined that the in-built pseudo random number generator in C++ produced some discrepancies in the flux values observed for large particle history and therefore a sophisticated 64 bit random number generator developed by Mersenne Twister is employed [22]. This allows for three different seeding options and is widely used in many applications. It has a long period which ensures that the random numbers generated are highly independent (uncorrelated), a desired feature for Monte Carlo simulations. [23].

3.7.3 Numerical solvers

Although the analog Monte Carlo simulation does not require any numerical solutions, some of the variance reduction techniques that are implemented require numerical solutions and hence secant based (more efficient but not robust) and bisection (robust but less efficient) based numerical solvers are included in the code.

3.7.4 Particle transport

In this section, the general flow of the program and specific methodologies adapted for sampling various parameters in the Monte Carlo simulation are detailed below.

1. The first step is to sample the source particle's location and energy group. This is done by forming a cumulative probability distribution function based on the source strength in each spatial cell and energy group. The location is first sampled, followed by the energy group directly using pseudo random numbers.

2. The second step is to sample an initial position and direction within the source cell. The initial direction is sampled isotropically using a rejection technique described in the LANL report LA-9721-MS [24]. The initial position of the particle within the cell in an arbitrary tetrahedral mesh is uniformly sampled by using a random set of barycentric co-ordinates.

3. Next, the optical distance to collision is uniformly sampled from a random

logarithmic distribution which is a function of the total cross section. From the direction and distance to collision the new position of the particle is determined.

4. If the new position is outside the cell, the particle is moved to the cell boundary and the distance traveled within the cell is scored in the tally. If the remaining distance the particle needs to travel is above a lower threshold, the particle transport continues into the new cell. If the distance to collision is below the threshold, a new distance to collision is sampled.

5. If the new position is within the cell, the particle undergoes a collision reaction. In order to avoid early termination of particles in highly absorbing problems, a simple survival biasing scheme (also known as forced collisions) is employed.

6. The particle's weight is adjusted after every collision to account for absorption and once the weight falls below a lower threshold, a rouletting operation is performed to determine the survival of the particle.

7. If the particle survives, the particle's new energy group and direction are sampled. The scattering bins that are calculated using XREP are used to compute the new direction while the energy group is determined directly from a cumulative probability distribution function of the group to group scattering cross sections.

8. The above process is repeated until the particle is terminated in rouletting scheme or reaches the problem boundary. After all the particle histories are

simulated, the region wise flux values and statistical error are determined from the path length tallies.

Chapter 4: LIFT Method Implementation

In this chapter some of the major implementation details of the LIFT method in Tortilla are presented. A more detailed description of the LIFT method can be found in Turner's thesis [9].

4.1 Biasing parameter calculation

The estimation of the biasing parameter $\boldsymbol{\rho}$ is one of the key components of the LIFT implementation; the more accurate the biasing parameter, the more effective the particle biasing and greater the efficiency of the method. One way to estimate an accurate biasing parameter would be to calculate it based on the adjoint deterministic solution. In the original LIFT implementation by Turner and Larsen [6], an expression based on the cell edge adjoint scalar fluxes was employed. In this method, an expression for the adjoint scalar flux in the cell is constructed and its gradient is calculated to generate the biasing parameter. A detailed procedure for implementing this method on a tetrahedral grid is presented in reference [25]. This approach calculates the biasing parameter from the adjoint scalar fluxes directly, but has its disadvantages when applied to heterogeneous problems. At region boundaries, the magnitude of the biasing parameter exhibits undamped oscillations and needs to be

limited to a maximum value for the effective total cross section ($\Sigma_t \;=\; \sigma_t - \mathbf{\Omega} \cdot \boldsymbol{\rho}$) to remain positive [6].

A new method to determine the biasing parameter from the adjoint current-flux ratio ($\widehat{\boldsymbol{J}}$ / $\widehat{\phi}$ method) was developed recently and was tested on a 1-D mono-energetic slab problem [10]. The results showed that this method is more efficient than the edge-flux based biasing parameter calculation and the magnitude of the biasing parameters are damped near the material region boundaries. We use this method to determine the magnitude of the biasing parameter $|\boldsymbol{\rho}|$ and determine its direction from the definition of adjoint current. This method is based on the fact that there exists a unique solution for each value of the scattering ratio c between 0 and 1 for the plane geometry transport problem in a half-space defined by,

$$\widehat{\psi}_c(x,\mu) \;=\; f_c\,(\mu)\, e^{\,|\boldsymbol{\rho}|\,x}. \tag{4.1}$$

The above equation defines the asymptotic exponential decay of the angular flux for a given value of c. Also, both the angular function $f_c\,(\mu)$ and the eigenvalue $\boldsymbol{\rho}$ depend on the scattering ratio c. The adjoint current-flux ratio is given by,

$$\frac{|\widehat{\boldsymbol{J}}|}{\widehat{\phi}} \;=\; \frac{\int\limits_{-1}^{1} \mu f_c\,(\mu)\, d\mu}{\int\limits_{-1}^{1} f_c\,(\mu)\, d\mu}, \tag{4.2}$$

We take the magnitude of the current $\widehat{\boldsymbol{J}}$ and biasing parameter $\boldsymbol{\rho}$ as they are vector quantities in a 3-D calculation. In the above equation the ratio $|\widehat{\boldsymbol{J}}|/\widehat{\phi}$ varies from 1 to

0 as c varies from 0 to 1. $|\boldsymbol{\rho}|$ decreases monotonically from σ_t to 0 as c increases from 0 to 1. Therefore, for every value of $|\widehat{\boldsymbol{J}}|/\widehat{\phi}$ an unique value of $|\boldsymbol{\rho}|$ can be specified. The angular function f_c can be defined in terms of a Legendre polynomial expansion for anisotropic scattering. The angular integrals in Equation (4.2) are then evaluated by neglecting all the higher order polynomials in the expansion of f_c to yield the expression:

$$\frac{|\widehat{\boldsymbol{J}}|}{\widehat{\phi}} \;=\; \frac{\Sigma_t}{|\boldsymbol{\rho}|}\,(1-c)\,. \tag{4.3}$$

A detailed derivation of the angular integration and calculation of the polynomial moments are given in [26].

The $|\widehat{\boldsymbol{J}}|/\widehat{\phi}$ method, though slightly more complex than the edge-flux based method, is more accurate and more importantly does not require a maximum limit on the magnitude of the biasing parameter. To calculate the biasing parameter $|\boldsymbol{\rho}|$, a table of 'k' discrete values of $|\boldsymbol{\rho}_k|$ on the interval $(0,\,\sigma_t)$ is generated and solved for the corresponding values of c_k using the dispersion law relation [10],

$$\int_{-1}^{1} f_c(\mu)d\mu \;=\; 1, \tag{4.4}$$

which after the polynomial expansion of f_c results in,

$$1 = c\left[\sum_{0}^{L_{max}} a_l q_l(\nu) r_l(\rho)\right], \tag{4.5}$$

where, the a_l are the scattering moments given by,

$$a_l = \frac{\sigma_{s,l}}{\sigma_{s,0}}, \tag{4.6}$$

the $r_l(\rho)$ are integrals of the form

$$r_l(\rho) = \frac{2l+1}{2} \int_{-1}^{1} \frac{P_l(\mu)}{1 - \frac{\mu}{\nu}} d\mu \ , \tag{4.7}$$

where P_l are the Legendre polynomials and $\nu = \frac{\sigma_t}{\rho}$. $q_l(\nu)$ are polynomials given by [26]:

$$q_0(\nu) = 1, \quad q_1(\nu) = (1-c)\nu, \quad q_2(\nu) = \left[\frac{3}{2}(1 - ca_1)(1-c)\right]\nu^2 - \frac{1}{2}, \ldots \tag{4.8}$$

The adjoint current-to-flux ratio at each discrete value of 'k' is then calculated in terms of $|\boldsymbol{\rho}_k|$ and c_k and a lookup table of $|\widehat{\boldsymbol{J}}|/\widehat{\phi}$ vs $|\boldsymbol{\rho}_k|$ for each cell n and energy group g is generated.

$$\left(\frac{|\widehat{\boldsymbol{J}}|}{\widehat{\phi}}\right)_{k,n,g} = \frac{\Sigma_{t,n,g}}{|\boldsymbol{\rho}_{k,n,g}|}\left(1 - c_k\right). \tag{4.9}$$

During the simulation, in each cell and energy group, the magnitude of the biasing parameter $\boldsymbol{\rho}_{n,g}$ is calculated by interpolating from the table using the corresponding adjoint quantities $(|\widehat{\boldsymbol{J}}_{n,g}|, \widehat{\phi}_{n,g})$.

Once the magnitude is determined, the direction of the biasing parameter $\boldsymbol{\rho}$ is calculated from the adjoint current $\widehat{\boldsymbol{J}}$. Since the biasing parameter guides the particles towards the detector (adjoint source), it assumes the direction of the adjoint flux

gradient within the cell. Traditionally, the adjoint particles are thought of traveling in reverse directions to that of the forward particles and therefore the direction of the adjoint current is the same as the direction of adjoint flux gradient within the cell. Therefore,

$$\boldsymbol{\rho}_{n,g} = |\boldsymbol{\rho}_{n,g}| \frac{\widehat{\boldsymbol{J}}_{n,g}}{|\widehat{\boldsymbol{J}}_{n,g}|}. \tag{4.10}$$

However, in most production codes like Attila, the adjoint calculations are performed by just interchanging the location of detector and source and transposing the scattering cross section matrices. In such cases care should be taken to reverse the direction of the current to obtain the correct adjoint current direction.

4.2 Modifications for a tetrahedral mesh

The original LIFT method was implemented in a three dimensional lattice with rectangular Monte Carlo cells. A fine mesh was used for the adjoint calculation and a coarse mesh for the Monte Carlo cells [15]. In our implementation we use the same unstructured tetrahedral mesh generated for the adjoint deterministic calculation in Attila. As a result, the evaluation of the spatial part of LIFT's angular flux expression is modified. The evaluation of the volume integral term in the normalization factor $\beta_{g,n}$ defined in Equation (2.14) is not straight forward for cells in an arbitrary unstructured tetrahedral mesh. A coordinate transformation of the vertices of the tetrahedral cells was employed to determine the volume integral terms. Figure 4.1 shows the unnormalized cell in global coordinates x,y and z and figure 4.2 shows the

normalized tet cell in the local coordinate system defined by ε, η and ζ.

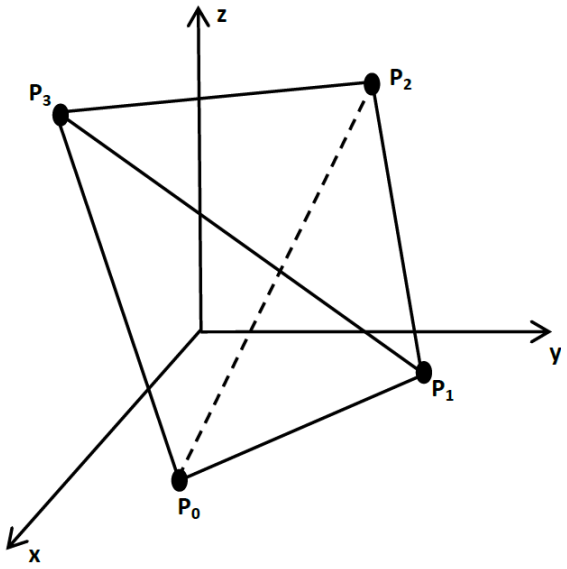

Figure 4.1: Unstructured Tet

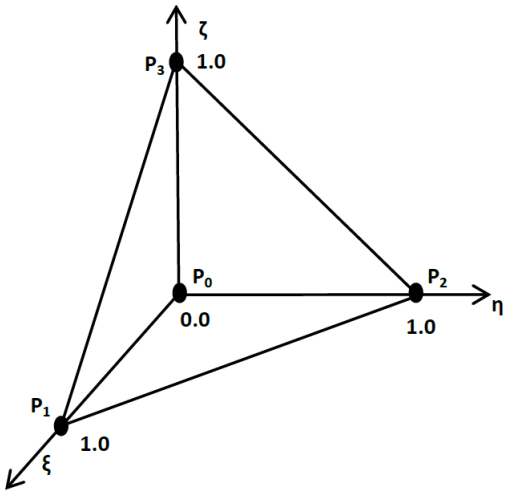

Figure 4.2: Normalized Tet

The volume integral in Equation (2.14) is given by,

$$\int_{V_n} e^{\rho_{g,n}\cdot r} dr = \int_{V_n} e^{\rho_{g,n}\cdot(x\widehat{i}+y\widehat{j}+z\widehat{k})} dxdydz \tag{4.11}$$

This integral is hard to evaluate directly, but after coordinate transformation to a normalized tetrahedron, the integration limits can be found fairly easily [27]. The transformed integral is then given by,

$$\int_{V_n} e^{\rho_{g,n}\cdot r} dr = \mathrm{Det}(\overline{\boldsymbol{J}}) e^{\rho_{g,n}\cdot r_0} \int_{\varepsilon=0}^{1} \int_{\eta=0}^{1-\varepsilon} \int_{\zeta=0}^{1-\varepsilon-\eta} e^{\rho_{g,n}\cdot \overline{\boldsymbol{J}}\delta} d\varepsilon d\eta d\zeta \tag{4.12}$$

where, r_0, r_1, r_2 and r_3 are the actual position vectors for nodes 1 to 4 of the tetrahedral cell. They are given by,

$$\begin{aligned} r_0 &\rightarrow (x_0\widehat{i} + y_0\widehat{j} + z_0\widehat{k}) \\ r_1 &\rightarrow (x_1\widehat{i} + y_1\widehat{j} + z_1\widehat{k}) \\ r_2 &\rightarrow (x_2\widehat{i} + y_2\widehat{j} + z_2\widehat{k}) \\ r_3 &\rightarrow (x_3\widehat{i} + y_3\widehat{j} + z_3\widehat{k}) \end{aligned}$$

$\overline{\boldsymbol{J}}$ is the transformation Jacobian matrix and is given by,

$$\overline{\boldsymbol{J}} = \begin{bmatrix} (x_1 - x_0) & (x_2 - x_0) & (x_3 - x_0) \\ (y_1 - y_0) & (y_2 - y_0) & (y_3 - y_0) \\ (z_1 - z_0) & (z_2 - z_0) & (z_3 - z_0) \end{bmatrix} \tag{4.13}$$

δ is represented by the transformed coordinates as given below,

$$\{\delta\} = (\varepsilon, \eta, \zeta)^T \tag{4.14}$$

4.3 Biasing methods

4.3.1 Source biasing with angle:

The distribution q_g for sampling the source particle for the transformed problem is given by,

$$q_g\left(r_n, \Omega\right) = \frac{Q_g(r_n)\widehat{\psi}_g\left(r_n, \Omega\right)}{\sum_g \sum_n \int_{4\pi} Q_g(r_n)\widehat{\psi}_g\left(r_n, \Omega\right) d\Omega dr_n} \tag{4.15}$$

This distribution is similar to the modified source distribution in the CADIS method (Equation 2.3), but has an angular dependence. The starting weight of the particle is adjusted to account for the non analog source sampling and is defined as,

$$w_s = \frac{R}{\widehat{\psi}\left(r_n, \Omega\right)}, \tag{4.16}$$

where R is the detector response given by,

$$R = \frac{1}{4\pi} \sum_g \sum_n \int_{D_{s,n}} \int_{4\pi} Q_g(r_n)\widehat{\psi}_g\left(r_n, \Omega\right) d\Omega dr_n.$$

4.3.2 Distance to collision:

The source particles are then transported within the problem geometry based on the parameter known as the distance to collision. Distance to collision is sampled from a probability distribution which is a function of the total cross section. The biased distance to collision probability function is given by,

$$p(s) = \Sigma_{t,g,n} e^{\Sigma_{t,g,n} s} \tag{4.17}$$

The distance to next collision s is sampled from the above function. The particle is now moved to see if it collides or escapes the cell. If the particle escapes, the particle is stopped at the boundary and a new distance to collision is sampled. The weight of the particle should be adjusted to account for its escape from the cell and is given by,

$$W_{esc} = e^{\rho_{g,n} \cdot (r_i - r_f)} \tag{4.18}$$

where, r_i is the initial particle position and $r_f$is the final particle position after it reaches the boundary.

4.3.3 Survival biasing:

If the particle undergoes collision within the cell, survival biasing is employed. It is a technique used to prevent excessive particle absorption in deep shielding problems by

allowing the particle to survive after collisions until it reaches the detector location. The particle's weight is again adjusted to account for its absorption probability. This method helps to reduce the number of histories required to achieve a given accuracy. The weight adjustment is given by the formula,

$$w_c = \frac{\frac{1}{4\pi}\sum_{g=1}^{G}\int_{4\pi}\Sigma_{s,g'\to g}\left(r_n,\Omega'\to\Omega\right)d\Omega}{\Sigma_{t,g'}\left(r_n,\Omega\right)} \tag{4.19}$$

Biased scattering After collision, the particle should undergo scattering and direction change. These two actions are also biased.The new energy group is sampled using the PDF,

$$p(g) = \frac{\int_{4\pi}\Sigma_{s,g'\to g}\left(r_n,\Omega'\to\Omega''\right)d\Omega''}{\sum_{g''=1}^{G}\int_{4\pi}\Sigma_{s,g'\to g''}\left(r_n,\Omega'\to\Omega''\right)d\Omega''} \tag{4.20}$$

The new direction is sampled using the PDF,

$$p\left(\Omega\right) = \frac{\Sigma_{s,g'\to g}\left(r_n,\Omega'\to\Omega''\right)}{\int_{4\pi}\Sigma_{s,g'\to g}\left(r_n,\Omega'\to\Omega''\right)d\Omega''} \tag{4.21}$$

Rejection sampling is used to sample the new energy group and directions. The formula for transformed scattering cross section is given by Equation (2.12).

A schematic of the various steps involved in implementing LIFT method in the MGMC code is given in figure 4.3. Although, the MGMC code can handle higher order scattering moments for analog and weight windows calculations using XREP

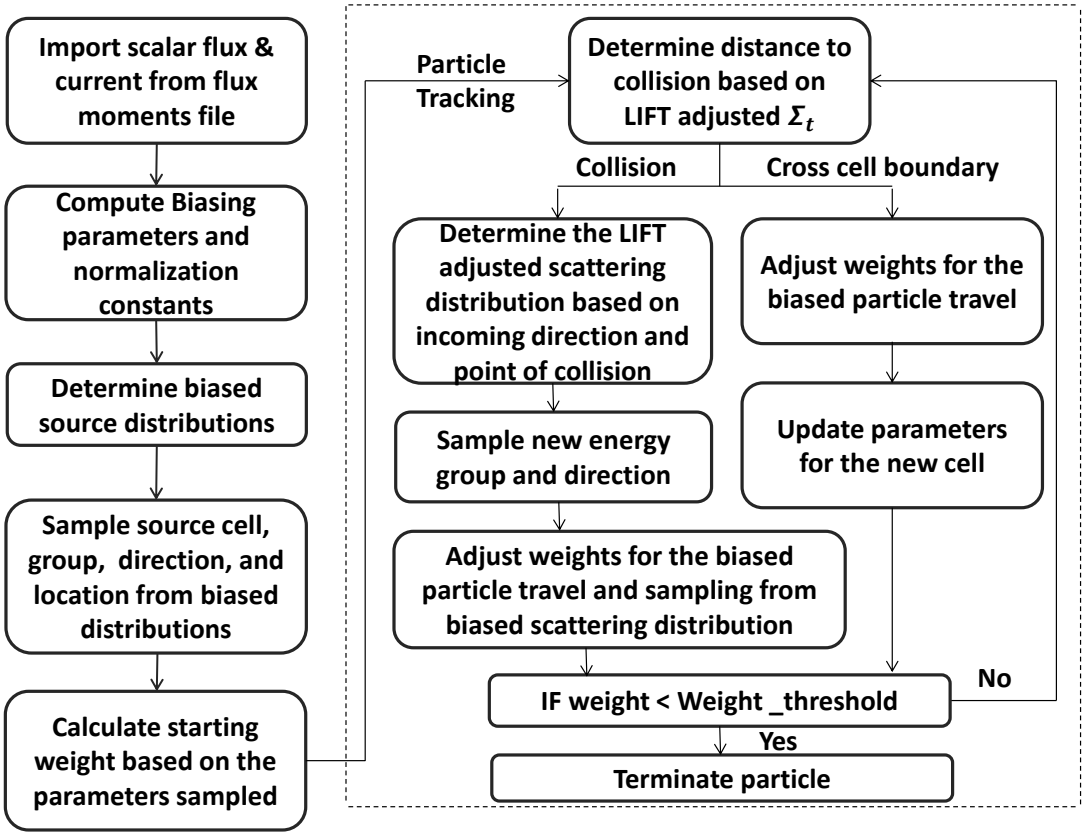

Figure 4.3: Schematic of the LIFT procedure as implemented in MGMC

[19], the LIFT method in its current form only allows for limited linear anisotropy as it calculates the effective scattering cross-sections based on the particle's incoming direction, energy group and the adjoint solution in the phase-space region [6]. In our implementation, all the calculations employ a simple rejection sampling scheme [9] to sample from a scattering distribution which allows for some anisotropy, such that the comparisons on the efficiencies of the different methods are not skewed.

In the original LIFT implementation a fixed coarse mesh Cartesian grid for the for-

ward Monte Carlo calculation was employed and the adjoint deterministic calculations were carried out using finer meshes with sizes optimized for the type of discretization scheme (several S_N and P_N methods were used). In this work, a single mesh with an optimal size is used for both forward and adjoint calculations in order to eliminate the overhead associated with employing two different meshes. However, for some test cases of the mulch problem, the CADIS based weight windows method performed better than the more sophisticated LIFT method. Since in the LIFT method, the non-analog distributions that determine the particle transport are calculated based on the particle's current spatial location, energy group and direction, using a fine mesh could potentially be a cause for the reduced efficiency of the LIFT method. Also, the LIFT method relies more heavily on the adjoint solution as it uses the 1^{st} flux moments (current) in addition to the 0^{th} moment (scalar flux) to achieve angular biasing as opposed to CADIS based weight windows that rely only on the 0^{th} moment of the adjoint. But for a more accurate adjoint solution, it is often required to use a fine mesh for the deterministic calculation. So, to test if the performance of the hybrid methods can be improved by using a coarser mesh for the forward calculation, an efficient nearest neighbor search algorithm and inverse distance weighting based interpolation scheme is developed to transfer the adjoint quantities from a fine arbitrary tetrahedral mesh to a coarser mesh. The algorithm for this method is presented in Appendix A.

Chapter 5: VVR Method Implementation

The VVR method has been implemented for three-dimensional source-detector type problems in Barrett's dissertation [7] in a relatively simple cube problem with water as the medium. In this implementation the angular domain is divided into octants and the forward and adjoint flux in each octant were then used to calculate the variational functional. Since this method is similar to a discrete ordinates (S_N) discretization of the angular domain, we will refer to it as the (S_8) implementation. The VVR implementation for criticality problems were based on spherical harmonics representation (P_N) of the flux moments [12] and therefore a VVR implementation using the first order spherical harmonic moments (P_1) is also performed. The P_1 implementation is computationally faster as the functional needs to be evaluated only 0^{th} and 1^{st} moments as opposed to the S_2 implementation in which the functional needs to be evaluated individually for each of the 8 octants. In both previous implementations, isotropic scattering cross sections were assumed. In this work, the VVR functionals are implemented with an approximation that accounts for linearly anisotropic scattering cross sections. Also, a modified form of the functional that speeds up the calculation of scattering source term in problems with many energy groups is adopted.

5.1 General form of the VVR functional for source detector problems

Recall from Equations (2.19), (2.17) and (2.21), for a forward problem defined by

$$L\Psi(\boldsymbol{r},\boldsymbol{\Omega}) = \frac{Q\left(\boldsymbol{r},\boldsymbol{\Omega}\right)}{4\pi}, \tag{5.1}$$

where L is the forward transport operator, Ψ is the exact forward angular flux and Q is the source and the equivalent adjoint problem defined by,

$$\widehat{L}\widehat{\Psi}(\boldsymbol{r},\boldsymbol{\Omega}) = \frac{\Sigma}{4\pi}, \tag{5.2}$$

where $\widehat{L}$ is the adjoint transport operator, $\widehat{\Psi}$ is the exact adjoint angular flux and Σ ($= \sigma_a\delta(r_d)$) is the detector response function for the forward problem. The variational functional for detector response R is given by,

$$\begin{aligned} F\left[\widehat{\Psi},\Psi\right] = & \frac{1}{4\pi}\int_V\int_{4\pi}\Sigma\Psi(\boldsymbol{r},\boldsymbol{\Omega})dVd\boldsymbol{\Omega} \\ & - \int_V\int_{4\pi}\widehat{\Psi}(\boldsymbol{r},\boldsymbol{\Omega})\left(L\Psi(\boldsymbol{r},\boldsymbol{\Omega}) - \frac{Q(\boldsymbol{r},\boldsymbol{\Omega})}{4\pi}\right)dVd\boldsymbol{\Omega} \\ & + \oint_S\int_{4\pi}(\boldsymbol{\Omega}\cdot\boldsymbol{n})\ \widehat{\Psi}(\boldsymbol{r}_b,\boldsymbol{\Omega})\ \Psi(\boldsymbol{r}_b,\boldsymbol{\Omega})dSd\boldsymbol{\Omega} \\ & \approx R. \end{aligned} \tag{5.3}$$

In the above functional, $\boldsymbol{r}_b$ denotes the problem boundary ($\boldsymbol{r}_b \in dV$). For problems with vacuum boundary conditions,

$$\widehat{\Psi}(\boldsymbol{r}_b, \boldsymbol{\Omega}) = 0, \text{for all } \boldsymbol{\Omega} \cdot \boldsymbol{n} > 0 \tag{5.4}$$
$$\Psi(\boldsymbol{r}_b, \boldsymbol{\Omega}) = 0, \text{for all } \boldsymbol{\Omega} \cdot \boldsymbol{n} < 0$$

and the functional reduces to,

$$F\left[\widehat{\Psi}, \Psi\right] = \int_V \int_{4\pi} \left[\frac{\Sigma}{4\pi}\Psi(\boldsymbol{r}, \boldsymbol{\Omega}) - \widehat{\Psi}(\boldsymbol{r}, \boldsymbol{\Omega}) \left(L\Psi(\boldsymbol{r}, \boldsymbol{\Omega}) - \frac{Q(\boldsymbol{r}, \boldsymbol{\Omega})}{4\pi}\right)\right] dV d\boldsymbol{\Omega} \tag{5.5}$$

The above functional is rewritten for a grid with J cells as:

$$F[\widehat{\psi}_j, \psi_j] = \sum_{j=1}^{J} \int_{V_j} \int_{4\pi} \left[\frac{\Sigma_j}{4\pi}\psi_j(\boldsymbol{r}, \boldsymbol{\Omega}) - \widehat{\psi}_j(\boldsymbol{r}, \boldsymbol{\Omega}) \left(L_j\psi_j(\boldsymbol{r}, \boldsymbol{\Omega}) - \frac{Q_j}{4\pi}\right)\right] dV d\boldsymbol{\Omega}, \tag{5.6}$$

where each cell has volume V_j. Expanding the transport operator L, the functional can be written as:

$$\begin{aligned} F[\widehat{\psi}_j, \psi_j] = & \frac{1}{4\pi} \sum_{j=1}^{J} \int_{4\pi} \int_{V_j} \Sigma\psi_j(\boldsymbol{r}, \boldsymbol{\Omega})\, dV d\boldsymbol{\Omega} \\ & - \sum_{j=1}^{J} \int_{4\pi} \int_{V_j} \widehat{\psi}_j(\boldsymbol{r}, \boldsymbol{\Omega}) \Big[\boldsymbol{\Omega}\nabla\psi_j(\boldsymbol{r}, \boldsymbol{\Omega}) + \sigma_{t,j}\psi_j(\boldsymbol{r}, \boldsymbol{\Omega}) \\ & - \frac{1}{4\pi} \int_{4\pi} \sigma_{s,j}\left(\boldsymbol{\Omega}' \to \boldsymbol{\Omega}\right) \psi_j(\boldsymbol{r}, \boldsymbol{\Omega}') d\boldsymbol{\Omega}' - \frac{Q_j}{4\pi}\Big] dV d\boldsymbol{\Omega}. \end{aligned} \tag{5.7}$$

Assuming the adjoint angular flux to be a piece-wise constant within each spatial cell and moving the adjoint angular flux inside the residual term,

$$
\begin{aligned}
F[\widehat{\psi}_j, \psi_j] = & \frac{1}{4\pi} \sum_{j=1}^{J} \int_{4\pi} \int_{V_j} \Sigma \psi_j(\boldsymbol{r}, \boldsymbol{\Omega}) \, dV \, d\boldsymbol{\Omega} \\
& - \sum_{j=1}^{J} \int_{4\pi} \widehat{\psi}_j(\boldsymbol{\Omega}) \int_{V_j} \boldsymbol{\Omega} \nabla \psi_j(\boldsymbol{r}, \boldsymbol{\Omega}) \, dV \, d\boldsymbol{\Omega} - \sum_{j=1}^{J} \sigma_{t,j} \int_{4\pi} \widehat{\psi}_j(\boldsymbol{\Omega}) \int_{V_j} \psi_j(\boldsymbol{r}, \boldsymbol{\Omega}) \, dV \, d\boldsymbol{\Omega} \\
& + \frac{1}{4\pi} \sum_{j=1}^{J} \int_{4\pi} \widehat{\psi}_j(\boldsymbol{\Omega}) \int_{V_j} \int_{4\pi} \sigma_{s,j} \left(\boldsymbol{\Omega}' \rightarrow \boldsymbol{\Omega} \right) \psi_j(\boldsymbol{r}, \boldsymbol{\Omega}') d\boldsymbol{\Omega}' \, dV \, d\boldsymbol{\Omega} \\
& + \sum_{j=1}^{J} \int_{4\pi} \widehat{\psi}_j(\boldsymbol{\Omega}) \int_{V_j} \frac{Q_j}{4\pi} \, dV \, d\boldsymbol{\Omega}.
\end{aligned} \tag{5.8}
$$

Applying the divergence theorem to the streaming term, the functional reduces to:

$$
\begin{aligned}
F[\widehat{\psi}_j, \psi_j] = & \frac{1}{4\pi} \sum_{j=1}^{J} \int_{4\pi} \int_{V_j} \Sigma \psi_j(\boldsymbol{r}, \boldsymbol{\Omega}) \, dV \, d\boldsymbol{\Omega} \\
& - \sum_{j=1}^{J} \int_{4\pi} \widehat{\psi}_j(\boldsymbol{\Omega}) \int_{s_j} (\boldsymbol{\Omega} \cdot \boldsymbol{n}) \psi_j(\boldsymbol{r}_s, \boldsymbol{\Omega}) \, dS \, d\boldsymbol{\Omega} \\
& - \sum_{j=1}^{J} \sigma_{t,j} \int_{4\pi} \widehat{\psi}_j(\boldsymbol{\Omega}) \int_{V_j} \psi_j(\boldsymbol{r}, \boldsymbol{\Omega}) \, dV \, d\boldsymbol{\Omega} \\
& + \frac{1}{4\pi} \sum_{j=1}^{J} \int_{4\pi} \widehat{\psi}_j(\boldsymbol{\Omega}) \int_{V_j} \left[\int_{4\pi} \sigma_{s,j} \left(\boldsymbol{\Omega}' \rightarrow \boldsymbol{\Omega} \right) \psi_j(\boldsymbol{r}, \boldsymbol{\Omega}') d\boldsymbol{\Omega}' \right] \, dV \, d\boldsymbol{\Omega} \\
& + \sum_{j=1}^{J} \int_{4\pi} \widehat{\psi}_j(\boldsymbol{\Omega}) \int_{V_j} \frac{Q_j}{4\pi} \, dV \, d\boldsymbol{\Omega}.
\end{aligned} \tag{5.9}
$$

where $\boldsymbol{n}$ is the outward unit normal of the surface of the cell.

5.2 VVR functional for a S_2 implementation

In the S_2 implementation, the angular dependence is simplified by dividing the angular domain into 8 octants and the VVR functional can be written as,

$$
\begin{aligned}
F[\widehat{\psi}_j, \psi_j] = & \frac{1}{4\pi} \sum_{j=1}^{J} \sum_{q=1}^{8} \int\limits_{\boldsymbol{\Omega}\in q} \int\limits_{V_j} \Sigma \psi_j(\boldsymbol{r}, \boldsymbol{\Omega})\, dV\, d\boldsymbol{\Omega} \\
& - \sum_{j=1}^{J} \sum_{q=1}^{8} \widehat{\psi}_{j,q} \int\limits_{\boldsymbol{\Omega}\in q} \int\limits_{S} (\boldsymbol{\Omega}\cdot\boldsymbol{n}) \psi_j(\boldsymbol{r}_s, \boldsymbol{\Omega})\, dS\, d\boldsymbol{\Omega} \\
& - \sum_{j=1}^{J} \sigma_{t,j} \sum_{q=1}^{8} \widehat{\psi}_{j,q} \int\limits_{\boldsymbol{\Omega}\in q} \int\limits_{V_j} \psi_j(\boldsymbol{r}, \boldsymbol{\Omega})\, dV\, d\boldsymbol{\Omega} \\
& + \frac{1}{4\pi} \sum_{j=1}^{J} \sum_{q=1}^{8} \widehat{\psi}_{j,q} \int\limits_{\boldsymbol{\Omega}\in q} \int\limits_{V_j} \left[\sum_{q'=1}^{8} \int\limits_{\boldsymbol{\Omega}\in q'} \sigma_{s,j}\left(\boldsymbol{\Omega}' \to \boldsymbol{\Omega}\right) \psi_j(\boldsymbol{r}, \boldsymbol{\Omega}') d\boldsymbol{\Omega}' \right] dV d\boldsymbol{\Omega} \\
& + \sum_{j=1}^{J} \sum_{q=1}^{8} \int\limits_{\boldsymbol{\Omega}\in q} \int\limits_{V_j} \frac{Q_j}{4\pi} \widehat{\psi}_{j,q}\, dV d\boldsymbol{\Omega}.
\end{aligned}
\tag{5.10}
$$

where q is the octant number and,

$$
\widehat{\psi}_{j,q} = \frac{\int\limits_{\boldsymbol{\Omega}\in q} \widehat{\psi}_j(\boldsymbol{\Omega}) d\boldsymbol{\Omega}}{\int\limits_{\boldsymbol{\Omega}\in q} d\boldsymbol{\Omega}} = \frac{\int\limits_{\boldsymbol{\Omega}\in q} \widehat{\psi}_j(\boldsymbol{\Omega}) d\boldsymbol{\Omega}}{\pi/2}. \tag{5.11}
$$

The adjoint angular flux in each octant is reconstructed using an exponential function given by,

$$\psi_j(\mathbf{\Omega}) = a_j e^{\boldsymbol{b}_j \cdot \mathbf{\Omega}}, \tag{5.12}$$

where a_j and $\boldsymbol{b}_j$ are calculated from the expressions for adjoint scalar flux $\widehat{\Phi}_j$ and current $\widehat{\boldsymbol{J}}_j$ given by,

$$\widehat{\Phi}_j = \int_{4\pi} a_j e^{\boldsymbol{b}_j \cdot \mathbf{\Omega}} d\mathbf{\Omega}, \tag{5.13}$$

$$\widehat{\boldsymbol{J}}_j = \int_{4\pi} \mathbf{\Omega}\ a_j e^{\boldsymbol{b}_j \cdot \mathbf{\Omega}} d\mathbf{\Omega}. \tag{5.14}$$

$|\boldsymbol{b}_j|$ is calculated by solving the expression for the ratio $|\widehat{\boldsymbol{J}}_j|/\widehat{\phi}_j$ numerically using a secant solver,

$$\frac{|\widehat{\boldsymbol{J}}_j|}{\widehat{\phi}_j} = \frac{e^{|\boldsymbol{b}_j|}\left(1 - \frac{1}{|\boldsymbol{b}_j|}\right) + e^{-|\boldsymbol{b}_j|}\left(1 + \frac{1}{|\boldsymbol{b}_j|}\right)}{e^{|\boldsymbol{b}_j|} - e^{-|\boldsymbol{b}_j|}} \tag{5.15}$$

and then $\boldsymbol{b}_j$ is determined using the formula,

$$\boldsymbol{b}_j = |\boldsymbol{b}_j| \frac{\widehat{\boldsymbol{J}}_j}{|\widehat{\boldsymbol{J}}_j|}. \tag{5.16}$$

a_j is calculated by substituting the value of $\boldsymbol{b}_j$ in Equation (5.13),

$$a_j = 2\pi \frac{\widehat{\Phi}_j |\boldsymbol{b}_j|}{e^{|\boldsymbol{b}_j|} - e^{-|\boldsymbol{b}_j|}}. \tag{5.17}$$

The average adjoint angular flux in each octant q is then determined by substituting the values of a_j and $\boldsymbol{b}_j$ in Equation (5.11). The angular integration is performed numerically using a Lebedev quadrature rule for unit sphere [28].

The forward quantities in the functional are determined from the Monte Carlo calculations using path length tallies for the fluxes and surface crossing tallies for the surface currents in each octant as defined below,

$$\psi_{j,q} = \int_{V_j} \int_{\mathbf{\Omega}\in q} \psi(\boldsymbol{r}, \mathbf{\Omega})\, d\mathbf{\Omega}\, dV_j, \tag{5.18}$$

$$\Gamma_{j,q} = \int_{S_j} \int_{\mathbf{\Omega}\in q} (\mathbf{\Omega}\cdot\boldsymbol{n})\psi(\boldsymbol{r}, \mathbf{\Omega})\, d\mathbf{\Omega}\, dS_j. \tag{5.19}$$

For linearly anisotropic scattering cross sections, the scattering source term in the VVR functional is written as:

$$\frac{1}{4\pi}\sum_{j=1}^{J}\sum_{q=1}^{8}\widehat{\psi}_{j,q}\int\limits_{\mathbf{\Omega}\in q}\int\limits_{V_j}\left[\sum_{q'=1}^{8}\int\limits_{\mathbf{\Omega}\in q'}\left[\sigma_{s,j,0}+3\sigma_{s,j,1}\mathbf{\Omega}'\cdot\mathbf{\Omega}\right]\psi_j(\boldsymbol{r},\mathbf{\Omega}')d\mathbf{\Omega}'\right]dVd\mathbf{\Omega}, \tag{5.20}$$

where $\sigma_{s,j,0}$ and $\sigma_{s,j,1}$ are the 0^{th} moment and 1^{st} moment of the scattering cross-section, respectively. Now, assuming a mean direction in each octant Ω_q such that,

$$\mathbf{\Omega}_q = \frac{\int\limits_{\mathbf{\Omega}\in q}\mathbf{\Omega}d\mathbf{\Omega}}{\int\limits_{\mathbf{\Omega}\in q}d\mathbf{\Omega}}, \tag{5.21}$$

we can write the scattering source term as,

$$\frac{1}{4\pi}\sum_{j=1}^{J}\sum_{q=1}^{8}\widehat{\psi}_{j,q}\left[\sum_{q'=1}^{8}\left[\sigma_{s0,j}+3\sigma_{s1,j}\mathbf{\Omega}_{q'}\cdot\mathbf{\Omega}_q\right]\int\limits_{\mathbf{\Omega}\in q}\int\limits_{V_j}\int\limits_{\mathbf{\Omega}\in q'}\psi_j(\boldsymbol{r},\mathbf{\Omega}')d\mathbf{\Omega}'dVd\mathbf{\Omega}\right]. \tag{5.22}$$

The integration of the forward angular flux in this expression can now be replaced using Equation (5.18) to yield,

$$\frac{1}{4\pi}\sum_{j=1}^{J}\sum_{q=1}^{8}\widehat{\psi}_{j,q}\left[\sum_{q'=1}^{8}\psi_{j,q'}\left[\sigma_{s0,j}+3\sigma_{s1,j}\boldsymbol{\Omega}_{q'}\cdot\boldsymbol{\Omega}_{q}\right]\int\limits_{\boldsymbol{\Omega}\in q}d\boldsymbol{\Omega}\right]. \tag{5.23}$$

$$\frac{1}{8}\sum_{j=1}^{J}\sum_{q=1}^{8}\widehat{\psi}_{j,q}\left[\sum_{q'=1}^{8}\psi_{j,q'}\left[\sigma_{s0,j}+3\sigma_{s1,j}\boldsymbol{\Omega}_{q'}\cdot\boldsymbol{\Omega}_{q}\right]\right]. \tag{5.24}$$

Substituting this expression for the scattering source term and replacing the integral adjoint and forward quantities in Equation (5.10) using the definitions in Equations (5.11), (5.18) and (5.19), the VVR functional can be written as,

$$\begin{aligned}F[\widehat{\psi}_j,\psi_j]=&\frac{\Sigma}{4\pi}\sum_{j=1}^{J}\sum_{q=1}^{8}\psi_{j,q}\\&-\sum_{j=1}^{J}\sum_{q=1}^{8}\widehat{\psi}_{j,q}\Gamma_{j,q}\\&+\sum_{j=1}^{J}\sigma_{t,j}\sum_{q=1}^{8}\widehat{\psi}_{j,q}\psi_{j,q}\\&-\frac{1}{8}\sum_{j=1}^{J}\sum_{q=1}^{8}\widehat{\psi}_{j,q}\left[\sum_{q'=1}^{8}\psi_{j,q'}\left[\sigma_{s0,j}+3\sigma_{s1,j}\boldsymbol{\Omega}_{q'}\cdot\boldsymbol{\Omega}_{q}\right]\right]\\&-\sum_{j=1}^{J}\frac{Q_j}{8}\sum_{q=1}^{8}\widehat{\psi}_{j,q}.\end{aligned} \tag{5.25}$$

5.3 VVR functional for a P_1 implementation

Spherical harmonics expansion of the angular fluxes for 1st polynomial degree is defined as:

$$\psi(\mathbf{\Omega}) = \sum_{n=0}^{1} \sum_{m=-n}^{n} \phi_{nm} Y_{nm}(\mathbf{\Omega}), \tag{5.26}$$

$$\widehat{\psi}(\mathbf{\Omega}) = \sum_{n=0}^{1} \sum_{m=-n}^{n} \widehat{\phi}_{nm} Y_{nm}(\mathbf{\Omega}), \tag{5.27}$$

for the forward and adjoint angular fluxes respectively. We define the spherical harmonic polynomials for a given direction vector $\mathbf{\Omega}$ $(= \Omega_x \hat{i} + \Omega_y \hat{j} + \Omega_z \hat{k})$ as,

$$Y_{00}(\mathbf{\Omega}) = \frac{1}{\sqrt{4\pi}},\ Y_{10}(\mathbf{\Omega}) = \sqrt{\frac{3}{4\pi}}\Omega_x,\ Y_{11}(\mathbf{\Omega}) = \sqrt{\frac{3}{4\pi}}\Omega_y,\ Y_{1-1}(\mathbf{\Omega}) = \sqrt{\frac{3}{4\pi}}\Omega_z. \tag{5.28}$$

The adjoint spherical harmonic moments in cell j are defined as,

$$\widehat{\phi}_{nm,j} = \int_{4\pi} \widehat{\psi}_j(\mathbf{\Omega}) Y_{nm}(\mathbf{\Omega})\ d\mathbf{\Omega}. \tag{5.29}$$

In this implementation, the adjoint flux moments are calculated deterministically using Attila and are imported into the MGMC code. For the forward solution, we define forward flux moments and forward surface current moments as below,

$$\phi_{nm,j} = \int_V \int_{4\pi} \psi(\boldsymbol{r}, \mathbf{\Omega}) Y_{nm}(\mathbf{\Omega})\ d\mathbf{\Omega}\ dV, \tag{5.30}$$

$$\Gamma_{nm,j} = \oint_S \int_{4\pi} (\mathbf{\Omega} \cdot \boldsymbol{n}) \psi(\boldsymbol{r}, \mathbf{\Omega}) Y_{nm}(\mathbf{\Omega})\ d\mathbf{\Omega}\ dS. \tag{5.31}$$

The surface current moments defined here will be useful when applying divergence theorem to the streaming operator in the transport term. Substituting the spherical harmonics expansion in the VVR functional, and using the orthogonal properties of the spherical harmonic moments, a reduced form of the VVR functional can be derived. Here we show only the derivation of the scattering source term with linearly anisotropic scattering and the rest of the terms in the functional are derived in the similar fashion to the previous work done by Densmore for criticality problems [12].

Considering the scattering source term in the general form of the VVR functional presented in Equation (5.9), we have:

$$\frac{1}{4\pi}\sum_{j=1}^{J}\int_{4\pi}\int_{V_j}\widehat{\psi}_j(\boldsymbol{\Omega})\left[\int_{4\pi}\sigma_{s,j}\left(\boldsymbol{\Omega}'\rightarrow\boldsymbol{\Omega}\right)\psi_j(\boldsymbol{r},\boldsymbol{\Omega}')d\boldsymbol{\Omega}'\right]\,dV\,d\boldsymbol{\Omega}, \tag{5.32}$$

Substituting the linearly anisotropic expansion for the scattering cross section, we get,

$$\frac{1}{4\pi}\sum_{j=1}^{J}\int_{4\pi}\int_{V_j}\widehat{\psi}_j(\boldsymbol{\Omega})\left[\int_{4\pi}(\sigma_{s0,j}+3\sigma_{s1,j}\boldsymbol{\Omega}\cdot\boldsymbol{\Omega}')\psi_j(\boldsymbol{r},\boldsymbol{\Omega}')d\boldsymbol{\Omega}'\right]\,dV\,d\boldsymbol{\Omega}, \tag{5.33}$$

Treating the 0^{th} and 1^{st} scattering cross section moments individually, we have for the 0^{th} moment,

$$\frac{1}{4\pi}\sum_{j=1}^{J}\sigma_{s0,j}\int_{4\pi}\int_{V_j}\widehat{\psi}_j(\boldsymbol{\Omega})\left[\int_{4\pi}\psi_j(\boldsymbol{r},\boldsymbol{\Omega}')d\boldsymbol{\Omega}'\right]\,dV\,d\boldsymbol{\Omega}, \tag{5.34}$$

From the definition of Y_{00} in Equation (5.28), we can write, $1 = \sqrt{4\pi}Y_{00}$ and introducing this into the above expression, we have,

$$\frac{1}{4\pi}\sum_{j=1}^{J}\sigma_{s0,j}\int_{4\pi}\int_{V_j}\sqrt{4\pi}Y_{00}\widehat{\psi}_j(\mathbf{\Omega})\left[\int_{4\pi}\psi_j(\boldsymbol{r},\mathbf{\Omega}')d\mathbf{\Omega}'\right]\,dV\,d\mathbf{\Omega}, \tag{5.35}$$

From the definition for adjoint spherical harmonic moments in Equation (5.29), we have

$$\frac{1}{\sqrt{4\pi}}\sum_{j=1}^{J}\sigma_{s0,j}\widehat{\phi}_{00,j}\int_{4\pi}\psi_j(\boldsymbol{r},\mathbf{\Omega}')d\mathbf{\Omega}'\,dV. \tag{5.36}$$

Again introducing $1 = \sqrt{4\pi}Y_{00}$ into the expression and using the definition in Equation (5.30), we get,

$$\sum_{j=1}^{J}\sigma_{s0,j}\widehat{\phi}_{00,j}\phi_{00,j}. \tag{5.37}$$

Now, considering the 1^{st} scattering moment, we have,

$$\frac{3}{4\pi}\sum_{j=1}^{J}\int_{4\pi}\int_{V_j}\widehat{\psi}_j(\mathbf{\Omega})\left[\int_{4\pi}\sigma_{s1,j}(\mathbf{\Omega}\cdot\mathbf{\Omega}')\psi_j(\boldsymbol{r},\mathbf{\Omega}')d\mathbf{\Omega}'\right]\,dV\,d\mathbf{\Omega}, \tag{5.38}$$

In the above expression, the constant $\frac{3}{4\pi}$ and the scalar product $\mathbf{\Omega}\cdot\mathbf{\Omega}'$ can be combined and written in terms of the spherical harmonics polynomials as,

$$\frac{3}{4\pi}(\mathbf{\Omega}\cdot\mathbf{\Omega}') = Y_{10}(\mathbf{\Omega})Y_{10}(\mathbf{\Omega}') + Y_{11}(\mathbf{\Omega})Y_{11}(\mathbf{\Omega}') + Y_{1,-1}(\mathbf{\Omega})Y_{1,-1}(\mathbf{\Omega}'), \tag{5.39}$$

using the definitions in Equation (5.28). Using this expression and the orthogonal property of spherical harmonic moments defined by,

$$\int_{4\pi} Y_{nm}(\mathbf{\Omega}) Y_{n'm'}(\mathbf{\Omega}) d\mathbf{\Omega} = 1, \text{if } n = n', m = m' \tag{5.40}$$

$$= 0, \text{elsewhere},$$

in Equation (5.38) and some algebra, the terms for the first moment of the scattering source is given by,

$$\sum_{j=1}^{J} \sigma_{s1,j} \sum_{-1}^{1} \widehat{\phi}_{1m,j} \phi_{1m,j}. \tag{5.41}$$

Similarly evaluating the other terms in the VVR functional, we get,

$$\begin{aligned} F[\widehat{\psi}_j, \psi_j] &= \frac{1}{4\pi} \sum_{j=1}^{J} \Sigma_j \phi_j \\ &- \sum_{j=1}^{J} \left[\sum_{n=0}^{1} \sum_{m=-n}^{n} \widehat{\phi}_{j,nm} \Gamma_{j,nm} + \sigma_{t,j} \sum_{n=0}^{1} \sum_{m=-n}^{n} \widehat{\phi}_{j,nm} \phi_{j,nm} - \sum_{n=0}^{1} \sigma_{s,n,j} \sum_{m=-n}^{n} \widehat{\phi}_{j,nm} \phi_{j,nm} \right] \\ &+ \frac{1}{4\pi} \sum_{j=1}^{J} Q_j \widehat{\phi}_j V_j. \end{aligned} \tag{5.42}$$

In the above functional, the quantities ϕ_j and $\widehat{\phi}_j$ are defined by,

$$\phi_j = \int_{4\pi} \sum_{n=0}^{1} \sum_{m=-n}^{n} \phi_{j,nm} Y_{nm}(\mathbf{\Omega})\, d\mathbf{\Omega} = \phi_{00}\sqrt{4\pi}, \tag{5.43}$$

$$\widehat{\phi}_j = \int_{4\pi} \sum_{n=0}^{1} \sum_{m=-n}^{n} \phi_{j,nm} Y_{nm}(\mathbf{\Omega})\, d\mathbf{\Omega} = \widehat{\phi}_{00}\sqrt{4\pi}. \tag{5.44}$$

5.4 Computational strategy

As discussed in the previous chapter, the VVR method is a unique variance reduction technique that does not modify the particle transport. But, to evaluate the variational functionals, the solution quantities need to be tallied in all regions of the problem domain. In Barrett's implementation (early 90's)[7], the memory write times for tallying quantities in all regions of the problem accounted for a significant part of the additional computational efforts necessary to implement the VVR method. However, the memory write times do not pose significant performance bottlenecks in our present day computers. But, the time to evaluate the VVR functional could still be significant relative to the overall simulation time and unless the variance reduction provided by employing the VVR method is considerable, the overall efficiency of the simulation using VVR may not be better. Two important implementation details to efficiently calculate the VVR functionals are presented below.

5.4.1 Batch means estimator

Usually, in Monte Carlo simulations for source-detector type problems the contributions made to the detector tallies are stored along with their squared moments and the mean estimate and its variance are calculated at the end of the simulation. For the VVR method however, the VVR estimates of the detector responses cannot be evaluated for every particle history as it would result in calculating the VVR functional in each spatial cell and energy group in the problem which would be computationally very expensive. To avoid this problem, the quantities are tallied in batches, similar to criticality problems, such that the additional computational time for evaluating the VVR functional remains reasonable. But, the number of particles simulated in each batch (batch size) and the number of batches required for a reasonable estimate is problem dependent. Also, if the number of batches used is very low, the application of Central Limit Theorem could lead to potential problems. Further, the variance calculated for the batch means estimator could have some bias due to correlation effects between the batches for small batch sizes [29]. The optimal batch size for each problem is determined based on trial and error such that the bias introduced in the variance estimate by using the batch means estimator is not significant and does not influence the FOM calculations.

5.4.2 Scattering source transformation

The VVR functional we derived in the previous sections did not include energy dependence. If we include energy dependence, the scattering source term in the VVR functional becomes computationally very expensive to calculate in its current form. This is because for every energy group, the scattering source contribution from all energy groups in the problem needs to be calculated every time the functional is evaluated. But, since the adjoint solution is pre-calculated and is essentially constant during the Monte Carlo calculation, we can take advantage of the properties of adjoint theory and compute the scattering source in each energy group only once. Lets consider the final form of the VVR functional for the P_1 implementation with energy dependence,

$$\begin{aligned} F[\widehat{\psi}_j, \psi_j] = & \frac{1}{4\pi} \sum_{j=1}^{J} \sum_{g=1}^{G} \Sigma_{j,g} \phi_{j,g} \\ & - \sum_{j=1}^{J} \sum_{g=1}^{G} \left[\sum_{n=0}^{1} \sum_{m=-n}^{n} \widehat{\phi}_{j,g,nm} \Gamma_{j,g,nm} + \sigma_{t,g,j} \sum_{n=0}^{1} \sum_{m=-n}^{n} \widehat{\phi}_{j,g,nm} \phi_{j,g,nm} \right] \\ & + \sum_{j=1}^{J} \sum_{n=0}^{1} \sum_{m=-n}^{n} \sum_{g=1}^{G} \widehat{\phi}_{j,g,nm} \sum_{g'=1}^{G} \sigma_{j,s,g' \to g,n} \phi_{j,g',nm} \\ & + \frac{1}{4\pi} \sum_{j=1}^{J} \sum_{g=1}^{G} Q_{j,g} \widehat{\phi}_{j,g} V_j. \end{aligned} \tag{5.45}$$

In the above functional subscript g indicates the energy group that is evaluated and g' represents the energy group from which particles are scattering. G is the total number of energy groups in the cross section library. In this form of the functional,

the scattering source of the forward problem is being weighted by the adjoint flux moments. The speed of evaluating this functional can be expressed as $\mathcal{O}(J \times G^2 \times (n+1)^2)$ where J is the number of cells, G is the number of groups and n is the P_N order. Now, we know from adjoint theory,

$$\langle \sigma_s \psi, \widehat{\psi} \rangle = \langle \psi, \widehat{\sigma}_s \widehat{\psi} \rangle, \tag{5.46}$$

where $\langle\ ,\ \rangle$ is the notation for inner product over the problem domain and $\widehat{\sigma}_s$, the adjoint scattering matrix is the transpose of the forward scattering cross section matrix σ_s. Using this relationship, we can rewrite the scattering source term in the functional as ,

$$\sum_{j=1}^{J} \sum_{n=0}^{1} \sum_{m=-n}^{n} \sum_{g=1}^{G} \widehat{\phi}_{j,g,nm} \sum_{g'=1}^{G} \sigma_{j,s,g\to g',n} \phi_{j,g,nm} = \sum_{j=1}^{J} \sum_{n=0}^{1} \sum_{m=-n}^{n} \sum_{g'=1}^{G} \phi_{j,g',nm} \sum_{g=1}^{G} \sigma_{j,s,g'\to g,n} \widehat{\phi}_{j,g,nm} \tag{5.47}$$

Since the adjoint flux is a constant, the scattering source contribution can be evaluated during the first evaluation of the VVR functional and for remaining evaluations, the scattering source just needs to be weighted by the new estimate of the forward flux moments calculated in Monte Carlo. The speed for evaluating the functionals after the first evaluation will be $\mathcal{O}(J \times G \times (n+1)^2)$. A similar transformation of the scattering source term is done for the S_2 implementation.

Chapter 6: Test Problems and Results

6.1 Mulch photon test problem

To test the hybrid methods on realistic problems, the "Mulch Box" benchmark test problem for photon transport is considered. The problem definition is taken from the PNNL report containing a list of benchmark problems for non-proliferation scenarios. It is a simple photon transport problem through a shield with total macroscopic cross-sections that are in the range between 0.5 and 1.0 cm^{-1} [30]. The geometry consists of a wooden crate containing the mulch material on a slab of concrete placed over soil. For this problem we use a "low density" mulch material (0.266 g/cc) consisting of C (48 wt%), H (6 wt%) and O (46 wt%). A HPGe detector modeled as a solid cylindrical volume ($62cm^3$) is placed to the east of the crate while a spherical Ba-133 photon source (65 cm^3) is placed to the west as shown in the Figure 6.1. The environment is modeled as nitrogen. A 50-group cross-section library generated from the CEPXS cross section processing utility [31] is used. The group structure of the cross section library and the source spectrum for Ba-133 are tabulated in appendices B and C, respectively. The adjoint problem is solved using a S_6, Triangular Chebyshev-Legendre quadrature set and linearly anisotropic scattering. DSA (Diffusion Synthetic Acceleration) is turned off in the Attila solver because the problem's geometry has

high aspect ratios and multiple heterogeneous material regions. Attila's standard DSA option does not provide significant acceleration benefits for problems with these characteristics [16]. A lumped finite element discretization is used for the solution of the adjoint transport equation to attempt to minimize the presence of negative adjoint angular fluxes. The calculations are performed on a single mesh of global size 0.3 cm, the same as that used in the PNNL report for the Attila calculation [30].

Figure 6.1: CAD model for mulch problem

Four variations of the Monte Carlo code are used to solve the forward calculation - (1) A standard Monte Carlo implementation with implicit absorption in which the particle weights are adjusted to account for absorption reactions (Analog+IA); (2) a hybrid Monte Carlo calculation with weight windows and source biasing (SB+WW); (3) a hybrid Monte Carlo calculation with the LIFT method (LIFT) and (4) a hybrid Monte Carlo method with LIFT along with the weight windows technique (LIFT+WW). Since implicit absorption is implemented by default in almost all production Monte

Carlo codes [2], calculations with this variance reduction technique are referred to as "analog" in the remainder of the dissertation. The VVR method is applied in five different configurations. The P_1 implementation of VVR is applied with the analog calculation (VVR P_1), source biasing with weight windows calculation (VVR P_1+SB+WW) and the LIFT calculation (VVR P_1+LIFT). The S_2 implementation of VVR is applied with the weight windows calculation (VVR S_2+WW) and LIFT with weight windows calculation (VVR S_2+LIFT+WW). The motivation for not using S_2 implementation with the analog calculation is explained in the next chapter.

Two different variance reduction goals are tested to evaluate the performance of the hybrid methods. In the first case, the goal is to optimize the detector flux across all energy groups by driving the adjoint calculation with a group-wise unit volume source placed in the HPGe detector. In this case, all the energy groups are given equal importance resulting in groups with larger particle fluxes having smaller variances than groups with smaller particle fluxes. For the second case, the detector absorption rate in the lower energy groups, those with the most significant absorption cross sections, is optimized. The reason for optimizing the reaction rates in specific energy groups is to test the performance of the variance reduction techniques when optimizing a specific energy and spatial tally as the detector efficiencies are often times a strong function of energy. This is achieved by setting the adjoint source equal to the absorption cross section in groups 46 to 49. Group 50 is neglected in the calculation as the absorption cross section of the detector is orders of magnitude greater than other groups and having an adjoint source in group 50 might mask the importance of the

groups above it. The efficiency of each method is quantified using the figure of merit (FOM) defined as,

$$FOM = \frac{1}{r^2 t}, \tag{6.1}$$

where r is the relative error and t is the CPU time (in seconds) for the calculation. For the hybrid methods, the time taken for the adjoint deterministic calculations is also included (180 seconds).

	Detector Flux	**Rel. Error**	**Time**	**FOM**
Analog+IA	4.01E-05	0.117	2472	0.033
VVR P_1	4.53E-05	0.097	2662	0.040
SB+WW	4.81E-05	0.031	824	1.263
VVR P_1+WW	5.04E-05	0.021	834	2.645
VVR S_2+WW	4.73E-05	0.023	928	2.064
LIFT	5.15E-05	0.005	998	49.482
LIFT+WW	5.17E-05	0.005	955	41.885
VVR P_1+LIFT	1.56E-05	0.013	1008	5.933
VVR S_2+LIFT+WW	3.18E-05	0.009	1041	11.600

Table 6.1: Results for optimizing the average detector flux in all groups

The results reported are for 100 batches with a batch size of 10^5 particles. The results for the case tallying the detector fluxes are shown in Table 6.1. The flux values for the analog calculation and the SB+WW method are within the 95% confidence interval of each other. The upper bound of the 95% confidence interval for the detector flux calculated by the Analog+IA calculation is approximately 3% lower than the lower bound of the 95% confidence interval calculated by the LIFT and LIFT+WW methods. This difference in the LIFT results is likely due to the nature of the exponential distribution in the LIFT expression which becomes highly peaked when the

magnitude of the biasing parameter is high or the cell size is large. A more detailed discussion on this result is included in the next chapter.

The computational time required by the hybrid methods to achieve similar statistical precision is significantly lower than the Analog+IA calculation. This results in FOM that are at least two orders of magnitude larger for the hybrid methods when compared to the analog calculation. The LIFT method has a FOM that is almost fifty times greater than the SB+WW method. LIFT also performed slightly better than the LIFT+WW method which indicates that the effect of weight windows when combined with LIFT is actually detrimental in this test problem. Since the adjoint problem is driven by an unit source in all groups, the deterministic solution quality is sufficiently accurate to bias the particles towards the detector region efficiently by the LIFT implementation alone and weight windows is actually an added computational burden.

The VVR method yields a marginal improvement in FOM over the analog estimate (0.03 vs 0.04). When applied with the SB+WW method, the VVR method yields an improvement in FOM that is twice the direct estimate for the P_1 implementation and about 1.6 times the FOM for the S_2 implementation of the functional. For the LIFT method and LIFT+WW method, the VVR estimate did not lead to any improvement in FOM. Also, the detector fluxes do not agree with the results from the other simulations. The reason for this is that the error correction term added to the detector response (in this case particle flux) is large when compared to the actual detector flux estimate. The LIFT method biases the particle transport greatly

that the errors in the solution in regions with lower importances (i.e regions from which the particles are more likely to reach the detector) are large. Since the VVR functional involves the global solution, the VVR estimate from the LIFT simulations is dominated by these imprecise correction terms and fails to provide any acceleration.

	Reaction Rates	**Rel. Error**	**Time**	**FOM**
Analog+IA	6.00E-05	0.0870	53197	0.002
VVR P_1	6.51E-05	0.287	2669	0.005
SB+WW	6.03E-05	0.01	1775	5.634
VVR P_1+WW	6.22E-05	0.009	1787	6.385
VVR S_2+WW	6.20E-05	0.009	1877	6.16
LIFT	6.00E-05	0.012	1586	4.168
LIFT+WW	6.10E-05	0.004	1824	40.047
VVR P_1+LIFT	-8.80E-05	0.075	1597	-
VVR S_2+LIFT+WW	9.80E-06	0.037	1918	-

Table 6.2: Results for optimizing the total detector reaction rates in groups 46-49

For the second case, the total detector reaction rate in the last four groups (46-49), where the absorption cross sections are large, is optimized and the results are shown in Table 6.2. This problem is optically thick as the cross-sections in the lower groups are larger and using the same particle histories as case 1 resulted in large errors for the analog calculation. Therefore, a longer simulation is performed for the Analog+IA method with 2×10^8 particle histories. The value of the total detector reaction rate in groups 46-49 for all the calculations lie within the 95% confidence interval of each other. The slight difference in the LIFT results that is observed in first case is not present in this calculation beacause the cross sections are much larger than the particle

fluxes in the lower energy groups of the detector. The relative error and the FOM are better for the SB+WW method than the LIFT method, while the LIFT+WW method has a FOM nine times (approx.) that of the SB+WW method. The reason for the relatively poor performance of the LIFT method in this case is the presence of higher weight fluctuations in the contributions made to the tally. The number of particles reaching the detector for the LIFT calculation is about a factor of five greater than the SB+WW calculation, as expected. But, the contributions made to the tally have a higher variance than the SB+WW calculation. The higher variance is caused by a small number of particles that contribute to the tally by traveling through phase space regions that have lower importances and hence a higher weight than the majority of the particles. Ideally, this situtation should not be possible as any particle track that leads to the detector will have higher importance. But because of the limitations in the LIFT approximate angular flux expression, these unlikely events do infrequently occur. For the LIFT+WW method, the weight fluctuations are suppressed to a large extent by weight windows and this results in improved efficiency.

VVR yields an improvement in FOM twice that of the direct analog estimate. When applied with the SB+WW method however, the improvement in FOM is very marginal ($\approx$ 10%) for both the P_1 and S_2 implementations. VVR with LIFT simulations demonstrate a similar behavior to case 1 and did not provide any acceleration solution convergence. Further, the VVR estimate becomes negative for the LIFT with VVR (P_1) estimate possibly due to the global errors in LIFT calculation being larger than the estimate itself and therefore we have not reported the FOM for the VVR+LIFT

calculations. Also, the estimations for the simulations with and without VVR are calculated in the same simulation and the additional computational time for evaluating the VVR functional is recorded and used for calculating the FOMs in the VVR calculations.

To test if the performance of the hybrid methods could be improved by employing a coarser mesh for the forward Monte Carlo calculation, a coarser mesh of global size 0.8 cm is used. The adjoint solution from the 0.3 cm fine mesh is transferred to the coarser mesh using the method described in the previous section. The computational cost of this transfer is negligible ($\approx$ 4 CPU seconds). The results for optimizing the

	Det. Flux	**Rel. Error**	**Time**	**FOM**
Analog+IA	4.02E-05	0.117	2447	0.030
VVR P_1	4.53E-05	0.097	2630	0.040
SB+WW	5.03E-05	0.068	768	0.282
VVR P_1+WW	5.36E-05	0.064	758	0.322
VVR S_2+WW	4.96E-05	0.068	768	0.278
LIFT	5.19E-05	0.007	910	22.427
LIFT+WW	5.27E-05	0.061	755	0.356
VVR P_1+LIFT	-1.79E-05	0.016	911	4.288
VVR S_2+LIFT+WW	1.57E-05	0.190	767	0.036

Table 6.3: Results for optimizing the average detector flux on a coarse mesh

total detector flux are shown in Table 6.3. The flux values for each of the four methods showed behavior similar to that of the fine mesh problems. When using a coarser mesh for optimizing the total detector flux, the analog calculation has approximately the same execution time as that of the fine mesh calculation. This shows that the size of the mesh does not have a significant impact on the performance of our analog Monte

Carlo calculation. For the hybrid methods, although there is a noticeable reduction in execution times, using a coarser mesh did not cause an increase in the FOMs. In fact, the coarse mesh hybrid calculations resulted in smaller FOMs than the fine mesh calculations. For the LIFT and SB+WW method, the FOM decreased by about 50%. But, for the LIFT+WW method, the FOM dropped below that of both SB+WW and LIFT method. This is because of the reduced precision in the importance map (adjoint solution), as the number of cells in the coarse mesh is approximately eight times fewer than the fine mesh. This causes the particles biased by the LIFT method traveling towards the detector with lower weights to be terminated before reaching the detector. This reduced precision is expected due to the approximation of the adjoint solution on a coarser mesh. But, the reduction in execution time must overcome the loss in precision to justify the use of a coarser mesh for the forward calculation. Although the computational cost of calculating the weights and the biasing parameters every time the particle crosses a spatial cell decreases using a coarse mesh, in a problem with significant scattering between the energy groups, the weights and the biasing parameters are recalculated every time the particle scatters from group to group. Even on a coarse mesh, the reduction in execution times may not be significant in this situation. Table 6.4 shows the results for the case in which the reaction rates in the lower groups (46-49) are optimized. The FOMs are again smaller on the coarse mesh than on the fine mesh. The execution times for the hybrid methods are smaller than those for the corresponding fine mesh calculations, but not significantly. The time to evaluate the VVR functionals in the coarse mesh for all the calculations is smaller by ≈ 100 seconds. The relative errors are slightly larger in the coarse mesh

	Reac. Rate	Rel. Error	Time	FOM
Analog+IA	6.14E-05	0.400	2470	0.003
VVR P_1	6.39E-05	0.290	2652	0.004
SB+WW	5.94E-05	0.016	1562	2.631
VVR P_1+WW	6.01E-05	0.015	1564	2.728
VVR S_2+WW	5.95E-05	0.015	1574	2.679
LIFT	6.01E-05	0.030	1485	0.748
LIFT+WW	6.11E-05	0.012	1267	5.481
VVR P_1+LIFT	-5.70E-05	0.024	1487	-
VVR S_2+LIFT+WW	-3.11E-05	0.074	1279	-

Table 6.4: Results for optimizing the total reaction rates in groups 46-49 on a coarser mesh

calculation resulting in smaller FOMs. Finally, although the FOMs for the hybrid methods are smaller in the coarse mesh calculations when compared to the fine mesh calculations, they are still significantly larger than those for the corresponding analog Monte Carlo calculations.

To study the effect of the quadrature set used for the S_N adjoint deterministic calculation on the performance of the variance reduction techniques, a quadrature mesh refinement study is performed. The first case for optimizing the detector flux on a geometric mesh of global size 0.3 cm is simulated using adjoint solutions generated from S_6, S_8 and S_{10} calculations in Attila. Further, the First Scattered Distributed Source (FSDS) solution technique is also utilized to generate highly accurate adjoint solutions. In Attila, the FSDS option breaks the solution into uncollided and collided components which are then added to yield the final solution [16]. The uncollided part of the solution is calculated using a higher order S_{18} calculation followed by a S_8 calculation to determine the collided part of the solution. The FOMs for the SB+WW

Quadrature	SB+WW	LIFT	VVR
$\boldsymbol{S}_6$	1.309	54.809	0.040
$\boldsymbol{S}_8$	1.458	52.162	0.038
$\boldsymbol{S}_{10}$	1.358	43.956	0.041
FSDS,$\boldsymbol{S}_{18}$-$\boldsymbol{S}_8$	1.375	38.388	0.040

Table 6.5: FOMs for adjoint solutions from higher order S_N and FSDS calculations

method, the LIFT method and the VVR (P_1) method when using the adjoint solutions calculated for different quadrature sets are given in Table 6.5. The Attila run times, although insignificant when compared to the Monte Carlo calculation times, are also included to calculate the FOMs. The results showed that the FOMs did not improve by using a higher quadrature refinement for the adjoint calculations. Further, when using the adjoint solution from the FSDS calculation there is a noticeable drop in the efficiency of the LIFT method. Figure 6.2 shows the adjoint flux maps for

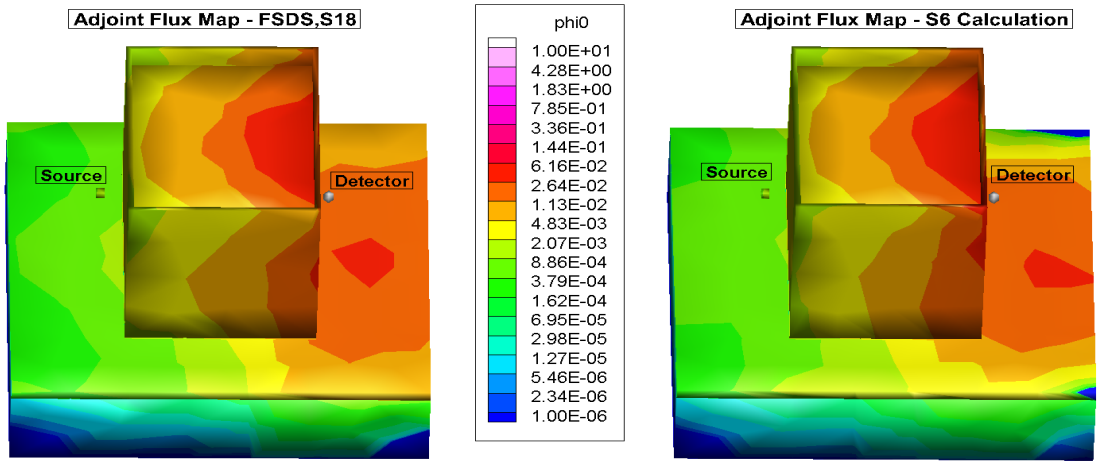

Figure 6.2: Adjoint flux for the S_6 calculation and the FSDS calculation with quadrature order 18 for the uncollided solution and 8 for the collided solution.

the S_6 calculation and the FSDS calculation. The only noticeable differences in the

adjoint flux values are observed on the edges of the concrete slab while the regions that are important for the particles to reach the detector such as the mulch region had similar flux values for both calculations. Therefore, for this mulch problem, the S_6 adjoint calculation is enough for the variance reduction methods to be efficient.

6.2 PANDA neutron test problem

The PANDA neutron problem (14.1.1) described in the PNNL report for benchmark nonproliferation test problems [30] is simulated. This is a highly scattering problem with a neutron source surrounded by a cylindrical polyethylene moderator region housing multiple Helium-3 detectors. This problem is relatively easier to solve using traditional Monte Carlo or deterministic techniques and therefore testing the hybrid variance reduction methods on this problem will provide the opportunity to study if there are still performance benefits to be had or if additional computational efforts associated with the hybrid methods are detrimental to the overall efficiency of the calculations. Two configurations of the PANDA neutron problem are simulated. The first case had 4 detectors and the second case had 8 detectors. The cylindrical moderator and detector geometries are approximated using a hexagonal geometry such that the surface areas and volumes are preserved. The neutron source is a unit volume source placed in the center of the problem geometry and the energy distribution comes from a Cf-252 spectrum. The space between the source and the moderator is modeled as nitrogen. The SCALE44 group neutron cross section library

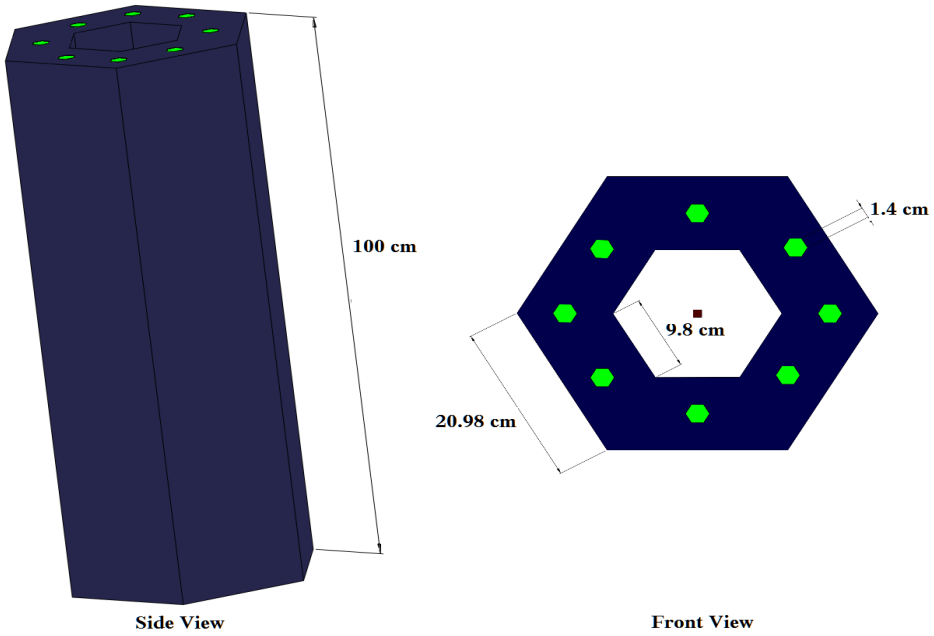

Figure 6.3: Neutron PANDA problem (CAD model)

[32] for thermal reactors in AMPX [33] format is used.

The adjoint calculation is performed with S_4 Triangle Chebychev Labotto quadrature rule and the maximum scattering order was limited to 1. The convergence criteria for the scattering source was set to 1e-02 and DSA acceleration was turned on. For this problem, using a higher order quadrature set did not yield any improvements in variance reduction as expected and only resulted in increased computational time for generating the adjoint solution. A global mesh size of 4 cm is used for both the adjoint and forward problems. The adjoint source had a strength of unity in each energy group and is placed in all the detectors in the problem such that the total average detector flux is optimized using the hybrid methods. The forward Monte Carlo simulation is performed with 10^5 particle histories and the estimates are tallied

in batches with a batch size of 2000 particles.

	Detector Flux	**Rel. Error**	**Time**	**FOM**
Analog+IA	2.75E-04	0.0071	12	1662
VVR P_1	2.75E-04	0.0067	379	59
SB+WW	2.75E-04	0.0074	368	50
VVR P_1+ WW	2.75E-04	0.0070	381	54
VVR S_2+WW	2.74E-04	0.0070	692	30
LIFT	2.75E-04	0.0052	374	99
LIFT+WW	2.71E-04	0.0070	377	54
VVR P_1+LIFT	2.62E-04	0.0050	399	100
VVR S_8+WW+LIFT	2.64E-04	0.0070	707	29

Table 6.6: Average detector flux for case 1 with 4 detectors

Table 6.6 shows the results for case 1 of the PANDA problem. As expected for this problem the errors in the detector flux estimate are very low even for fewer particle histories due to the nature of the problem. The FOM for the analog calculation is two orders of magnitude higher than for the hybrid methods; however, the reason for the poor performance of the variance reduction techniques is due to the time taken for calculating the adjoint solution in Attila. For a highly scattering problem with multiple detectors (which act as source locations to the adjoint calculation), the time taken for the scattering source to converge in the deterministic calculation even with a bigger convergence criteria and S_4 quadrature order is significant (≈ 350 seconds) in comparison to the time taken for the forward Monte Carlo simulations. The forward problem is relatively much easier to solve due to a single unit source location in the center of the problem geometry with optically thin material between the detector and the source. The computational time for the Monte Carlo calculations, even with

variance reduction techniques implemented, are relatively short ($\approx$ 12 seconds for analog and $\approx$ 35 seconds for hybrid methods except for calculations with VVR S_2 method which took $\approx$ 300 seconds). The time to evaluate the VVR functionals in the S_2 implementation is significantly longer than the P_1 implementation due to particles reaching all the angular bins of the phase space in every batch. The error in the LIFT method is slightly smaller than the analog calculation (0.5% for LIFT and 0.7% for analog) while the SB+WW method and the LIFT+WW method had identical error to the analog calculation. The VVR method provided slight improvements over the direct estimates for the analog and SB+WW calculations. Also, unlike the mulch problem where the VVR estimate for the LIFT calculations are unphysical, the VVR+LIFT estimates fell within the 95% confidence interval of the LIFT results, but did not provide any noticeable improvement in error. For the PANDA problem, the errors in the solution are smaller across all regions of the phase space and therefore the errors in the unimportant regions of the problem do not dominate the LIFT estimate as they did in the mulch problem.

	Detector Flux	**Rel. Error**	**Time**	**FOM**
Analog+IA	2.63E-04	0.0059	17	1690
VVR P_1	2.63E-04	0.0057	395	78
SB+WW	2.61E-04	0.0059	365	80
VVR P_1+WW	2.61E-04	0.0048	414	107
VVR S_2+WW	2.61E-04	0.0047	1042	44
LIFT	2.61E-04	0.0042	389	147
LIFT+WW	2.60E-04	0.0035	401	204
VVR P_1+LIFT	2.54E-04	0.0038	438	160
VVR S_2+WW+LIFT	2.53E-04	0.0038	1054	65

Table 6.7: Average detector flux for case 2 with 8 detectors

Table 6.6 shows the results for case 2 of the PANDA problem. The performance of the methods is similar to that observed in case 1. The errors in the flux estimates calculated by the LIFT calculations improved slightly over the analog calculation (0.6% error for analog and $\approx$ 0.4% error for the LIFT methods). But the time required to calculate the adjoint solution deterministically was large, such that the overall performance did not improve.

Figures 6.4, 6.5 and 6.6 show the reduction in statistical error as a function of the particle histories in batches with and without VVR for the analog, weight windows and LIFT simulations for case 1 of the PANDA problem. The simulation results using a higher quality adjoint solution from a S_8 deterministic calculation is used in these plots. There was a very small improvement in the VVR estimates by using a higher quality adjoint, but the time for calculating the adjoint solution was very expensive ($\approx$ 700 seconds) and as a result the FOM did not improve. The error reduction in the VVR estimate over the analog estimate is significant at the beginning of the simulation and after a few batches the error of the analog calculation settles to a value slightly above the VVR error. However, since the reduction in error is inversely proportional to the square of the particle histories, the VVR method continues to provide better error reduction than the direct analog estimate. VVR combined with the weight windows calculations also shows a significant reduction in error compared to the direct estimate in the initial batches and smaller fluctuations than the direct estimates. Also, the S_2 VVR implementation provides better error reduction than the P_1 method which is expected due to the finer angular resolution it provides for

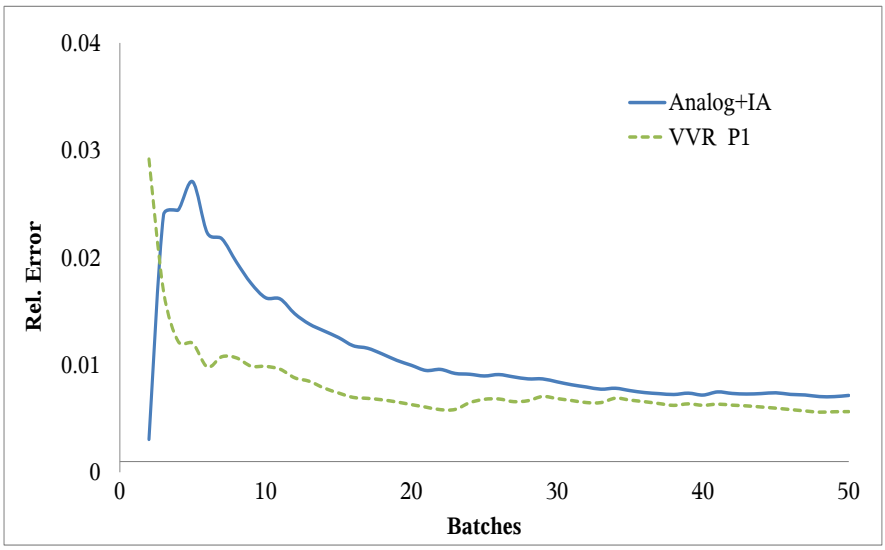

Figure 6.4: Reduction in statistical error for analog Monte Carlo and with VVR.

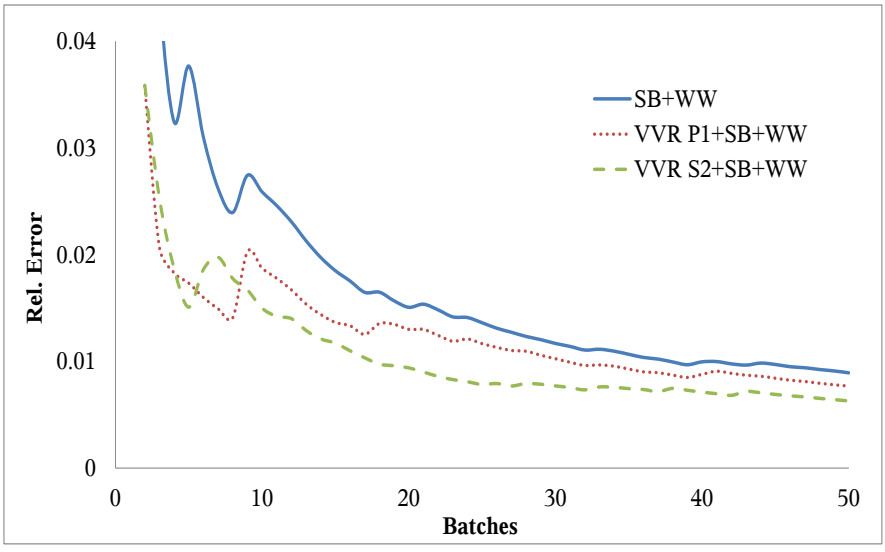

Figure 6.5: Reduction in statistical error for source biasing with weight windows and with VVR.

this highly scattering problem. VVR combined with LIFT calculations did not show

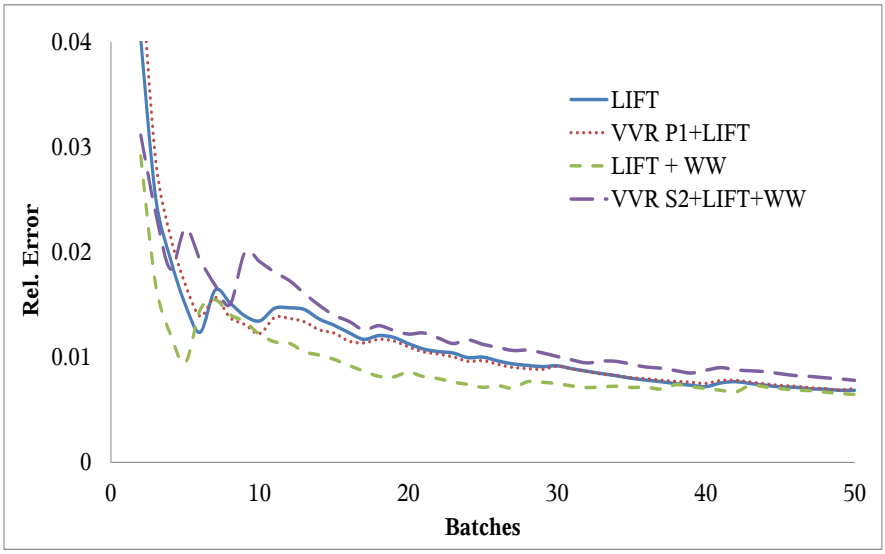

Figure 6.6: Reduction in statistical error for LIFT, LIFT with weight windows and VVR.

a significant reduction in error over the direct flux estimates, as is evident in the overlapping curves in figure 6.6. The LIFT+WW estimate had the smallest error indicating that the VVR method fails to provide any benefits when coupled with the LIFT method.

6.3 UF_6 spent fuel cask test problem

The UF_6 spent fuel cask test problem represents a more realistic test problem for which experimental results exist [30]. In this research, we have only tested the per-

formance of the hybrid methods using the photon spectrum for the decay products of U235-U238. For a complete benchmarking of the hybrid methods with experimental results, the photon spectrum from the fission products of other fissionable elements like Pu-239 and the contribution from neutron spectrum of the decay products should be simulated. Nevertheless, this initial testing using the U235-U238 spectrum helped to gain further insights on the performance and limitations of the hybrid methods.

A 200 group photon cross section library generated using CEPXS [31] is used. Figures 6.7 and 6.8 show the top view and side view of the problem. A more detailed description of the geometry can be found in the PNNL report [30]. In the CAD model, the cylindrical detector geometry is approximated using a hexagonal geometry while still preserving the surface areas and volume so that number of cells in the geometric mesh is reasonable. The adjoint calculation is driven from the NaI sensor of the middle detector with a unit energy spectrum for all energy groups. The adjoint solution generated will help to optimize the average flux in the NaI sensor of the detector. The UF_6 in the cask is modeled as a solid block in the lower half of the cylinder while the upper half of the cylinder is modeled as Nitrogen. A geometric mesh with a global size of 0.1 m is used. The detector and UF_6 cylinder had finer mesh regions. The deterministic solution is generated by performing a FSDS calculation in Attila with a S_{18} quadrature set for the uncollided flux and a S_8 quadrature set for the collided flux. The Triangular Chebychev Legendre quadrature rule is used.

The results for optimizing the average detector flux are shown in Table 6.8. For this problem, the results from all the calculations are within the 95% confidence

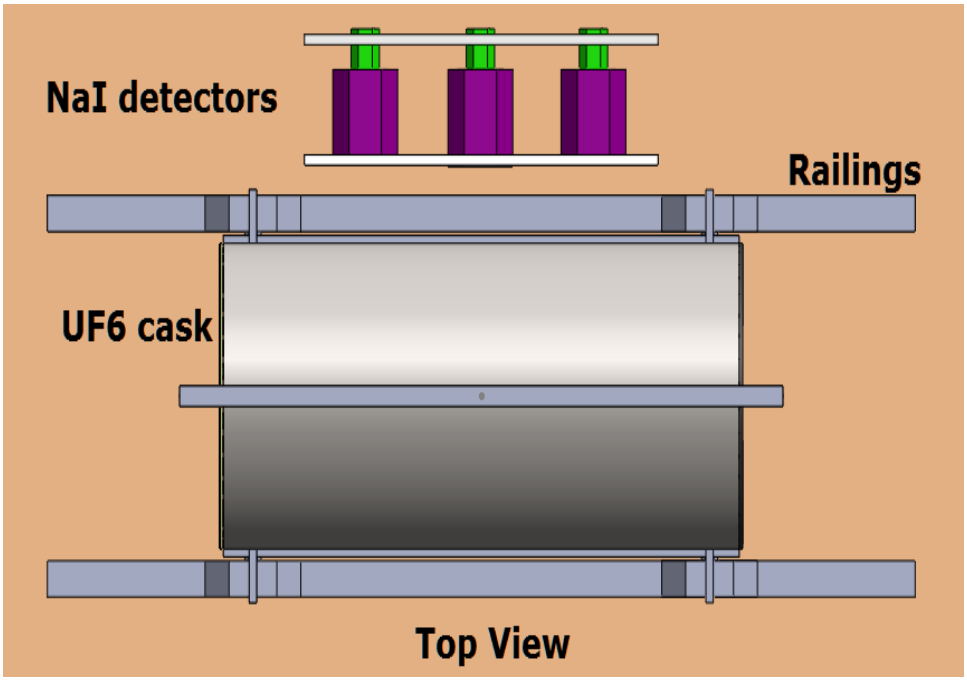

Figure 6.7: UF_6 problem top view

interval of each other. The hybrid methods resulted in small relative errors for a 100 batches with a batch size of 10^4 particles. But, the relative error for the analog calculation and the VVR P_1 method remained large even after increasing the number of particle histories by a factor of 100. The computational time for the deterministic calculation is approximately 750 seconds. The Monte Carlo calculation times for the hybrid methods except for the VVR method, are actually smaller than the time taken for the adjoint calculation and they still resulted in small relative errors when compared to the analog calculation. One important difference in the application of the LIFT method and the LIFT+WW method for the UF_6 problem is that the spatial biasing part (path length stretching based on the exponential transform) of the LIFT technique is not employed. This is due to the numerical restrictions caused by large

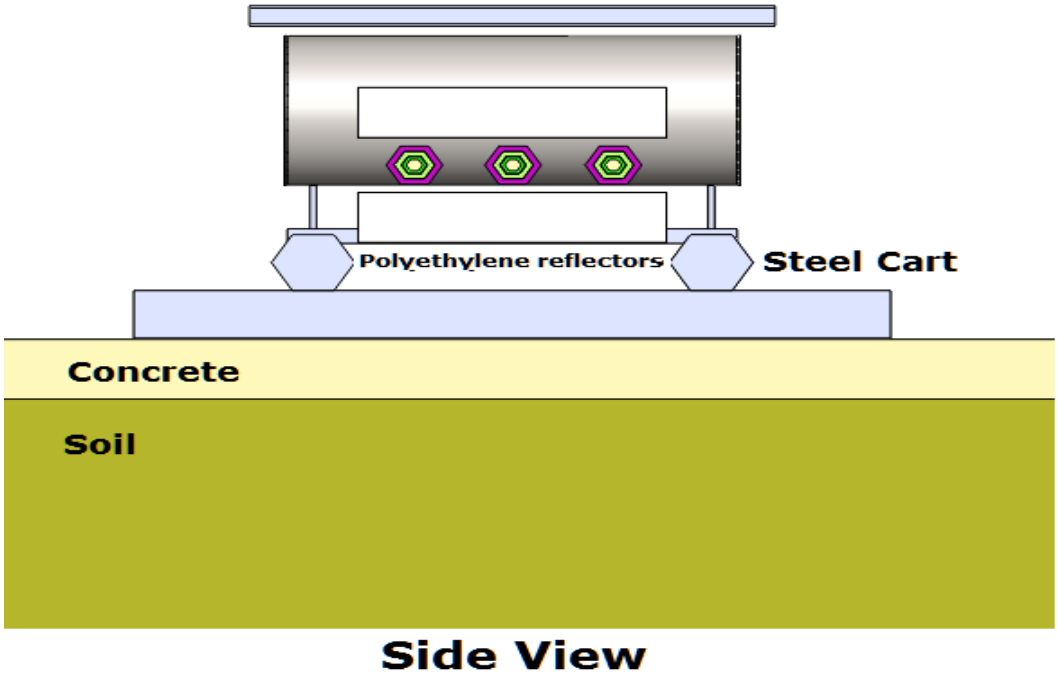

Figure 6.8: UF_6 problem side view

total cross-sections which resulted in large biasing parameters in certain regions of the phase-space. The exponential function used for the spatial representation in LIFT returns errors for regions with large biasing parameters when the values are outside the numeric limits of C++. A more detailed discussion of this issue is presented in the next chapter.

But, with only employing the angular biasing part of the LIFT method, it performed better with a FOM almost twice that of the SB+WW method. The LIFT+WW method performed the best for this problem with FOM that is thrice that of the LIFT method. The reason for the efficiency of the LIFT method even without spatial biasing is likely due to the nature of the problem. The problem has a large source

	Detector Flux	**Rel. Error**	**Time**	**FOM**
Analog+IA	5.65E-03	0.335	10076	0.0009
VVR P_1	6.14E-03	0.275	11014	0.0012
SB+WW	3.59E-03	0.100	865	0.1156
VVR P_1+WW	4.55E-03	0.087	981	0.1347
VVR S_2+WW	4.52E-03	0.088	1792	0.0721
LIFT	4.47E-03	0.058	1284	0.2315
LIFT+WW	4.82E-03	0.042	901	0.6292
VVR P_1+LIFT	4.48E-03	0.066	1403	0.1651
VVR S_2+WW+LIFT	4.81E-03	0.043	1824	0.2951

Table 6.8: Average detector flux in the NaI sensor for the middle detector

region and a small detector region without any shielding material in between. We think that source biasing plays a major role in sampling source particles in regions that are more likely to reach the detector and when combined with angular biasing, the LIFT and the LIFT+WW methods perform efficiently even without any spatial biasing.

The VVR method, when applied with the analog calculation and the SB+WW method provided marginally smaller relative errors than the direct estimates. However, the additional time for evaluating the VVR functional resulted in a marginal improvement in FOM for the P_1 method and is detrimental to the efficiency of the S_2 method. When VVR is applied with the LIFT and the LIFT+WW methods, the resulting flux values agreed with the flux values from the other calculations, but the errors remained higher than the corresponding direct estimates. This indicates that the benefit provided by employing the VVR functional is not helping the LIFT based calculations.

Chapter 7: Analysis and Conclusions

A platform to test sophisticated hybrid methods on realistic problems with reduced overhead and user effort has been successfully developed. A multi group Monte Carlo code that works seamlessly within the framework of the deterministic radiation transport tool Attila has been developed using an object-oriented approach. The use of an adjoint solution calculated using deterministic transport with the linear discontinuous spatial differencing scheme for variance reduction techniques has been studied. Two advanced automated variance reduction techniques, the LIFT method and the VVR method, along with the CADIS based weight windows and source biasing technique have been implemented in the test bench. The performance of the methods, both individually and in combination, in comparison with analog Monte Carlo calculation with only implicit capture has been studied for two realistic nonproliferation test problems. The benefits and limitations of the two hybrid methods from this research are described in the sections below.

7.1 Benefits and limitations of the LIFT method

The results of the hybrid methods for the mulch problem clearly showed that they are significantly better than the analog Monte Carlo calculations where the source

and detectors are separated by several mean free paths and the detector volume is relatively small when compared to the overall problem dimensions. The FOMs for the hybrid methods were at least an order of magnitude greater than those for the analog simulations in all the cases of the mulch problem. Also, the LIFT method performs significantly better than the weight windows method when the detector flux across all groups is optimized. The weight windows method performs better than the LIFT method when the detector absorption reaction rate in the four lower energy groups is optimized. The weight fluctuations in the LIFT method that were encountered while optimizing the reaction rates were mitigated when CADIS-based weight windows is applied. The time spent for calculating the adjoint solution in Attila is significantly shorter than the time needed to get reasonable estimates in the forward Monte Carlo simulation and the total run times for LIFT calculations were still less than those of the analog calculation.

The application of the LIFT method to the mulch test problem has helped to identify some of the limitations of the method that were not clearly understood when it was applied to simple test problems [15]. One of the major limitations is the ability to include higher order scattering moments in the LIFT method. For the scattering distributions to remain positive, the cosine of the scattering angle is limited to one-third, allowing for limited linear anisotropy as suggested in the earlier implementation [15]. However, even after restricting the value of the scattering cosine, negative adjoint angular fluxes were encountered for the mulch calculation. The cause of the negative angular fluxes was traced to the linearly anisotropic factor present in LIFT's angular

flux definition. In Equation (2.13) a spatially continuous expression based on the adjoint scalar flux and material properties is used to construct the local adjoint angular flux. The anisotropic factor in the angular flux expression, defined in Equation (2.15) is,

$$b_g\left(\mathbf{\Omega}\right) = 1 + 3\overline{\mu}_{g\to g}\frac{\sigma_{t,g} - \sigma_{s0,g\to g}}{\left|\boldsymbol{\rho}_g\right|^2}\boldsymbol{\rho}_g \cdot \mathbf{\Omega} \tag{7.1}$$

For the linearly anisotropic factor to be positive, (i.e, $b_g > 0$ for all $\mathbf{\Omega} \in 4\pi$), the following condition should be satisfied,

$$\overline{\mu}_{g\to g}\frac{\sigma_{t,g} - \sigma_{s0,g\to g}}{\left|\boldsymbol{\rho}_g\right|} < \frac{1}{3}.$$

This condition imposes additional limitations in representing the adjoint angular flux. Even with restricted average scattering cosines ($\mu_{g\to g} < 1/3$), in certain energy groups where $\sigma_{t,g} >> \sigma_{s0,g\to g}$, it is possible for the above condition to be violated. To preserve positivity, one can either restrict the magnitude of the biasing parameter ρ_g or use an isotropic representation of adjoint angular flux. For the mulch problem, an isotropic representation of adjoint flux performed better than having limits on the biasing parameter. This condition arises due to the LIFT expression being derived from the assumption of a 1-group problem for each energy group [6]. In future work, it would be beneficial to investigate the possibility of including an additional term to account for the presence of other energy groups while still retaining the simplistic nature of the expression.

Another potential cause for bias in the LIFT results was identified due to the nature

of the exponential function used to represent the spatial distribution of the adjoint flux locally within each cell. The spatial component of the adjoint flux within each cell is expressed by:

$$\widehat{\psi}(\boldsymbol{r}) = \widehat{\phi}_n V_n \frac{e^{\,\boldsymbol{\rho}\,\cdot\,(\boldsymbol{r}-\boldsymbol{r}_n)}}{\int\limits_{V_n} e^{\,\boldsymbol{\rho}\,\cdot\,(\boldsymbol{r}-\boldsymbol{r}_n)} d^3\boldsymbol{r}}, \tag{7.2}$$

where n is the spatial cell number and ϕ_n is the average adjoint scalar flux, V_n is the volume and $\boldsymbol{r}_n$ is the centroid of the cell. The range of the exponential function when the value of the biasing parameter is large or the spatial cell is large is many orders of magnitude. Since the expression is normalized by its integral in the denominator, it is a density function which becomes highly peaked when the argument of the exponential is large. The quantities sampled from such a peaked distribution tend to be biased as they do not converge to the expected value of the distribution in reasonable number of samples. We performed a calculation with only the angular biasing part of the LIFT method, ignoring the exponential function and the results fell within the 95% confidence interval of the analog estimate. In the test problem, where the reaction rates in the lower groups are optimized, we also tallied the total detector fluxes for comparison to the first case (The type of adjoint solution should only affect the efficiency of the calculation and not the result). We discovered that the weight windows method converged to flux values similar to those obtained when the flux in all groups was optimized, although the error was higher as expected. For the LIFT cases, the detector fluxes were actually lower than the analog estimate (4.4E-05 [LIFT], 4.7E-05 [Analog+IA]) by approximately the same amount as they were higher for the first case (5.1e-05 [LIFT]). The fluxes for the calculation with

only angular biasing and without the exponential transform agreed well for both cases (4.93E-05 [case 1], 4.99E-05 [case 2]). In future work, it is desirable to have a spatial function that exhibits smaller variations and at the same time remains positive and easy to sample. A linear form of the spatial function was considered for highly scattering problems in previous work [9], but was not recommended due to significantly increased computational time. It would be beneficial to investigate this further for realistic test problems using Tortilla.

Although the LIFT method exhibits minimal weight fluctuations for the most part when compared to the analog calculation and weight windows calculations, some outliers in the weight contributions to the detector are encountered, especially for the problem in which the four group reaction rates were optimized. These outliers occur very rarely and are due to inaccurate reconstruction of the importance function in certain regions of the phase space in LIFT. This affects the performance of the LIFT method as the error slightly increases whenever there is a larger than usual contribution to the tally, but eventually this effect diminishes.

For the PANDA problem, the LIFT method did not provide any performance benefits over the analog Monte Carlo calculation. The reduction in error for the LIFT calculations was very slight in comparison with the analog calculation. But, the time to calculate the adjoint solution in Attila for the PANDA cases was much longer than the Monte Carlo run times, yielding a much smaller Figure of Merit. This is because there are multiple detector locations in the problem and a highly scattering moderator region. The computational cost to generate the adjoint solution is much

more significant than the marginal reduction in error and time savings in the Monte Carlo calculation achieved by applying the LIFT technique. There could still be some potential benefits from applying the LIFT method if the adjoint solution is calculated via Monte Carlo calculation for such highly scattering problems.

For the UF_6 problem, the LIFT method was applied with a restriction. This restriction was due to the nature of the exponential function used for the spatial representation of adjoint angular flux. In regions with large total cross section, the biasing parameter becomes large and this causes numerical handling issues. A restriction on the biasing parameter's magnitude ($|\boldsymbol{\rho}|$) was made such that the value returned by the exponential function remains within the allowed numerical limits of C++, but this resulted in reduced biasing in optically thick regions of the problem. The computational time for calculating the modified distributions necessary to achieve path-length stretching in the LIFT implementation remains constant irrespective of the amount of the magnitude of $|\boldsymbol{\rho}|$. Therefore restricting $|\boldsymbol{\rho}|$ yields a method with poor efficiency. However, the LIFT method with only the angular biasing part performed efficiently for this problem and using an alternative spatial representation could be potentially more efficient.

7.2 Benefits and limitations of the VVR method

For the mulch problem and the UF_6 problem, the VVR method yielded an improvement in the FOM when used in conjunction with the analog calculation, and the

source biasing and weight windows (SB+WW) calculation. Also, the time required to evaluate the VVR functional was a small fraction of the overall computational time for the P_1 implementation and the more involved S_2 implementation. These improvements in error are similar to what Barrett reported for her 3D calculations [7]. For the PANDA problem, even though the VVR estimates had smaller variances than the direct estimates, the time required to evaluate the functionals was a significant portion of the run time, especially for the S_2 implementation. This is because in a highly scattering problem, all the angular bins in the phase space contribute to every batch and therefore evaluating the VVR functional becomes expensive. For the mulch and UF_6 problems, it is likely that many regions of the phase space did not contribute to estimates in a given batch and therefore the time for evaluating the functional was not significant.

The VVR method did not provide error reduction and produced highly biased estimates when applied with the LIFT method for all the problems. To study this behavior of the VVR method, recall from chapter 2 Equation (2.22c), we can write the VVR estimate for detector response as:

$$F[\psi, \widehat{\psi}] = R + d\psi d\widehat{\psi} \tag{7.3}$$

$d\widehat{\psi}$ is the truncation error from the deterministic adjoint calculation and is a constant in phase space for all particle batches. The quantity $d\psi$ represents the error in the forward solution globally. The VVR method achieves error reduction by using the estimate of the global forward solution and adjoint solution to add a correction term

to the direct detector response estimate which yields an VVR estimate that is closer to the actual detector response. Provided a reasonable estimate of the global forward and adjoint solutions are available, the correction term will ensure that the variability in the VVR estimate is smaller than the variability of the direct detector response estimate. However, in a complex problem (both in geometry and energy) with a small detector response, the second order correction term can become dominant and cause bias in the result. Two conditions under which the bias occurs are encountered during this research. The first condition occurred when VVR was used with an analog Monte Carlo calculation for the mulch problem. When a small batch size was used, the detector flux remained zero for several batches and the errors in the global flux estimate were also large resulting in a negative value for the detector response. One way to mitigate this would be to increase the batch size such that the error in all parts of the phase space is reasonable and the detector response is not zero. But this would mean, for a fixed number of particle histories, the number of batches will be very small and calculating the variance from a small number of samples would result in a bias. The other solution is to use VVR with other variance reduction techniques. When variance reduction techniques are applied, more particles would reach the detector and even for a smaller batch size, a reasonable estimate of the detector response will be calculated, avoiding the possibility of negative estimates of the detector response. This was the case when VVR was applied with weight windows.

However, when VVR was applied with the LIFT method, negative VVR estimates were again encountered as seen in the results for the mulch problem. The reason

for this is that the LIFT method biases the particles heavily such that the errors in the high importance regions are much smaller, while the error in the low importance regions remain large resulting in global forward error estimates that are much larger ($d\psi >> 0$) and the magnitude of the second order error correction is dominant ($|d\psi d\widehat{\psi}| >> R$). This is the reason for the negative VVR estimates for the LIFT calculations of the mulch problem. For the PANDA problem, the $d\psi d\widehat{\psi}$ term remained smaller than the detector response due to the fact that most regions of the problem had high importance and therefore a reasonable forward estimate was available in all regions of the phase space. It also appears that the weight windows method did not exhibit such behavior because of splitting which will cause some particles to still reach the lower importance regions of phase space, while in the LIFT method the particles traveling in the low important regions are more likely to get terminated. To illustrate this phenomenon, the relative errors in the total detector flux for the different material regions of the mulch problem from the analog, the weight windows and the LIFT calculations are plotted in Figure 7.1. As observed from the plot, the errors are much smaller than the LIFT method for the analog and weight windows calculations in all regions of the problem except the detector, whereas the error for the detector region is smallest for the LIFT calculation. The plot also shows that the errors for the weight windows method are relatively smaller in all regions of the problem unlike the LIFT method which yields significant error reduction only in the detector region. Therefore, it is fair to conclude that the VVR method works optimally when applied with weight windows calculations.

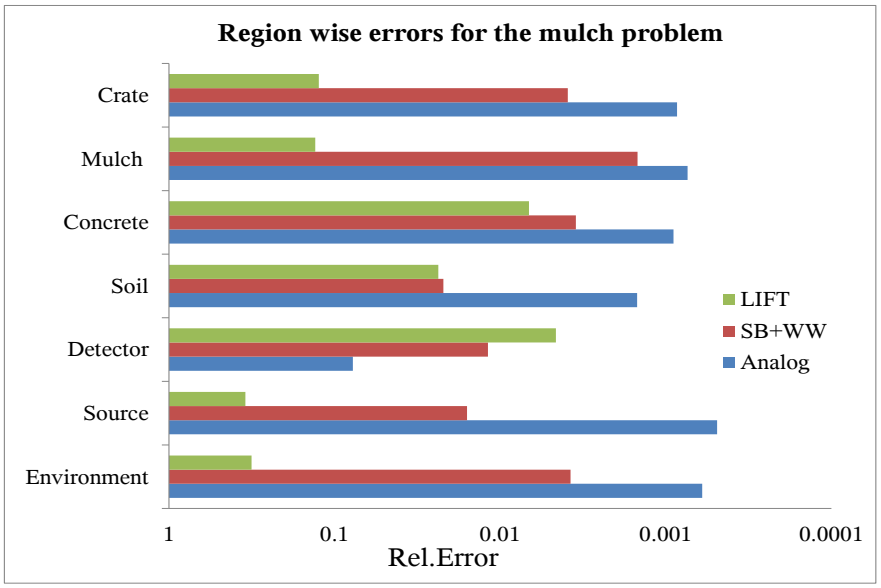

Figure 7.1: Region wise errors for the mulch problem case 1

7.3 Future work

In the future, besides serving as a platform for improving the existing hybrid methods, Tortilla can be modified to test other more complex test problems and new hybrid transport algorithms. Two such future extensions are discussed below.

7.3.1 Continuous energy Monte Carlo

In the MGMC code, as the name indicates, the energy domain is discretized using the multigroup approach. There are applications in which a continuous representation of

the energy domain is essential and most of the Monte Carlo calculations in production codes like MCNP are performed in continuous energy domain. Some of the variance reduction methods like the CADIS based weight windows lend themselves directly for continuous energy Monte Carlo calculations. However, methods like LIFT have been tested only in Multigroup calculations and the feasibility of extending these methods to continuous energy calculations needs to be analyzed.

One of the major challenges in implementing a continuous energy calculation in Tortilla is the ability to read from continuous energy cross section libraries and using them in the particle transport routines. Since our test bench is developed around a deterministic tool, the APIs can read only those cross sections that are in the multigroup format. A separate module to import the continuous cross-section data into the Monte Carlo code is required. One possible approach would be to make use of the cross section APIs that are developed for the open source continuous energy Monte Carlo code, OpenMC [34]. The OpenMC code package is developed in C++ to solve reactor criticality calculations. The cross section format used by OpenMC is ACE. The ACE format contains continuous-energy cross sections for the following types of reactions: elastic scattering, fission (or first-chance fission, second-chance fission, etc.), inelastic scattering, (n,xn), (n,γ), and various other absorption reactions. OpenMC also includes sophisticated look up techniques to fetch the appropriate cross section data for particle transport.

Another major modification that should be done to the test bench to accommodate continuous energy Monte Carlo would be the sampling distributions for secondary

angle and energy after every scattering collision. OpenMC employs various techniques to sample the scattering angle and energy depending on the reaction type. They include simply isotropic sampling, sampling from equiprobable bins and sampling from tabular distributions. The routines that perform these methods are available and could be modified to work with Tortilla.

7.3.2 Implementing adaptive Monte Carlo algorithms

There have been significant research efforts in the area of adaptive Monte Carlo methods that could achieve exponential convergence and these methods have only been tested on limited problem with fewer dimensions [35] [36]. Tortilla, due to its modular nature can serve as the right platform to implement and test these methods on complex problems. Further, the deterministic solution available from Attila could be used as the initial starting point for the learning algorithms. This in turn would result in fewer batches to achieve the desired error levels rather than starting with zero information. Also, adaptive methods can be used to improve the accuracy of the biasing parameters from the forward solution in the previous batches. The major modification in implementing these algorithms would be in configuring the appropriate tallies, the sampling distributions based on the methods used and of course the algorithms itself.

Bibliography

[1] R. Sanchez and N.J. McCormick. A Review of Neutron Transport Approximations. *Nuclear Science & Engineering*, 80:481–535, 1982.

[2] MCNP – A General Monte Carlo N-Particle Transport Code, Version 5, Volume I: Overview and Theory, 2003. LA-UR-03-1987, Los Alamos National Laboratory.

[3] A. Haghighat and J.C. Wagner. Monte Carlo Variance Reduction with Deterministic Importance Functions. *Progress in Nuclear Energy*, 42:25–53, 2003.

[4] D.E. Peplow. Monte Carlo Shielding Analysis Capabilities with MAVRIC. *Nuclear Technology*, 174:289–313, 2011.

[5] J.C. Wagner, D.E.Peplow, S.W. Mosher and T.M. Evans. Review of Hybrid (Deterministic/Monte Carlo) Radiation Transport Methods, Codes, and Applications at Oak Ridge National Laboratory. DE-AC05-00OR22725. Published by Atomic Energy Society of Japan.

[6] S.A. Turner and E.W. Larsen. Automatic Variance Reduction for Three-Dimensional Monte Carlo simulations by the Local Importance Function Transform 1: Analysis. *Nuclear Science & Engineering*, 127:22–35, 1997.

[7] C.L. Barrett. A Variationally-Based Variance Reduction Method for Monte Carlo Particle Transport Problems, 1999. Ph.D. Thesis, Department of Nuclear Engineering, University of Michigan.

[8] T M. Evans , A.S. Stafford, K.T. Clarno. Denovo: A New Three-Dimensional Parallel Discrete Ordinates Code in SCALE. *Nuclear Technology*, 2010.

[9] S.A. Turner. Automatic Variance Reduction for Monte Carlo simulations via the Local Importance Function Transform, 1996. University of Michigan, Ann Arbor, MI, USA.

[10] K.P. Keady and E.W. Larsen. A Modified Monte Carlo "Local Importance Function Transform" Method. In *International Conference on Mathematics and Computational Methods Applied to Nuclear Science & Engineering M&C (2013)*, Sun Valley, Idaho, 2013.

[11] J.D. Densmore and E.W. Larsen. A New Variational Variance Reduction Method for Monte Carlo Source-Detector Problems. *Transactions of the American Nuclear Society*, 83:336, 2000.

[12] J.D. Densmore. Variational Variance Reduction for Monte Carlo Reactor Analysis. Ph.D. Thesis, Department of Nuclear Engineering, University of Michigan (2002).

[13] C.L. Barrett and E.W. Larsen. A Variationally-Based Variance Reduction Method for Monte Carlo Neutron Transport Calculations. *Annals of Nuclear Energy*, 28:457, 2001.

[14] J.D. Densmore and E.W. Larsen. Variational Variance Reduction for Monte Carlo Eigenvalue and Eigenfunction Problems. *Nuclear Science & Engineering*, 146:121, 2004.

[15] S.A. Turner and E.W. Larsen. Automatic Variance Reduction for Three-Dimensional Monte Carlo simulations by the Local Importance Function Transform – 2: Numerical results. *Nuclear Science & Engineering*, 127:36–53, 1997.

[16] J. McGhee and T. Wareing and D. Barnett. Attila User's Manual, 2007. Transpire, Inc.

[17] J.E Morel, T.A. Wareing and K. Smith. A Linear-Discontinuous Spatial Differencing Scheme for Sn Radiative Transfer Calculations. *Journal of Computational Physics*, 128:445–462, 1996.

[18] E.E. Lewis and W.F., Jr. Miller. *Computational Methods of Neutron Transport*. John Wiley and Sons, NY, 1984.

[19] R.S. Baker. The XREP Module, 1995. LA-UR-95-4473, Los Alamos National Laboratory.

[20] B. Karlsson. Smart Pointers in Boost. C/C++ Users Journal. `www.cuj.com/documents/s=8014/cuj0204karlsson/`.

[21] H. Sutter. Generalized Function Pointers C/C++ Users Journal. `www.cuj.com/documents/s=8464/cujcexp0308sutter/`.

[22] T. Nishimura. Tables of 64-bit Mersenne Twisters. *ACM Transactions on Modeling and Computer Simulation*, 10:348–457, 2000.

[23] M. Matsumoto and T. Nishimura. Mersenne Twister: A 623-Dimensionally Equidistributed Uniform Pseudo-Random Number Generator. *ACM Transactions on Modeling and Computer Simulation*, 8, 1998.

[24] C.J. Everett and E.D. Cashwell. A Third Monte Carlo Sampler, 1983. LA-9721-MS.

[25] E. Somasundaram and T.S. Palmer. Implementation of Hybrid Variance Reduction Methods in A Multigroup Monte Carlo code for Deep Shielding Problems. In *International Conference on Mathematics and Computational Methods Applied to Nuclear Science & Engineering M&C (2013)*, Sun Valley, Idaho, 2013.

[26] J.R. Mika. Neutron Transport with Anisotropic Scattering. *Nuclear Science & Engineering*, 11:415–427, 1961.

[27] C. Hollauer. Modeling of Thermal Oxidation and Stress Effects. PhD Thesis, Technical University of Vienna, Vienna, Austria (2007).

[28] Sphere Lebedev Rule - Quadrature Rules for the Unit Sphere. `http://people.sc.fsu.edu/~jburkardt/f_src/sphere_lebedev_rule/sphere_lebedev_rule.html`.

[29] M. Shermana and D. Goldsmanb. Large-Sample Normality of the Batch-Means Variance Estimator. *Operations Research Letters*, 30:319–326, 2003.

[30] M.W. Shaver, E.A. Miller, R.S. Wittman and B.S. McDonald. Transport Test Problems for Hybrid Methods Development. Pacific Northwest National Laboratory (PNNL), Richland, WA, USA (2011).

[31] Physics Guide to CEPXS: A Multigroup Coupled Electron-Photon Cross-Section Generating Code, 1989. Ver 1.0, Sandia National Laboratory report SAND89-1685.

[32] M.D. DeHart and S.M. Bowman. Validation of the SCALE Broad Structure 44-Group ENDF/B-V Cross-Section Library for Use in Criticality Safety Analyses, 1992. Oak Ridge National Laboratory.

[33] Tech. Report PSR-315, 1992. AMPX-77, Radiation Safety Information Computational Center.

[34] P.K. Romano and B. Forget. The OpenMC Monte Carlo Particle Transport Code. *Annals of Nuclear Energy*, 51:274–281, 2013.

[35] R.Kong and J. Spanier. A New Proof of Geometric Convergence for General Transport Problems Based on Sequential Correlated Sampling Methods. *Journal of computational physics*, 227(23):9762–9777, 2008.

[36] L. Li and J. Spanier. Approximation of Transport Equations by Matrix Equations and Sequential Sampling. *Monte Carlo Methods and Applications*, 1997.

[37] D.M. Mount and S. Arya. ANN: A Library for Approximate Nearest Neighbor Searching. `http://www.cs.umd.edu/~mount/ANN/`, 2010.

APPENDICES

Appendix A: Algorithm for transferring solution quantities between arbitrary tetrahedral meshes

The algorithm used for transferring the adjoint solution quantities from a fine arbitrary tetrahedral mesh to a coarser mesh is described below:

1. The fine mesh solution quantities for each tetrahedral cell (average flux and currents) and the vertices (cell vertex flux and currents) are extracted from the flux moments output generated by Attila. Since Attila uses a linear discontinuous spatial discretization scheme, the vertices in the mesh can have more than one solution quantity. In such cases, the average of all the solutions at the vertex is calculated.

2. A stack of the position vectors of the "cell centers" and "vertices" (we will refer to them as nodes) in the fine mesh along with their corresponding solution quantities is created.

3. A stack of the position vectors of the cell centers in the coarse mesh is also created.

4. A "K-D search tree" [37] is then constructed using the stack of position vectors in the fine mesh.

5. For each cell in the coarse mesh, the longest distance between the cell center and its vertices is determined.

6. The nearest neighbor search algorithm is then invoked to perform a radius search for the distance calculated in the previous step. The search algorithm returns the nodes in the fine mesh that lie within the imaginary sphere constructed around the tetrahedral cell.

7. The nodes returned by the search algorithm are then evaluated to determine if they are within the cell boundaries. The nodes that lie within the coarse mesh cell and those that lie outside are stored in two different arrays along with their respective distances to the cell center.

8. The search procedure is repeated for all the coarse mesh cells and resulting fine mesh nodes are stored in arrays as described in the previous step.

9. For all the coarse mesh cells that have at least one spatial node of the fine mesh within the cell boundaries, the coarse mesh solution is determined by a direct average of the fine mesh solution quantities at all the nodes within the cell.

10. For coarse mesh cells that have no fine mesh nodes within the cell, an inverse distance weighting based interpolation scheme is employed. The nodes that were determined to lie outside the cell boundaries are invoked and their distances to the coarse mesh cell center are used to determine weights for each node. The coarse mesh solution is then calculated by taking a weighted average of the

solution quantities at the nodes.

11. If there are no fine mesh nodes within the imaginary sphere around the coarse mesh cell, the nearest neighbor search is invoked again to determine a fixed number of closest nodes in the fine mesh.

12. The solution for that cell is then interpolated by a similar inverse distance weighting scheme described previously.

The search algorithm used in the above procedure comes from the ANN(Approximate Nearest Neighbor) C++ library [37]. ANN supports data structures and algorithms for both exact and approximate nearest neighbor searching in arbitrarily high dimensions. The library implements a number of different data structures, based on kd-trees and box-decomposition trees, and employs a couple of different search strategies. Currently we only employ the exact search features of this library and the search times are very quick for our purposes. For larger problems, the approximate search algorithms could be turned on, provided the overhead is high enough to justify the resulting loss in accuracy. For inverse distance weighting, we employ the fairly standard procedure to determine the weights based on the distances. The formula for determining the weights is given by,

$$w_i(r, r') = \frac{1}{dist\,(r, r')^p}, \tag{A.1}$$

where p determines the influence of the points that are away from the point at which the solution is interpolated. Higher values for p places stronger influence on points

closer to the interpolation point while a weaker influence on the points further away. The solution quantities (ϕ_r) are determined using the formula below,

$$\phi_r = \frac{\sum_{i \in N} w_i \phi_i}{\sum_{i \in N} w_i}. \tag{A.2}$$

Appendix B: Mulch - Source spectrum and energy group structure

Table B1: Ba-133 Source spectrum

Energy Group	Source Strength
3	8.9160E-02
6	6.2150E-01
12	1.8400E-01
16	7.0850E-02
21	4.5960E-03
26	5.9780E-03
35	3.4180E-01
37	3.1800E-02
41	2.1690E-02
44	4.3950E-02
46	1.7970E-01
48	1.0130E+00

Table B2: Energy group structure for the mulch problem

Group	High E. (MeV)	Low E. (MeV)	Group	High E. (MeV)	Low E. (MeV)
1	4.0000E-01	3.9215E-01	26	1.7067E-01	1.5986E-01
2	3.9215E-01	3.8434E-01	27	1.5986E-01	1.4924E-01
3	3.8434E-01	3.8334E-01	28	1.4924E-01	1.3883E-01
4	3.8334E-01	3.6985E-01	29	1.3883E-01	1.2865E-01
5	3.6985E-01	3.5650E-01	30	1.2865E-01	1.1869E-01
6	3.5650E-01	3.5550E-01	31	1.1869E-01	1.0898E-01
7	3.5550E-01	3.4488E-01	32	1.0898E-01	9.9532E-02
8	3.4488E-01	3.3435E-01	33	9.9532E-02	9.0361E-02
9	3.3435E-01	3.2392E-01	34	9.0361E-02	8.1489E-02
10	3.2392E-01	3.1358E-01	35	8.1489E-02	8.0489E-02
11	3.1358E-01	3.0335E-01	36	8.0489E-02	8.0112E-02
12	3.0335E-01	3.0235E-01	37	8.0112E-02	7.9112E-02
13	3.0235E-01	2.9379E-01	38	7.9112E-02	7.1416E-02
14	2.9379E-01	2.8530E-01	39	7.1416E-02	6.3934E-02
15	2.8530E-01	2.7689E-01	40	6.3934E-02	5.6682E-02
16	2.7689E-01	2.7589E-01	41	5.6682E-02	4.9677E-02
17	2.7589E-01	2.6357E-01	42	4.9677E-02	4.2942E-02
18	2.6357E-01	2.5139E-01	43	4.2942E-02	3.6506E-02
19	2.5139E-01	2.3937E-01	44	3.6506E-02	3.5506E-02

20	2.3937E-01	2.2750E-01	45	3.5506E-02	3.5467E-02
21	2.2750E-01	2.1579E-01	46	3.5467E-02	3.4467E-02
22	2.1579E-01	2.0425E-01	47	3.4467E-02	3.1473E-02
23	2.0425E-01	1.9288E-01	48	3.1473E-02	3.0473E-02
24	1.9288E-01	1.8168E-01	49	3.0473E-02	2.5000E-02
25	1.8168E-01	1.7067E-01	50	2.5000E-02	1.0000E-03

Appendix C: PANDA - Source spectrum and energy group structure

Table C1: Cf-252 spectrum

Energy Group	Source Strength	Energy Group	Source Strength
1	1.78E-02	24	7.70E-09
2	2.31E-02	25	3.75E-09
3	3.98E-02	26	2.25E-09
4	2.12E-01	27	2.50E-10
5	6.14E-02	28	2.50E-10
6	1.47E-02	29	2.50E-10
7	8.23E-02	30	5.00E-10
8	1.28E-01	31	2.50E-10
9	1.46E-01	32	2.50E-10
10	1.65E-01	33	2.50E-10
11	8.57E-02	34	5.00E-10
12	1.19E-02	35	5.00E-10
13	8.43E-04	36	3.00E-10
14	9.86E-04	37	2.00E-10
15	7.90E-05	38	1.00E-10

16	4.50E-06	39	1.00E-10
17	7.00E-07	40	4.70E-11
18	2.00E-07	41	1.53E-10
19	1.90E-08	42	2.50E-11
20	2.10E-08	43	4.50E-11
21	1.25E-08		
22	1.75E-08		
23	1.23E-08		

Table C2: Energy group structure for the PANDA problem

Group	High E. (MeV)	Low E. (MeV)	Group	High E. (MeV)	Low E. (MeV)
1	2.00E+01	8.19E+00	24	1.77E-06	1.00E-06
2	8.19E+00	6.43E+00	25	1.00E-06	6.25E-07
3	6.43E+00	4.80E+00	26	6.25E-07	4.00E-07
4	4.80E+00	3.00E+00	27	4.00E-07	3.75E-07
5	3.00E+00	2.48E+00	28	3.75E-07	3.50E-07
6	2.48E+00	2.35E+00	29	3.50E-07	3.25E-07
7	2.35E+00	1.85E+00	30	3.25E-07	2.75E-07
8	1.85E+00	1.40E+00	31	2.75E-07	2.50E-07
9	1.40E+00	9.00E-01	32	2.50E-07	2.25E-07

10	9.00E-01	4.00E-01	33	2.25E-07	2.00E-07
11	4.00E-01	1.00E-01	34	2.00E-07	1.50E-07
12	1.00E-01	2.50E-02	35	1.50E-07	1.00E-07
13	2.50E-02	1.70E-02	36	1.00E-07	7.00E-08
14	1.70E-02	3.00E-03	37	7.00E-08	5.00E-08
15	3.00E-03	5.50E-04	38	5.00E-08	4.00E-08
16	5.50E-04	1.00E-04	39	4.00E-08	3.00E-08
17	1.00E-04	3.00E-05	40	3.00E-08	2.53E-08
18	3.00E-05	1.00E-05	41	2.53E-08	1.00E-08
19	1.00E-05	8.10E-06	42	1.00E-08	7.50E-09
20	8.10E-06	6.00E-06	43	7.50E-09	3.00E-09
21	6.00E-06	4.75E-06			
22	4.75E-06	3.00E-06			
23	3.00E-06	1.77E-06			

Appendix D: UF_6 - Source spectrum and energy group structure

Table C1: U235-U238 photon spectrum

Energy Group	Source Strength	Energy Group	Source Strength
1	3.69E-16	93	1.31E-03
2	3.65E-16	94	5.80E-05
3	3.60E-16	95	5.16E-05
4	3.55E-16	96	5.80E-05
5	3.51E-16	97	3.89E-04
6	3.46E-16	98	5.80E-05
7	3.41E-16	99	9.76E-04
8	3.37E-16	100	9.60E-04
9	3.31E-16	101	5.80E-05
10	3.27E-16	102	3.98E-04
11	3.22E-16	103	5.80E-05
12	3.17E-16	104	1.23E-03
13	3.12E-16	105	1.20E-03
14	3.07E-16	106	1.18E-03
15	3.02E-16	107	1.15E-03

16	2.40E-04	108	1.13E-03
17	3.93E-04	109	1.10E-03
18	3.87E-04	110	5.80E-05
19	3.80E-04	111	1.14E-03
20	3.72E-04	112	1.11E-03
21	3.65E-04	113	1.08E-03
22	8.09E-06	114	5.80E-05
23	5.01E-05	115	1.04E-03
24	8.09E-06	116	1.02E-03
25	2.85E-04	117	9.95E-04
26	8.09E-06	118	9.73E-04
27	2.64E-04	119	9.52E-04
28	2.61E-04	120	9.30E-04
29	8.09E-06	121	9.07E-04
30	2.18E-04	122	8.85E-04
31	8.09E-06	123	8.62E-04
32	3.06E-04	124	8.39E-04
33	3.01E-04	125	4.19E-02
34	2.96E-04	126	4.07E-02
35	2.92E-04	127	3.95E-02
36	2.87E-04	128	3.82E-02
37	8.09E-06	129	3.70E-02
38	1.73E-04	130	3.57E-02

39	1.71E-04	131	3.44E-02
40	8.09E-06	132	2.98E-03
41	3.00E-04	133	3.06E-02
42	2.96E-04	134	2.98E-02
43	8.09E-06	135	2.89E-02
44	2.93E-04	136	2.98E-03
45	4.63E-03	137	2.59E-02
46	4.64E-03	138	2.52E-02
47	4.56E-03	139	2.45E-02
48	4.48E-03	140	2.39E-02
49	4.39E-03	141	2.98E-03
50	4.32E-03	142	2.07E-02
51	1.33E-04	143	2.98E-03
52	4.44E-03	144	2.42E-02
53	4.36E-03	145	2.36E-02
54	4.28E-03	146	2.28E-02
55	4.20E-03	147	2.21E-02
56	4.11E-03	148	2.98E-03
57	4.02E-03	149	2.09E-02
58	1.33E-04	150	2.03E-02
59	3.48E-03	151	1.96E-02
60	3.42E-03	152	2.98E-03
61	1.33E-04	153	1.58E-02

62	1.29E-04	154	1.53E-02
63	1.33E-04	155	1.49E-02
64	2.43E-03	156	2.98E-03
65	1.33E-04	157	3.19E-03
66	6.63E-06	158	2.98E-03
67	1.33E-04	159	2.29E-03
68	3.58E-04	160	2.98E-03
69	1.33E-04	161	1.76E-02
70	2.96E-03	162	1.69E-02
71	1.33E-04	163	2.98E-03
72	1.91E-03	164	4.68E-03
73	1.33E-04	165	2.98E-03
74	2.27E-04	166	6.31E-04
75	1.33E-04	167	2.98E-03
76	4.22E-04	168	2.60E-03
77	1.33E-04	169	2.98E-03
78	2.94E-03	170	1.28E-02
79	2.89E-03	171	1.24E-02
80	1.33E-04	172	2.98E-03
81	6.13E-04	173	1.26E-02
82	1.33E-04	174	1.21E-02
83	1.98E-03	175	2.98E-03
84	1.33E-04	176	9.95E-03

85	3.99E-04	177	9.64E-03
86	1.33E-04	178	9.33E-03
87	5.55E-04	179	2.98E-03
88	5.80E-05	180	1.00E-02
89	4.65E-04	181	9.67E-03
90	5.80E-05	182	9.30E-03
91	1.10E-03	183	8.93E-03
92	5.80E-05	184	1.49E-04

Table C2: Energy group structure for the UF_6 problem

Group	High E. (MeV)	Low E. (MeV)	Group	High E. (MeV)	Low E. (MeV)
1	3.00E+00	2.94E+00	101	6.99E-01	6.98E-01
2	2.94E+00	2.88E+00	102	6.98E-01	6.91E-01
3	2.88E+00	2.82E+00	103	6.91E-01	6.90E-01
4	2.82E+00	2.76E+00	104	6.90E-01	6.69E-01
5	2.76E+00	2.70E+00	105	6.69E-01	6.49E-01
6	2.70E+00	2.64E+00	106	6.49E-01	6.28E-01
7	2.64E+00	2.59E+00	107	6.28E-01	6.08E-01
8	2.59E+00	2.53E+00	108	6.08E-01	5.89E-01
9	2.53E+00	2.47E+00	109	5.89E-01	5.70E-01
10	2.47E+00	2.42E+00	110	5.70E-01	5.69E-01
11	2.42E+00	2.37E+00	111	5.69E-01	5.49E-01

12	2.37E+00	2.31E+00	112	5.49E-01	5.30E-01
13	2.31E+00	2.26E+00	113	5.30E-01	5.12E-01
14	2.26E+00	2.21E+00	114	5.12E-01	5.11E-01
15	2.21E+00	2.16E+00	115	5.11E-01	4.93E-01
16	2.16E+00	2.11E+00	116	4.93E-01	4.75E-01
17	2.11E+00	2.06E+00	117	4.75E-01	4.58E-01
18	2.06E+00	2.01E+00	118	4.58E-01	4.41E-01
19	2.01E+00	1.97E+00	119	4.41E-01	4.25E-01
20	1.97E+00	1.92E+00	120	4.25E-01	4.09E-01
21	1.92E+00	1.88E+00	121	4.09E-01	3.93E-01
22	1.88E+00	1.87E+00	122	3.93E-01	3.78E-01
23	1.87E+00	1.87E+00	123	3.78E-01	3.63E-01
24	1.87E+00	1.87E+00	124	3.63E-01	3.48E-01
25	1.87E+00	1.83E+00	125	3.48E-01	3.34E-01
26	1.83E+00	1.83E+00	126	3.34E-01	3.21E-01
27	1.83E+00	1.80E+00	127	3.21E-01	3.07E-01
28	1.80E+00	1.77E+00	128	3.07E-01	2.95E-01
29	1.77E+00	1.77E+00	129	2.95E-01	2.82E-01
30	1.77E+00	1.74E+00	130	2.82E-01	2.70E-01
31	1.74E+00	1.74E+00	131	2.70E-01	2.59E-01
32	1.74E+00	1.70E+00	132	2.59E-01	2.58E-01
33	1.70E+00	1.66E+00	133	2.58E-01	2.47E-01
34	1.66E+00	1.63E+00	134	2.47E-01	2.37E-01

35	1.63E+00	1.59E+00	135	2.37E-01	2.28E-01
36	1.59E+00	1.55E+00	136	2.28E-01	2.27E-01
37	1.55E+00	1.55E+00	137	2.27E-01	2.18E-01
38	1.55E+00	1.53E+00	138	2.18E-01	2.10E-01
39	1.53E+00	1.51E+00	139	2.10E-01	2.01E-01
40	1.51E+00	1.51E+00	140	2.01E-01	1.93E-01
41	1.51E+00	1.47E+00	141	1.93E-01	1.92E-01
42	1.47E+00	1.44E+00	142	1.92E-01	1.85E-01
43	1.44E+00	1.43E+00	143	1.85E-01	1.84E-01
44	1.43E+00	1.40E+00	144	1.84E-01	1.76E-01
45	1.40E+00	1.36E+00	145	1.76E-01	1.68E-01
46	1.36E+00	1.33E+00	146	1.68E-01	1.61E-01
47	1.33E+00	1.29E+00	147	1.61E-01	1.53E-01
48	1.29E+00	1.26E+00	148	1.53E-01	1.52E-01
49	1.26E+00	1.23E+00	149	1.52E-01	1.45E-01
50	1.23E+00	1.19E+00	150	1.45E-01	1.38E-01
51	1.19E+00	1.19E+00	151	1.38E-01	1.32E-01
52	1.19E+00	1.16E+00	152	1.32E-01	1.31E-01
53	1.16E+00	1.13E+00	153	1.31E-01	1.26E-01
54	1.13E+00	1.09E+00	154	1.26E-01	1.20E-01
55	1.09E+00	1.06E+00	155	1.20E-01	1.15E-01
56	1.06E+00	1.03E+00	156	1.15E-01	1.14E-01
57	1.03E+00	1.00E+00	157	1.14E-01	1.13E-01

58	1.00E+00	1.00E+00	158	1.13E-01	1.12E-01
59	1.00E+00	9.74E-01	159	1.12E-01	1.12E-01
60	9.74E-01	9.48E-01	160	1.12E-01	1.11E-01
61	9.48E-01	9.47E-01	161	1.11E-01	1.05E-01
62	9.47E-01	9.47E-01	162	1.05E-01	9.89E-02
63	9.47E-01	9.46E-01	163	9.89E-02	9.79E-02
64	9.46E-01	9.27E-01	164	9.79E-02	9.64E-02
65	9.27E-01	9.26E-01	165	9.64E-02	9.54E-02
66	9.26E-01	9.26E-01	166	9.54E-02	9.52E-02
67	9.26E-01	9.25E-01	167	9.52E-02	9.42E-02
68	9.25E-01	9.22E-01	168	9.42E-02	9.33E-02
69	9.22E-01	9.21E-01	169	9.33E-02	9.23E-02
70	9.21E-01	8.99E-01	170	9.23E-02	8.80E-02
71	8.99E-01	8.98E-01	171	8.80E-02	8.38E-02
72	8.98E-01	8.84E-01	172	8.38E-02	8.28E-02
73	8.84E-01	8.83E-01	173	8.28E-02	7.86E-02
74	8.83E-01	8.81E-01	174	7.86E-02	7.45E-02
75	8.81E-01	8.80E-01	175	7.45E-02	7.35E-02
76	8.80E-01	8.77E-01	176	7.35E-02	7.02E-02
77	8.77E-01	8.76E-01	177	7.02E-02	6.69E-02
78	8.76E-01	8.54E-01	178	6.69E-02	6.38E-02
79	8.54E-01	8.32E-01	179	6.38E-02	6.28E-02
80	8.32E-01	8.31E-01	180	6.28E-02	5.94E-02

81	8.31E-01	8.26E-01	181	5.94E-02	5.62E-02
82	8.26E-01	8.25E-01	182	5.62E-02	5.31E-02
83	8.25E-01	8.10E-01	183	5.31E-02	5.01E-02
84	8.10E-01	8.09E-01	184	5.01E-02	4.91E-02
85	8.09E-01	8.06E-01	185	4.91E-02	4.56E-02
86	8.06E-01	8.05E-01	186	4.56E-02	4.23E-02
87	8.05E-01	7.96E-01	187	4.23E-02	3.91E-02
88	7.96E-01	7.95E-01	188	3.91E-02	3.60E-02
89	7.95E-01	7.87E-01	189	3.60E-02	3.31E-02
90	7.87E-01	7.86E-01	190	3.31E-02	3.03E-02
91	7.86E-01	7.67E-01	191	3.03E-02	2.76E-02
92	7.67E-01	7.66E-01	192	2.76E-02	2.51E-02
93	7.66E-01	7.43E-01	193	2.51E-02	2.26E-02
94	7.43E-01	7.42E-01	194	2.26E-02	2.04E-02
95	7.42E-01	7.41E-01	195	2.04E-02	1.83E-02
96	7.41E-01	7.40E-01	196	1.83E-02	1.63E-02
97	7.40E-01	7.34E-01	197	1.63E-02	1.45E-02
98	7.34E-01	7.33E-01	198	1.45E-02	1.28E-02
99	7.33E-01	7.16E-01	199	1.28E-02	1.13E-02
100	7.16E-01	6.99E-01	200	1.13E-02	1.00E-02