

ABSTRACT

HIGGINS, NOAH MACKENZIE. Mitigation of Ray Effects in the THOR Discrete Ordinates Neutron Transport Code by Coupling with MCNP's Cartesian Grid Flux Output Tallies. (Under the direction of Dr. Yousry Azmy).

Among the more commonly-used numerical methods of finding solutions to the neutron transport equation is the deterministic method of discrete ordinates (S_N). Due to the angular discretization employed in this method, a physically erroneous phenomenon known as ray effects is prevalent in problem configurations having low-scattering media and spatially-localized sources. Remedies to ray effects have been sought since their discovery to enable use of the S_N method in such cases in order to benefit from its computational efficiency. Among these is an approach in which the neutron flux distribution from the first few particle collisions of a Monte Carlo (MC) run is extracted and imported into an S_N source calculation.

The MCNP transport code is among the most widely-used software packages in the area of neutron transport simulation. It has the ability to tally neutron scalar fluxes having an arbitrary collision count in a user-specified Cartesian grid. A new coupling scheme in which these fluxes are mapped into distributed fixed sources suitable for the open source S_N code THOR has been developed in this work. MCNP is favored over earlier, specialty MC codes, designed to be compatible with a specific coupled S_N code and employing identical spatial discretization schemes. Our new method therefore demonstrates the ability to use a more general mapping procedure between different meshes, for example, from a Cartesian to tetrahedral mesh, while also employing a general-purpose, widely-used MC code. Our ultimate objective is to make MC coupling an option to mitigate ray effects in any S_N solver.

Two configurations which are severely impacted by ray effects in standard S_N calculations, even for quadrature orders as high as S_{16} , show drastic improvement when applying the new MCNP-based coupling scheme. As neutrons with more collisions

simulated in MCNP are used to generate MC-collided sources for THOR, the computed responses of the detectors in the considered configurations approach those of continuous energy MCNP calculations. When utilizing neutron flux with as few as three collisions modeled by MCNP, ray effects are effectively mitigated in THOR's solution with a low quadrature order of S_2 .

© Copyright 2023 by Noah Mackenzie Higgins

All Rights Reserved

Mitigation of Ray Effects in the THOR Discrete Ordinates Neutron Transport Code by
Coupling with MCNP's Cartesian Grid Flux Output Tallies

by
Noah Mackenzie Higgins

A thesis submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the degree of
Master of Science

Nuclear Engineering

Raleigh, North Carolina
2023

APPROVED BY:

Dr. Dmitriy Anistratov

Dr. Semyon Tsynkov

Dr. Yousry Azmy
Chair of Advisory Committee

DEDICATION

To the memory of my late grandfather, Eugene.

BIOGRAPHY

Noah Higgins was raised in Clemmons, North Carolina prior to his admission to North Carolina State University as an undergraduate student in 2018. There, he obtained his B.S. in Nuclear Engineering in May 2022. He continued to study for his M.S. in Nuclear Engineering with support from a University Nuclear Leadership Program Graduate Fellowship. He has held internship positions at Idaho National Laboratory and Radiant Industries, Inc. Noah was married in February 2023 while completing this thesis.

ACKNOWLEDGEMENTS

Dr. Yousry Azmy has been my advisor since late 2019 and has helped guide my academic progress and research work tremendously, and he deserves my utmost thanks for his support through the years. I would additionally like to thank Dr. Dmitriy Anisratov and Dr. Semyon Tsynkov for taking the time to serve on my committee. Dr. Raffi Yessayan, now at Los Alamos National Laboratory, was my first student mentor as a sophomore while he was completing his doctoral dissertation. He provided a thorough background of the theory and use of the THOR discrete ordinates code utilized in this work. The assistance of Dr. Nicholas Herring, who has recently taken up a position at The University of Texas at Austin, has proven invaluable thanks to his improvements and documentation of THOR as well as his insights regarding previous work done in the area of coupling between Monte Carlo and discrete ordinates solution schemes. My friends and colleagues in the Department of Nuclear Engineering also deserve my thanks for the camaraderie and countless hours of study sessions we have shared throughout our academic careers. Finally, I would like to give thanks to my wife, Brittany, for never once failing to accompany, support, and believe in me.

This material is based upon work supported under a Department of Energy, Office of Nuclear Energy University Nuclear Leadership Program Graduate Fellowship.

TABLE OF CONTENTS

List of Tables	vii
List of Figures	viii
Chapter 1 Introduction	1
1.1 Motivation	1
1.2 Problems of Interest	2
1.3 Overview	4
Chapter 2 Background	5
2.1 Neutron Transport Theory	5
2.2 Computational Solution Methods	7
2.2.1 Monte Carlo	7
2.2.2 Discrete Ordinates	10
2.3 Ray Effects	14
2.3.1 Overview	14
2.3.2 Review of Mitigation Techniques	16
Chapter 3 MCNP-THOR Coupling Scheme	21
3.1 Introduction	21
3.2 Review of Existing Monte Carlo Coupling Schemes	22
3.2.1 DOMINO	22
3.2.2 CADIS	22
3.2.3 First Collision Source Method	23
3.2.4 General Collision-Count Source Method	25
3.3 MCNP	27
3.3.1 FMESH Cartesian Mesh Overlay	28
3.4 THOR	29
3.4.1 THOR Meshes	30
3.4.2 Scattering Order	31
3.5 Coupling Scheme	32
3.5.1 Centroidal Mapping	32
3.5.2 A Note on Mesh Agreement	34
3.5.3 Cross Section Application	35
3.5.4 Induced Fission	36
3.5.5 Summing Fluxes & Detector Response Calculation	37
3.5.6 Postprocessing	38
Chapter 4 Experiment & Modeling	39
4.1 Experiment	39
4.1.1 BeRP Ball Description	39
4.1.2 Experimental Configuration	41
4.2 Models	43

4.2.1	MCNP Models	43
4.2.2	THOR Models	45
4.2.3	Helium Modeling	48
Chapter 5	Results & Validation	50
5.1	Introduction	50
5.1.1	Explanation of Important Model & Running Parameters	51
5.1.2	Postprocessing	52
5.2	Unshielded Case	52
5.2.1	Uncoupled S_N Solution	53
5.2.2	Unoptimized Coupling Case	58
5.2.3	Optimized Coupling Case	63
5.3	Shielded Case	71
Chapter 6	Conclusions	77
6.1	General Remarks	77
6.2	Proposals for Future Work	78
6.2.1	Parallelization & Runtime Improvements	78
6.2.2	Alternative Mapping Techniques	79
6.2.3	Optimization of MC Collision Count	80
6.2.4	Full Validation	80
6.2.5	Investigation of Higher Quadrature Orders	81
6.2.6	Analysis of Mesh Similarity Requirements	81
6.2.7	Implementation Using OpenMC	82

LIST OF TABLES

Table 2.1	Description of the terms in Eq. 2.1.	6
Table 4.1	Estimated isotopic composition of the plutonium metal comprising the BeRP ball on 6 January 2009. Copied here from Ref. [31].	40
Table 5.1	Raw detector responses of the uncoupled problem at the various quadrature orders.	57

LIST OF FIGURES

Figure 2.1	Cross-sectional view of a heatmap of flux computed with an S_2 quadrature calculation of k_{eff} for the Godiva benchmark in a vacuum.	15
Figure 4.1	Experimental configuration of the BeRP ball and detectors on August 3, 2017. Copied here from Ref. [6].	41
Figure 4.2	Simplified CAD model of the experimental setup showing the BeRP ball (left) and SNAP detector (right). The shielded configuration is shown. Copied here from Ref. [6]	44
Figure 4.3	Annotated view of the simplified model of the unshielded case as used in THOR. The view is a slice through the middle of both the BeRP ball and detector. Mesh boundaries are hidden.	46
Figure 4.4	Enhanced view of a cross section of the unshielded detector model with visible mesh boundaries.	46
Figure 4.5	The simplified model of the shielded case as used in THOR. The view is a slice through the middle of both the BeRP ball and detector. Mesh boundaries are hidden.	47
Figure 4.6	Enhanced view of a cross section of the shielded detector model with visible mesh boundaries.	48
Figure 4.7	ENDF/B-VIII.0 total cross section for He-3 (barns) [33].	49
Figure 4.8	ENDF/B-VIII.0 (n, p) cross section for He-3 (barns) [33].	49
Figure 5.1	Heatmap of the flux output from THOR in the uncoupled test problem using S_2 quadrature, with the plane taken at the center of both the source and detector.	54
Figure 5.2	Flux output from THOR in the uncoupled test problem using increasing quadrature orders. The view is down the problem's midplane. The first figure uses an S_4 quadrature set. The second uses S_8 . The third uses S_{16}	55
Figure 5.4	Cross-sectional heatmap of MCNP flux using the uniform <i>FMESH</i> grid. Both the THOR mesh and <i>FMESH</i> grid are overlaid.	59
Figure 5.5	First collision total source (summed over energy groups) found using the unoptimized grid for coupling to THOR.	60
Figure 5.6	Detector response (CPS) computed using P_0 , P_1 , and P_3 scattering cross sections in THOR with the unoptimized <i>FMESH</i> grid. MCNP reference solution provided for comparison [6].	61
Figure 5.7	Optimized <i>FMESH</i> grid (white gridlines) overlaid on the geometry of the unshielded case.	63
Figure 5.8	First collision source found using the better-optimized grid for coupling to THOR.	65
Figure 5.9	Detector response (CPS) computed using P_0 , P_1 , and P_3 scattering cross sections in THOR with the optimized <i>FMESH</i> grid. The continuous energy MCNP reference solution is provided for comparison [6].	66
Figure 5.10	Line plots of the flux of neutrons having any number of collisions down the centerline of the problem. The left set of blue lines encompasses the BeRP ball, while the right set contains the cadmium-to-cadmium width of the detector.	68

Figure 5.11 Comparison of the n -collided to uncollided MCNP flux remaining in the BeRP ball, indicating a mitigation of ray effects with increasing collision counts.	70
Figure 5.12 Average relative error of thermal flux in the <i>FMESH</i> cells making up the detector.	71
Figure 5.13 <i>FMESH</i> grid overlaid on the geometry of the shielded problem.	72
Figure 5.14 First collision total source found for the shielded case.	73
Figure 5.15 Detector response (CPS) computed using P_0 , P_1 , and P_3 scattering cross sections in THOR for the shielded case. MCNP reference solution provided for comparison [6].	74

CHAPTER 1

INTRODUCTION

1.1 Motivation

Efficient numerical simulation of neutron transport is of interest for many areas in the field of nuclear science and engineering. While whole-core light water reactor modeling can use the deterministic diffusion approximation to save on computational costs, its many limitations of applicability necessitate the use of transport-based methods such as Monte Carlo and deterministic schemes in other applications. While accurate, Monte Carlo simulations, which rely on the direct simulation of neutron histories, can prove prohibitively expensive for problems having large domains and/or unlikely events, such as shielding and whole-facility dose computations [1]. On the other hand, deterministic methods such as the discrete ordinates method incur errors related to spatial, angular, and energy discretization in order to preserve computational efficiency, but they can exactly solve for neutron flux over the entire discretized domain in phase space [2].

One such error in the discrete ordinates method's solution is known as ray effects. Ray effects are observed in problems with strongly-localized neutron sources. Insufficient

scattering in the problem's materials can result in discrete rays emanating from the localized source. They manifest as non-physical oscillations in the spatial distribution of the neutron flux. The rays of high flux travel uninterrupted from the localized source along the set of discretized angles, as insufficient scattering prevents them from spreading out into nearby spatially-discretized cells [3].

The computational efficiency of the discrete ordinates method when compared to Monte Carlo then produces a desire to mitigate the troublesome ray effects. Numerous techniques have been investigated to varying success. One such technique involves the coupling of Monte Carlo and discrete ordinates codes. By using Monte Carlo to accurately simulate neutrons undergoing up to a given number of collisions, a more broadly-distributed fixed source may be supplied to the discrete ordinates calculation. Distributing the source in this manner improves the probability that areas of interest in the problem are properly illuminated by neutron flux, since rays in standard S_N can either miss entire regions of a problem, or, conversely, intersect areas of interest and overpredict a computed response. The method was first explored by computing only the uncollided neutron flux in MC, and it was shown to reduce ray effects in two-dimensional problems [4]. It was expanded to allow for any number of MC collisions to be used to generate the S_N source, successfully mitigating ray effects in three-dimensional problems [5][6]. This work applies such an MC- S_N coupling scheme between the widely-used MCNP Monte Carlo code [7] and the open source THOR discrete ordinates code developed at North Carolina State University [8][9] in order to demonstrate the broad applicability of the method.

1.2 Problems of Interest

Problems having both significant shielding and heavily-localized sources often characterize nonproliferation applications. Source-detector type problems are of particular

interest, as special nuclear material (SNM), especially plutonium, emits significant quantities of gamma and neutron radiation. In pursuit of mitigating nuclear security risk, detection of such radiation signatures from SNM is of paramount interest [10]. Modeling and simulation of source-detector problems can therefore inform the development and applicability of radiation detectors which are valuable for nuclear nonproliferation efforts.

Such problems, however, are intuitively prone to encountering the drawbacks of both the Monte Carlo and discrete ordinates simulation methods. Cases having significant shielding of SNM to mask radiation signatures can hamper the Monte Carlo method, as particles which successfully penetrate a shield are rare, and those proceeding to reach and successfully interact with a detector are even more so. Numerous variance reduction techniques exist in Monte Carlo to improve the statistical confidence in the estimated quantities in such cases, but the suitability of such methods can be strongly problem-dependent and require significant tuning by the user [11]. Additionally, discrete sources in the form of lumped masses of SNM cause source-detector problems to be prone to ray effects in discrete ordinates calculations, since the air medium separating the source and detector in such problems is low-scattering.

Both solution schemes to the neutron transport equation then appear inadequate for problems of great interest to nuclear nonproliferation applications. Indeed, for two source-detector problems analyzed in this work, ray effects are severe, with the standard S_N method completely failing to converge to a consistent solution with increasing N for N between 2–16 as seen in Table 5.1. At quadrature sets beyond S_{16} , the discrete ordinates method, itself favored over Monte Carlo for its computational efficiency, becomes prohibitively expensive. Therefore, the coupling scheme between MCNP and THOR shall prove particularly valuable for localized source systems such as source-detector problems in the area of nuclear nonproliferation.

1.3 Overview

The remainder of this thesis is split into five chapters. In Chapter 2, a background of neutron transport theory and descriptions of the Monte Carlo and discrete ordinates computational solution schemes are presented. The erroneous ray effects phenomenon, present in the discrete ordinates method, is also introduced in greater detail along with a review of existing mitigation schemes intended to improve accuracy of the discrete ordinates solutions. Chapter 3 begins with a discussion of existing coupling schemes between Monte Carlo and discrete ordinates codes before detailing our novel coupling scheme between MCNP and THOR. This coupling is intended to reduce the impact of ray effects in THOR calculations. In Chapter 4, descriptions are provided of a series of experiments performed in 2017 and of MCNP and THOR models used to simulate these experiments. These models provide computational exercises which are susceptible to severe ray effects, the demonstrable elimination of which serves to validate the coupling scheme. Chapter 5 demonstrates the severity of ray effects in these problems as well as the improved results obtained upon applying the coupling scheme. The scheme demonstrates a successful mitigation of ray effects for both problems of interest considered in this work. Finally, in Chapter 6, conclusions regarding the efficacy and applicability of the general collision-count source method as applied in this work are drawn, and proposals for future work to further improve the method are presented.

CHAPTER 2

BACKGROUND

2.1 Neutron Transport Theory

The neutron transport equation is a linear integro-differential equation with independent variables generally in space and time and in the neutron's energy and angle of travel [12]. In full detail, it takes the following form:

$$\begin{aligned} & \frac{1}{v(E)} \frac{\partial \psi(\mathbf{r}, E, \boldsymbol{\Omega}, t)}{\partial t} + \boldsymbol{\Omega} \cdot \nabla \psi(\mathbf{r}, E, \boldsymbol{\Omega}, t) + \Sigma_t(\mathbf{r}, E) \psi(\mathbf{r}, E, \boldsymbol{\Omega}, t) \\ &= \int_{4\pi} d\boldsymbol{\Omega}' \int_0^\infty dE' \Sigma_s(\mathbf{r}, E' \rightarrow E, \boldsymbol{\Omega}' \rightarrow \boldsymbol{\Omega}) \psi(\mathbf{r}, E', \boldsymbol{\Omega}', t) \\ &+ \frac{\chi(E)}{4\pi} \int_{4\pi} d\boldsymbol{\Omega}' \int_0^\infty dE' v_f(\mathbf{r}, E') \Sigma_f(\mathbf{r}, E') \psi(\mathbf{r}, E', \boldsymbol{\Omega}', t) + s(\mathbf{r}, E, \boldsymbol{\Omega}, t). \end{aligned} \quad (2.1)$$

The equation describes the local balance of the neutron population, $\frac{\psi(\mathbf{r}, E, \boldsymbol{\Omega}, t)}{v(E)}$, throughout the phase space comprised of space (\mathbf{r}), energy (E), direction of motion ($\boldsymbol{\Omega}$), and time (t). Σ is used to represent macroscopic cross sections in this thesis. Consult Table 2.1 for a brief description of the terms appearing in Eq. 2.1. Fission is assumed to emit particles isotropically in this formulation.

Table 2.1: Description of the terms in Eq. 2.1.

Term	Description
$\frac{1}{v(E)} \frac{\partial \psi(\mathbf{r}, E, \boldsymbol{\Omega}, t)}{\partial t}$	Change Rate
$\boldsymbol{\Omega} \cdot \nabla \psi(\mathbf{r}, E, \boldsymbol{\Omega}, t)$	Streaming
$\Sigma_t(\mathbf{r}, E) \psi(\mathbf{r}, E, \boldsymbol{\Omega}, t)$	Total Collisions
$\int_{4\pi} d\boldsymbol{\Omega}' \int_0^\infty dE' \Sigma_s(\mathbf{r}, E' \rightarrow E, \boldsymbol{\Omega}' \rightarrow \boldsymbol{\Omega}) \psi(\mathbf{r}, E', \boldsymbol{\Omega}', t)$	Scattering
$\frac{\chi(E)}{4\pi} \int_{4\pi} d\boldsymbol{\Omega}' \int_0^\infty dE' v_f(\mathbf{r}, E') \Sigma_f(\mathbf{r}, E') \psi(\mathbf{r}, E', \boldsymbol{\Omega}', t)$	Fission Production
$s(\mathbf{r}, E, \boldsymbol{\Omega}, t)$	External Source

In most applications not involving rapid temporal changes in the neutron flux, the neutron transport equation is recast in a steady-state form [13]:

$$\begin{aligned}
 & \boldsymbol{\Omega} \cdot \nabla \psi(\mathbf{r}, E, \boldsymbol{\Omega}) + \Sigma_t(\mathbf{r}, E) \psi(\mathbf{r}, E, \boldsymbol{\Omega}) \\
 &= \int_{4\pi} d\boldsymbol{\Omega}' \int_0^\infty dE' \Sigma_s(\mathbf{r}, E' \rightarrow E, \boldsymbol{\Omega}' \rightarrow \boldsymbol{\Omega}) \psi(\mathbf{r}, E', \boldsymbol{\Omega}') \\
 &+ \frac{\chi(E)}{4\pi} \int_{4\pi} d\boldsymbol{\Omega}' \int_0^\infty dE' v_f(\mathbf{r}, E') \Sigma_f(\mathbf{r}, E') \psi(\mathbf{r}, E', \boldsymbol{\Omega}') + s(\mathbf{r}, E, \boldsymbol{\Omega}).
 \end{aligned} \tag{2.2}$$

This work considers the time-independent Eq. 2.2 as its focus, since the problems of interest in the validation exercise are steady-state configurations. From here, we may conveniently represent the equation using operator notation:

$$L\psi(\mathbf{r}, E, \boldsymbol{\Omega}) = P\psi(\mathbf{r}, E', \boldsymbol{\Omega}') + s(\mathbf{r}, E, \boldsymbol{\Omega}), \tag{2.3}$$

where L represents particle streaming and all particle collisions (the loss term), and P represents particle production through scattering and fission.

The time-independent, three-dimensional neutron transport equation, Eq. 2.2, does not possess an analytical solution except for unrealistically simplified configurations. The complexity of its operators and the dependence of the variable ψ on space, energy, and angle precludes such solutions in practical applications. However, when considering Eq. 2.3, one may notice that it represents a fundamentally simple relationship — that of a local balance between particle loss and particle gain. So although it cannot be solved

analytically, several computational schemes exist to numerically evaluate solutions to the neutron transport equation.

2.2 Computational Solution Methods

Broadly, the computational methods employed to find numerical solutions to Eq. 2.2 may be separated into two classes: stochastic and deterministic methods. Monte Carlo, a stochastic method, simulates and tracks individual particle histories in a continuous domain. It uses pseudorandom number generators to model the interactions these particles experience based on corresponding probability distributions, which are quantified by the applicable macroscopic cross sections [11]. The average behavior in a system can then be inferred by aggregating the histories of the particles simulated. Deterministic methods are a broader class. Among the popular deterministic methods, which solve directly for quantities such as flux using some form of discretization in the independent variables, is the discrete ordinates, or S_N , method.

The Monte Carlo method is discussed in Section 2.2.1, and the discrete ordinates method is introduced and formulated in Section 2.2.2. We then dedicate the remainder of this chapter to a detailed description of the ray effects phenomenon found in the S_N method, its cause, and existing methods of ray effects mitigation.

2.2.1 Monte Carlo

The Monte Carlo (MC) approach to solving the neutron transport equation relies on a probabilistic description of particle interactions. Consider, for example, the many possibilities for interactions which a neutron encounters as it travels. It may stream, fail to collide with any other particles, and leak out of the system under consideration. Or, it may undergo a collision after traveling some distance which is influenced by the mean free path in the medium. When undergoing a collision, the collision may be a

scattering or absorption event depending on the probability of each interaction type in the material in which the neutron is traveling. The particle's energy and direction of motion changes stochastically within allowable limits following a scattering event. It may either be captured or induce fission, if the material so allows, in an absorption event. The number of fission neutrons born in a fission event is also probabilistic.

Concisely, almost all aspects of the history of a neutron from birth to ultimate absorption or leakage from the system permit probabilistic descriptions. The probability that an event has a given outcome is described by a probability density function, or PDF [11]. The PDF and its associated cumulative distribution function, or CDF, may be used to evaluate the outcomes of events using random numbers. In a simple example, a fair coin toss, the coin has a 50% chance to land on either face. The CDF of this event may be defined such that, for an input $0 \leq x \leq 0.5$, the coin lands on heads, while for an input $0.5 < x \leq 1$, the coin lands on tails. Using a pseudorandom number generator with a uniform distribution between 0 and 1, this function could be sampled to simulate the outcome of the coin toss. This simple principle can be applied to the prediction of particle interactions, as individual particle histories may be simulated using the exact same sort of sampling, just with more complex CDFs. The location of a particle's birth, its starting energy and initial direction of motion, and the probability of undergoing various interactions with different outcomes can thus all be modeled using the MC method.

Some mention should be made of pseudorandom number generators (PRNGs) themselves. Random numbers are required to model phenomena of a stochastic nature as described. Therefore, for MC calculations, a computer must be able to generate a random string of numbers. This process is not as trivial as it appears, since only quantum computers [14] and random physical phenomena such as radioactive decay [15] can generate truly random numbers. PRNGs are therefore the current standard used for MC calculations. They are designed with the goal of algorithmically generating sequences of numbers which have distributions that are statistically indistinguishable from truly

random numbers. So, with the PRNG taken as a black box — that is, with the generation algorithm unknown to the MC application — and assuming sufficient "randomness" in the algorithm itself, pseudorandom number generators are suitable for MC calculations [15]. Additionally, one notable advantage of pseudorandom number sequences over truly random numbers is the ability to repeat the same sequence by supplying the PRNG with a seed, which facilitates debugging and verification efforts in MC codes by eliminating variations in the strings of numbers generated for different executions.

Of course, when considering the results of a numerical experiment, a single evaluation of a particle's history cannot give an accurate prediction of the outcome of an event. Intuitively, the history of a single neutron tracked using MC in no way represents the behavior of all neutrons in a system under consideration. Therefore, the fundamental challenge of using MC for radiation transport problems lies in ensuring accurate estimations of the aggregate behavior of particles in the system based on finitely-many tracked particle histories. MC accuracy suffers particularly for problems having complex geometry and/or improbable occurrences [1]. For example, accurate estimations of dose in shielding problems using MC can prove prohibitively expensive or even impossible without employing variance reduction (VR) techniques to decrease statistical uncertainty [16]. Accordingly, for problems which are ill-suited for analog MC, that is, MC using no VR, either very large numbers of particles must be simulated or VR methods tailored to the studied configuration must be employed. Both of these approaches are expensive, can require trial-and-error to reduce variance to acceptable levels, and can not with certainty always alleviate the problem at hand. Therefore, while MC is generally regarded as an excellent method for estimating the solution to the exact neutron transport equation, other solution schemes such as S_N can prove more efficient for many applications.

2.2.2 Discrete Ordinates

Theory & Applications

MC, using pseudorandom numbers, can operate on continuous variables down to the level of machine precision, but it incurs statistical error proportional to the number and weight of particles reaching different areas of the problem domain. The S_N method, on the other hand, discretizes the angular variable Ω to be able to find exact numerical solutions to the discretized form of the neutron transport equation [2]. That is, error in MC arises from the ability or inability of the histories of simulated neutrons to accurately describe the average behavior of neutrons in the system, but the problem domain is exactly represented. Error in S_N arises from the discretization of space, angle, and energy in the problem domain, but the solution to the discretized problem can be exact. Therefore, where MC yields an estimate of the solution to the exact neutron transport equation Eq. 2.2, S_N and other deterministic methods provide exact solutions to approximations of the transport equation.

In S_N , finitely-many discrete angles are chosen along which the angular flux ψ is computed. The resulting equations couple ψ appearing in the loss terms to ψ in the production terms, the left and right hand sides of Eq. 2.2, respectively. Typically, the solution scheme involves successive iterations of the computation of ψ . The production operator P in Eq. 2.3 operates on the former iteration of ψ , while the loss operator L operates on the current iteration. Once the scalar flux $\phi(\mathbf{r}, E) = \int_{4\pi} d\Omega \psi(\mathbf{r}, E, \Omega)$ changes little between iterations, the solution is considered converged. Of course, false convergence can occur in cases involving slow convergence, such as in problem configurations having very high-scattering media, as the change in $\phi(\mathbf{r}, E)$ from one iteration to the next can be sufficiently small to erroneously trigger a solver's convergence check [17]. Quadrature sets, which are predefined sets of ordinates and respective angular weights, are used to numerically compute integrals, such as those involved in evaluating secondary particle

production.

The S_N method proves valuable in solving problems with complex domains and/or unlikely events, which are weaknesses of the MC method. This is because, during the iterative solution process, each cell in the domain is swept over. Rather than requiring individual particle histories to adequately illuminate the entire problem domain, flux in each cell and along each discrete ordinate is solved for exactly, with factors such as quadrature order and spatial mesh refinement influencing accuracy of the solution. For this reason, S_N is well-suited for problems having large domains, significant shielding, or other properties which would render analog MC unsuitable.

Formulation of the S_N Method

At the heart of the S_N formulation is the construction of appropriate quadrature sets. It has been noted that physical symmetry and the arrangement of discrete directions on latitudes on the unit sphere are crucial elements in the development of quadrature sets [18]. A description of a fully-symmetric quadrature set, in which rotational symmetry across all eight octants of the unit sphere is assumed, will be presented here referencing [13] and [18]. However, alternative quadrature sets with varying degrees of rotational symmetry exist.

The direction unit vector Ω may point in any direction. Consider a unit sphere having surface coordinates defined by (μ, η, ζ) , the cosines of Ω . For a quadrature rule of order N , there are $\frac{N}{2}$ points placed along each positive axis of (μ, η, ζ) . When each of these points is projected out to the unit sphere, the number of intersections per octant is then given as $\frac{N(N+2)}{8}$. The set of vectors pointing from the origin to these points on the unit sphere define the set of discrete ordinates in an octant. Rotated about the remaining octants, these are then the set of angles along which the S_N method solves for neutron flux in each spatially-discretized cell.

Due to the symmetry requirements in this formulation, there is in fact only one degree

of freedom for a given angular cosine. The recursive relationship between the remaining discrete ordinates in an octant is given by:

$$u_n^2 = u_1^2 + (n - 1) \frac{2(1 - 3u_1^2)}{N - 2}, \quad i = 1, 2, \dots, \frac{N}{2}. \quad (2.4)$$

That is, of the $\frac{N}{2}$ points along one axis, only the first can be chosen freely. Additionally, with complete rotational invariance, the points along each axis will be identical, meaning the choice of a single angular cosine determines all others in the set.

Weights for the points defined by the quadrature rule must additionally be chosen. They are normalized,

$$\sum_{n=1}^{N(N+2)/8} w_n = 1, \quad (2.5)$$

and their formulation, the details of which shall prove extraneous for this description, may be found in [18].

Once the quadrature rule is defined by a set of discrete ordinates and weights, the neutron transport equation Eq. 2.2 may be discretized in angle [13]. That is, a set of equations known as the discrete ordinates equations are derived from Eq. 2.2, given in general geometry:

$$\mathbf{\Omega}_n \cdot \nabla \psi_n(\mathbf{r}, E) + \Sigma_t(\mathbf{r}, E) \psi_n(\mathbf{r}, E) = Q_n(\mathbf{r}, E), \quad (2.6)$$

where $Q_n(\mathbf{r}, E)$ includes the scattering, fission, and fixed source terms. If we assume isotropic scattering and fission production, an isotropically-emitting source $S(\mathbf{r}, E)$ which includes both the fixed source and fission production, and constrain the discrete ordinates equations in Cartesian geometry, we receive:

$$\mu_n \frac{\partial \psi_n(\mathbf{r}, E)}{\partial x} + \eta_n \frac{\partial \psi_n(\mathbf{r}, E)}{\partial y} + \zeta_n \frac{\partial \psi_n(\mathbf{r}, E)}{\partial z} + \Sigma_t(\mathbf{r}, E) \psi_n(\mathbf{r}, E) = Q(\mathbf{r}, E), \quad (2.7)$$

$$Q(\mathbf{r}, E) = \int_0^\infty dE' \frac{\Sigma_s(\mathbf{r}, E' \rightarrow E)}{8} \sum_{n=1}^{N(N+2)} w_n \psi_n(\mathbf{r}, E') + S(\mathbf{r}, E),$$

where we can see that the angular integral in Eq. 2.2 has been replaced by a sum over the quadrature angles using their appropriate weights. The continuous variables \mathbf{r} and E may now be discretized. Numerous spatial discretization schemes exist, including the diamond difference and step relationships, allowing the derivative terms to be computed numerically by using differences in neutron flux across the cells in each dimension. Energy may be discretized into any desired energy group structure using averaged cross sections.

An iterative solution scheme may then yield the angular flux ψ_n along the discrete ordinates in the discretized spatial cells and energy groups. This scheme solves for ψ on the left hand side of equation Eq. 2.6 in terms of a known right hand side:

$$\mathbf{\Omega}_n \cdot \nabla \psi_n^{\ell+1}(\mathbf{r}, E) + \Sigma_t(\mathbf{r}, E) \psi_n^{\ell+1}(\mathbf{r}, E) = Q_n^\ell(\mathbf{r}, E), \quad (2.8)$$

where this generalized $Q_n(\mathbf{r}, E)$ permits anisotropy when compared to the definition given in Eq. 2.7. By repeatedly solving this set of equations beginning with Q_n^0 generated from an initial guess ψ_n^0 , a value of the scalar flux ϕ may be found which varies little from iteration to iteration. Within a certain tolerance for this variation, then, an approximate solution to Eq. 2.8 may be found. This solution, which, within the specified tolerance, is the exact solution to the discrete ordinates equation, is then an approximate solution to Eq. 2.2 to within the discretization error. In more general formulations, neither scattering production nor the fixed source necessarily need to be isotropic, but appropriate scattering cross sections must be available for the treatment of anisotropy.

2.3 Ray Effects

2.3.1 Overview

First discovered in the mid-1960s, ray effects in the discrete ordinates formulation, as well as their cause and some simple mitigation schemes, were discussed by Lathrop in 1968 [3]. Lathrop explains the origin of ray effects in the steady-state S_N equations (see Eq. 2.6) by comparing to the steady-state continuous neutron transport equation (see Eq. 2.2). These two equations are identical except for the discretization in angle used in the S_N formulation. This angular discretization, which the S_N method uses to directly compute flux, causes oscillations of non-physically high and low flux to form in the absence of significant scattering.

As an example, for a purely-absorbing medium or a vacuum with a discrete point source in space, we expect a flux distribution which is symmetrically distributed about the source and which monotonically attenuates proportionally to the square of the distance. Equation 2.2 has such a solution, which a Monte Carlo simulation of such a problem would also produce with some statistical error. However, consider the spatially-discretized cells used in the S_N method which are far from the source. In such cells, the discrete angles along which flux is computed may not coincide with a ray emanating from the point source. Without scattering (in a pure absorber or vacuum), flux will continue "traveling" along these rays. Accordingly, the cells which do not coincide with the discrete ordinates emanating from the source will not be illuminated at all with flux. Only those cells with quadrature angles pointing directly to the source will have a non-zero solution, which is overestimated in pursuit of maintaining the neutron balance. These are the rays after which the ray effects phenomenon is named.

Figure 2.1 shows an example case with severe ray effects. The problem is a k_{eff} calculation for the Godiva criticality benchmark [19], a model of which is included in THOR's modeling tutorial [9]. To modify the problem in order to display the ray effects, a

200cm×200cm×200cm cube of vacuum is placed around the fissile material. The modified problem employs an S_2 quadrature set to maximize the visibility of the rays. Half of the problem domain is visible in Figure 2.1, where the wide surface is the geometry's midplane passing through the center of the sphere and perpendicular to the cube's sides. The figure displays a heatmap of the fundamental eigenmode (flux), with darker red indicating a very high flux, and darker blue indicating a very low flux. The elements of the sphere have an expected flux distribution, while the solution in the vacuum is dominated by ray effects.

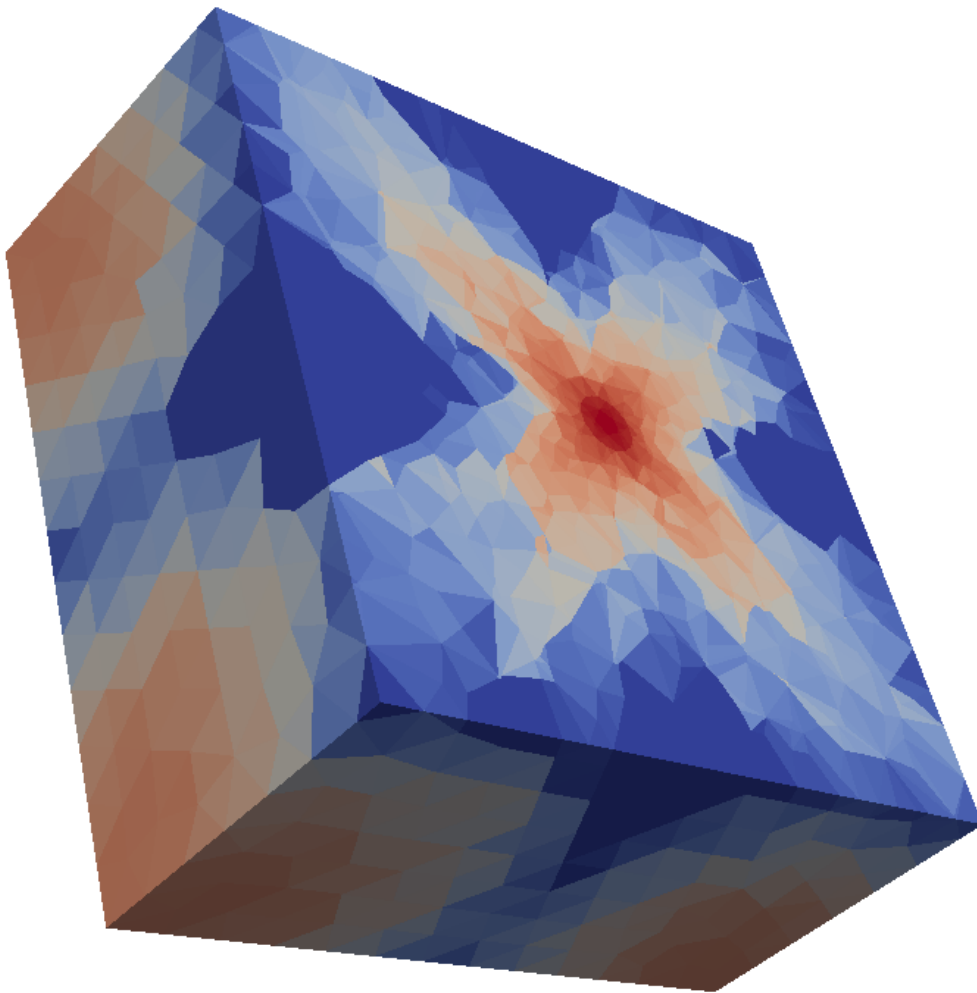


Figure 2.1: Cross-sectional view of a heatmap of flux computed with an S_2 quadrature calculation of k_{eff} for the Godiva benchmark in a vacuum.

2.3.2 Review of Mitigation Techniques

Numerous techniques exist to reduce the impact of ray effects in the S_N method. Of these mitigation methods, we discuss: raising the quadrature order, P_{N-1} equivalence, modulated P_{N-1} equivalence, artificial scattering, and the rotated S_N method. The benefits and drawbacks of these mitigation schemes are presented. Existing Monte Carlo coupling methods — the basis of the remainder of this work — are then discussed, and comparison to these other mitigation techniques will demonstrate their value.

Raising Quadrature Order

Ray effects arise from the angular discretization of the neutron transport equation. If few angles are evaluated, many cells are likely to fail to be adequately illuminated by angular intersection with a localized source, making ray effects more pronounced. As the number of discrete ordinates approaches infinity, the solution approaches that of the continuous-angle transport equation, and ray effects disappear.

However, raising the quadrature order does not fundamentally solve the issue of ray effects, primarily due to the additional computational expense incurred. Since each additional angle results in an additional equation of the form Eq. 2.6 to be iteratively solved in each spatial cell, raising the quadrature order can dramatically increase computational time. And, as noted by Lathrop [3] comparing S_2 and S_4 quadrature sets, nonphysical ray effects are still very apparent — if muted — by increasing the quadrature order within reasonable ranges. Therefore, more rigorous, less expensive methods of ray effects mitigation are sought-after.

P_{N-1} Equivalence

It is additionally noted in [3] that the discrete ordinates equations are equivalent to the spherical harmonics (P_N) equations of order P_{N-1} with Mark boundary conditions for plane geometry when using a Gaussian-Legendre quadrature set. These equations

do not display ray effects. In more general geometry, however, the equivalence between the S_N and P_{N-1} equations is necessarily not retained. In S_2 , this is due to the presence of an additional moment in the S_2 equations when compared to the moments of the P_1 equations:

$$Q = 3 \sum_{n=1}^4 w_n \mu_n \eta_n \psi_n. \quad (2.9)$$

This term may be eliminated by modifying the S_2 equations:

$$\mu_n \frac{\partial \psi_n}{\partial x} + \eta_n \frac{\partial \psi_n}{\partial y} + \Sigma_t \psi_n = Q_n + \eta_n \frac{\partial Q}{\partial x} + \mu_n \frac{\partial Q}{\partial y}, \quad n = 1, 2, 3, 4, \quad (2.10)$$

resulting in an equation having moments equivalent to the P_1 equations and results which match the diffusion solution with no ray effects in a numerical experiment. However, for higher-order S_N -to- P_{N-1} equivalence, a new treatment is required for each N , which may require special quadrature sets. That is, S_N -to- P_{N-1} equivalence is analytically intensive and may not be possible for all N in all geometric configurations.

Modulated P_{N-1} Equivalence

A modulated P_1 -equivalent S_2 method was examined in [20]. The iterative method is formulated in two-dimensional Cartesian geometry with isotropic scattering in the following equation:

$$\mathbf{A}_d \frac{\partial \vec{\psi}^{\ell+1}}{\partial x} + \mathbf{B}_d \frac{\partial \vec{\psi}^{\ell+1}}{\partial y} + \Sigma_t \vec{\psi}^{\ell+1} = \Sigma_s \vec{\phi}^\ell - \mathbf{A}_r \frac{\partial \vec{\psi}^\ell}{\partial x} - \mathbf{B}_r \frac{\partial \vec{\psi}^\ell}{\partial y} + \vec{Q}, \quad (2.11)$$

where $\vec{\psi}^\ell$ indicates a vector of the angular flux on four different direction cosines for iteration ℓ , $\vec{\phi}^\ell$ indicates the scalar flux, \vec{Q} is a vector of source components also along these direction cosines, and the coefficient matrices are given:

$$\begin{aligned}
\mathbf{A}_d &= \frac{1}{\sqrt{3}} \begin{bmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, & \mathbf{A}_r &= \frac{\alpha}{4\sqrt{3}} \begin{bmatrix} 1 & -1 & -1 & 1 \\ 1 & -1 & -1 & 1 \\ -1 & 1 & 1 & -1 \\ -1 & 1 & 1 & -1 \end{bmatrix}, \\
\mathbf{B}_d &= \frac{1}{\sqrt{3}} \begin{bmatrix} -1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, & \mathbf{B}_r &= \frac{\alpha}{4\sqrt{3}} \begin{bmatrix} 1 & -1 & -1 & 1 \\ -1 & 1 & 1 & -1 \\ 1 & -1 & -1 & 1 \\ -1 & 1 & 1 & -1 \end{bmatrix}.
\end{aligned} \tag{2.12}$$

For $\alpha = 0$, the standard S_2 equations are obtained. For $\alpha = 1$, the P_1 -equivalent S_2 equations are obtained. By modulating $\alpha \in [0, 1]$, a balance can be achieved between the ray effects-prone but efficient standard S_2 method and the ray effects-free but slow-to-converge P_1 -equivalent S_2 method. It was found that the modulated P_1 -equivalent S_2 method presented in [20] reduces ray effects in several test cases by raising α near 0.9–0.99, but that such a method's efficacy is highly problem-dependent.

Artificial Scattering

Another proposed method of addressing ray effects, called artificial scattering S_N or as - S_N , was proposed in [21]. The right hand side of the transport equation is modified by adding an artificial scattering term,

$$\begin{aligned}
\frac{\partial \psi(\mathbf{r}, \boldsymbol{\Omega}, t)}{\partial t} + \boldsymbol{\Omega} \cdot \nabla \psi(\mathbf{r}, \boldsymbol{\Omega}, t) + \Sigma_t(\mathbf{r})\psi(\mathbf{r}, \boldsymbol{\Omega}, t) \\
= \Sigma_s(\mathbf{r})(S^+ \psi)(\mathbf{r}, \boldsymbol{\Omega}, t) + \Sigma_{as}(\mathbf{r})(S_{as} \psi)(\mathbf{r}, \boldsymbol{\Omega}, t) + q(\mathbf{r}, t),
\end{aligned} \tag{2.13}$$

where $(S^+ \psi)$ is the in-group scattering kernel operator which describes the gain of particles from all directions into direction $\boldsymbol{\Omega}$. The artificial scattering term is defined:

$$\Sigma_{as}(\mathbf{r})(S_{as}\psi)(\mathbf{r}, \mathbf{\Omega}, t) = \Sigma_{as}(\mathbf{r}) \int_{S^2} s_\epsilon(\mathbf{\Omega}' \cdot \mathbf{\Omega}) (\psi(\mathbf{r}, \mathbf{\Omega}', t) - \psi(\mathbf{r}, \mathbf{\Omega}, t)) d\Omega', \quad (2.14)$$

where

$$s_\epsilon(\mu) = \frac{2}{\sqrt{\pi}\epsilon \text{Erf}(\frac{2}{\epsilon})} e^{-(1-\mu)^2/\epsilon^2}, \quad (2.15)$$

$$\epsilon = \beta/N_q. \quad (2.16)$$

β is a user-defined constant and N_q is the number of ordinates. As $N_q \rightarrow \infty$, as - S_N converges to S_N .

The authors note similarities to the class of S_N -to- P_{N-1} equivalence methods, with the artificial scattering term in place of alterations to the moments such as that found in Eq. 2.9. The as - S_N method was found to have dramatically lower error than standard S_N calculations, particularly for lower N . For a line-source problem, noted by Lathrop [3] as particularly prone to ray effects, as - S_2 outperformed standard S_4 . This dramatically improves on memory usage, but the added complexity of the method does not lead to substantial runtime improvements. This is especially apparent with higher N , where standard S_N can execute in shorter runtime but with higher memory consumption when compared to as - S_N .

Rotated S_N

The rotated S_N , or rS_N , method was recently proposed both to avoid the potential for negative particle densities observed in P_N methods and to mitigate ray effects in the standard S_N method [22]. The angular quadrature set is rotated after every time step, and an interpolation procedure is used to find the solution on the new ordinates. This effectively mimics the addition of quadrature points without the added computational

cost. The authors demonstrate consistency with the standard S_N equations as the angular grid is refined, while also noting that the rotation and interpolation processes add diffusivity to the method.

In two dimensions, the rS_N method reduces ray effects in the commonly-cited line-source test. The standard S_N method very poorly approximates the anticipated rotationally-symmetric solution of an isotropically-emitting line source in a vacuum. rS_N is able to outperform standard S_N with less than one third the number of quadrature points for this problem. The method also reduces ray effects in the so-called lattice problem, where there is not rotational symmetry and the system is heterogeneous.

However, rS_N , as presented in [22], does have limitations. In two dimensions, the ability of rS_N to mitigate ray effects is limited due to only having one available axis of rotation. Its efficacy is dependent on a compatible quadrature set. Finally, the method is not fully compatible with steady-state transport sweeps and requires modification of the underlying code of S_N solvers.

CHAPTER 3

MCNP-THOR COUPLING SCHEME

3.1 Introduction

We have discussed the theoretical background, uses, and limitations of the MC method in Section 2.2.1. This probabilistic solution method is inherently immune to ray effects, but it fails for problems having large, complex domains and when modeling unlikely phenomena. As discussed in Sections 2.2.2 and 2.3.1, the S_N method has the opposite benefits and drawbacks. Naturally, then, there have been numerous attempts to couple MC and S_N to enjoy the advantages of both methods.

The remainder of this work will focus on such coupling methods. In Section 3.2, we will briefly review existing MC- S_N applications and their influence on this work. In Section 3.3, the notable MC code MCNP is discussed. THOR, an open-source S_N code developed at North Carolina State University, is introduced in Section 3.4. Finally, in Section 3.5, a new collision-based coupling scheme between MCNP and THOR is formulated. Numerical results and a demonstration of ray effects mitigation when using the method are presented later in Chapter 5.

3.2 Review of Existing Monte Carlo Coupling Schemes

Because the collision-based MCNP-THOR coupling scheme was heavily influenced by existing methods, the theory and some examples of MC- S_N coupling will provide a foundation for understanding the remainder of this work.

3.2.1 DOMINO

The earliest examples of MC- S_N coupling were envisioned primarily as a way to improve upon existing analog Monte Carlo programs. DOMINO, the Discrete Ordinates Monte Carlo Interface Operation, is an early example [23]. Developed at Oak Ridge National Laboratory, it bridged the DOT two-dimensional S_N code and the MORSE MC code. Angular fluxes from DOT generated probability distributions for use by MORSE as either a source or scoring function for either forward or adjoint runs. Whichever solution method proved most amenable to the problem could then be chosen. The style of coupling used in DOMINO resembles in many ways modern MC variance reduction techniques, such as the automated variance reduction presented in [16].

3.2.2 CADIS

A more modern technique of MC- S_N coupling uses the S_N adjoint function to automatically calculate variance reduction parameters for MC. This method, referred to as the Consistent Adjoint Driven Importance Sampling (CADIS) technique, has been developed for applications for which analog MC proves unsuitable, and generally delivers significant computational efficiency [24].

CADIS applies the knowledge that the adjoint solution to the transport equation, easily computed with S_N codes, quantifies the anticipated contribution of particles born anywhere in the phase space covered by the problem's configuration to a desired response function. A biased source distribution is then computed from the adjoint solution, allowing

for MC source sampling to be proportional to the expected contribution to the response function. This source biasing improves upon analog MC by favoring the simulation of particles which are likely to contribute to the response which is being measured. The weight-window method, also defined with the adjoint solution and consistent with the source biasing, is used to additionally speed up MC calculations. This scheme has proven most suitable for optimization of the calculation of the response in a single region, such as a single detector.

An extension of CADIS, called Forward Weighted CADIS or FW-CADIS, uses a typical forward S_N calculation to estimate the response. An adjoint function is then computed, informed by the forward S_N results, and it is used in the same way as standard CADIS to improve the efficiency of the MC run. By first running forward S_N , the method is able to produce a more uniform MC particle density throughout the entire problem domain, making FW-CADIS as a coupling method very valuable for large problems with multiple localized tally regions or significant shielding, such as whole-facility dose calculations, a problem investigated by the authors of [24].

3.2.3 First Collision Source Method

The aforementioned MC- S_N coupling schemes focused on improving statistics in and shortening MC runtimes by using deterministic calculations to produce sources or to inform source biasing. However, the opposite goal — that is, mitigating ray effects in S_N calculations by using MC as the starting point — has been investigated through collision-based coupling approaches. By using S_N for the bulk of the computations, its enhanced efficiency over MC can be taken advantage of.

Alcouffe first proposed a method of mitigating ray effects in localized source problems in [4]. We may split the angular flux into uncollided (ψ_u) and collided (ψ_c) components such that the total angular flux $\psi = \psi_u + \psi_c$. Consider the transport equation Eq. 2.1. It may be applied equally for uncollided neutrons and for collided neutrons. However,

since an uncollided neutron by definition becomes a collided neutron after a scattering event, the scattering term is dropped for the uncollided flux and instead is found in the equation for collided flux, and we receive two equations. For the uncollided flux:

$$\frac{1}{v(E)} \frac{\partial \psi_u}{\partial t} + \mathbf{\Omega} \cdot \nabla \psi_u + \Sigma_t(\mathbf{r}, E) \psi_u = Q(\mathbf{r}, E, \mathbf{\Omega}, t), \quad (3.1)$$

the source term Q is the sole source of uncollided neutrons. For the collided flux, given by:

$$\begin{aligned} \frac{1}{v(E)} \frac{\partial \psi_c}{\partial t} + \mathbf{\Omega} \cdot \nabla \psi_c + \Sigma_t(\mathbf{r}, E) \psi_c = \\ = \int_{4\pi} d\mathbf{\Omega}' \int_0^\infty dE' \Sigma_s(\mathbf{r}, E' \rightarrow E, \mathbf{\Omega}' \rightarrow \mathbf{\Omega}) \psi(\mathbf{r}, E', \mathbf{\Omega}', t), \end{aligned} \quad (3.2)$$

scattering always results in collided neutron production, hence the ψ term with no subscript. In Alcouffe's implementation of this first collision source method, MC particles were killed upon making their first collision. Then, all the uncollided particles were tallied to produce ψ_u , and a source — referred to here as a first collision source — for the S_N problem was produced. Once the S_N calculation found ψ_c , the two could be added together to produce the solution ψ . He found a significant, but not complete, reduction of ray effects using the first collision method, with the first collision source S_4 method producing better results than typical S_{16} for a point source problem.

This coupling scheme addresses one of the two major contributing factors which produce ray effects in S_N solutions, namely, a heavily localized source. By producing a more distributed source for the ψ_c calculation, the S_N method has a better chance at properly illuminating discrete spatial cells which do not have discrete ordinates tracing from the cell to the localized source.

3.2.4 General Collision-Count Source Method

Introduction

With the reduction in ray effects demonstrated by Alcouffe using the first collision source method, one may question whether allowing particles to survive for one or more collisions would provide further benefit. After all, in larger problems with optically thick high-scattering regions surrounding a localized source, the uncollided neutrons may still very well yield a relatively localized source when computing the response in a distant location of interest. Herring in [6] and [5] investigated a general collision-count — referred to also as an n -collided — source method for MC- S_N coupling using the Hammer MC code [25] and THOR S_N code [8]. Hereafter, this scheme shall be referred to as the "general collision-count source method," or GCCSM.

Formulation

Hammer is a multigroup, object-oriented MC transport code with the ability to tally on unstructured tetrahedral grids [25]. The code is tailored to support ray effect mitigation in THOR, which uses such a grid for spatial discretization. The same cross sections may be utilized by the two codes for a seamless compatibility. With the production of n -collided sources as a key goal for Hammer, it has the ability to kill particles after an arbitrary, user-specified collision count, and it can tally particles by collision count. Herring [6] extended Alcouffe's derivation of the first collision source method by splitting the flux by particle collision count. Note that while Herring's formalism used angular fluxes, in practice, scalar fluxes have been used in the implementations of the GCCSM. So, the notation will be adapted here. For a first collision source, the scalar flux is split according to the collision count,

$$\phi = \sum_{i=0}^{\infty} \phi^{(i)}, \quad (3.3)$$

$$L\phi^{(0)} = Q, \quad (3.4)$$

$$L\phi^{(i)} = S\phi^{(i-1)}, \quad i > 0, \quad (3.5)$$

where the superscript (i) indicates flux of particles having collision count i , and (L, S, Q) are the loss operator, secondary production operator, and fixed source, respectively, similar to the operators found in Eq. 2.3. For an n -collided source, then, a modification of the first collision source method yields the following equations:

$$\phi_{MC} = \sum_{i=0}^{n-1} \phi^{(i)}, \quad (3.6)$$

$$L\phi^{(n)} = Q_n = S\phi^{(n-1)}, \quad (3.7)$$

$$\phi_{n+} = \sum_{i=n}^{\infty} \phi^{(i)}, \quad (3.8)$$

such that $\phi = \phi_{MC} + \phi_{n+}$. In this case, ϕ_{MC} is the scalar flux computed by Hammer having collision counts less than n , making it the flux from all neutrons simulated in MC. ϕ_{n+} is the scalar flux of neutrons having collision counts of n or more. It is computed in THOR by using Q_n in the same way as a typical fixed source. Therefore, $(n - 1)$ -collided MC flux can generate the source used to compute ϕ_{n+} in the S_N calculation, and for $n = 1$, the method is seen to be equivalent to Alcouffe's.

Ray Effects Mitigation

For a plutonium sphere in air test problem, Herring found a demonstrable reduction in visible ray effects with increasing the MC collision count of neutrons used for generating the S_N source [6]. Using pure S_N , severe ray effects persist for this problem with a quadra-

ture order as high as S_{16} . Additionally, the first collision source method in combination with an S_2 quadrature set, while improving upon standard S_2 , still has notable ray effects; this possibility was noted in [4] due to the inability to directly illuminate large patches of cells caused by killing the flux upon first collision.

For each increase in the collision count of MC particles used to generate the n -collision source, however, the flux solution further approaches the anticipated flux distribution, with rays in the domain's corners eventually disappearing upon reaching a fifth-collided source, which is generated using neutron flux from particles having four MC collisions. That is, from purely visual results, S_2 in THOR using a fifth collided source from Hammer outperforms standard S_{16} for a test problem which is highly susceptible to ray effects [6].

The success of the GCCSM in mitigating ray effects serves as the primary motivation for this work. The method, while dramatically improving accuracy of the S_N calculation, was limited in applicability due to the reliance on the unique compatibility between Hammer and THOR. The two codes use the same meshes, energy group structures, and even multigroup cross sections. If the GCCSM can so improve upon the first collision source method, then applying it with a more widely-available code such as MCNP and with less reliance upon identical meshes serves as a logical next step. The remainder of this thesis will therefore focus on a new implementation of the GCCSM with a coupling scheme between MCNP and THOR. The goal of this work is to demonstrate a broader applicability and simpler implementation when compared to the Hammer-THOR coupling scheme discussed in [5] and [6] as well as the robustness of the GCCSM itself.

3.3 MCNP

MCNP is an MC code for particle transport, including neutron transport [7]. It uses Constructive Solid Geometry (CSG) to represent materials, surfaces, and volumes in space. It is one of the most-utilized particle transport codes and is typically considered the gold

standard for neutron transport applications. Due to its broad availability, confidence in its accuracy, and wide range of features and useful options including the *FMESH* Cartesian mesh overlay, it was chosen as the candidate code of choice for the new implementation of the GCCSM. Specifically, MCNP6.2 is the version used.

3.3.1 *FMESH* Cartesian Mesh Overlay

MCNP uses CSG to generate the geometry within which particle simulations of the type described in Section 2.2.1 are performed. The geometric description is functionally continuous, so particles can be found at any point in space. Accordingly, by default, there is no spatially-discretized mesh for events to be tallied on as is the case in S_N codes such as THOR, where the angular neutron flux is computed per-cell, for each angle, for each energy group, and even at each time step if applicable. However, for the GCCSM to be applied, there must be a way to use the neutron fluxes generated by MCNP to find a next-collision source for each cell in the THOR model.

Fortunately, MCNP has the *FMESH* card. Using *FMESH*, a user can overlay a Cartesian grid on top of the continuous geometry description. This grid need not be uniform, and it can have varied refinement according to the demands of the problem and objectives of the calculation. For example, areas of significant importance to the overall response function or areas in which the flux shape rapidly changes can have a higher refinement than areas of little concern. *FMESH* outputs scalar flux values and relative statistical errors in each cell using track-length tallying by default. Therefore, the particles traveling within the continuous domain are tallied on a discretized grid, allowing for coupling to S_N .

The *FMESH* card is not limited to spatial discretization. Although MCNP, by default, operates using a continuous energy spectrum for neutron simulations, *FMESH* can tally fluxes within user-defined energy group bins at the end of execution. This allows for compatibility with THOR's multigroup energy structure, as MCNP fluxes can be binned in groups with the same energy boundaries as those used in the THOR calculation.

This capability enables multiplication by THOR's multigroup scattering cross sections to produce fixed sources for the S_N calculation. When compared to Hammer, which uses the same multigroup cross sections as THOR throughout the MC calculation, less error from energy discretization may be expected from MCNP, since its particles are given a continuous energy treatment prior to binning.

Finally, the *FMESH* card has the capability to tally fluxes of neutrons having a specified collision count. A user may choose to tally either neutrons having a single collision count or a range of counts. For an $(n - 1)$ -collided flux MCNP output, an n -collided source may be produced for the THOR calculation. The $[0, n - 1]$ -collided flux may also be tallied for later addition when computing a total, fully collided response value. The *FMESH* card in MCNP therefore enables a straightforward implementation of the GCCSM using a user-specified Cartesian mesh and energy group structure.

3.4 THOR

THOR is a discrete ordinates transport code using unstructured tetrahedral meshes and based on the Arbitrarily High Order Transport Method of the Characteristics type, or AHOT-C [8]. The AHOT-C method is a method of short characteristics solver. It can utilize high-order scattering cross sections for modeling of anisotropic scattering [26]. The unique implementation of AHOT-C for unstructured grids is referred to as AHOT-C-UG and is described in [27]. THOR is under development by the Nuclear Computational Science Group at North Carolina State University (NCSU-NCSG). THOR version 1.0.2 is the version used [9].

THOR is capable of solving fixed source and eigenvalue problems. For eigenvalue calculations, it can apply power iteration and Jacobian-Free Newton Krylov (JFNK) solvers. THOR has parallel execution capabilities and iterative accelerators primarily implemented by Yessayan and as described in [28]. These capabilities are designed to

allow THOR to rapidly perform transport calculations on complex geometries amenable to discretization on unstructured, tetrahedral meshes.

3.4.1 THOR Meshes

In order to represent complex geometries accurately and to allow potential interoperability with multiphysics codes, THOR uses unstructured grids with tetrahedral elements [8] [27]. The tetrahedral elements are better able to represent curved surfaces and geometric boundaries than a Cartesian grid. This is because points functioning as the tetrahedra's vertices, also called nodes, can be placed in any location in space so long as there are no gaps between elements. This is particularly beneficial at curved material boundaries where cross sections rapidly change, where cells in a Cartesian grid may be very likely to significantly overlap both materials and require a corrective treatment ("smearing") of cross sections. Since the grid is not structured, tetrahedra may vary in shape, size, and orientation as deemed desirable by the meshing software of choice and which may be influenced by the user. Therefore, the unstructured mesh can allow for finer elements, oriented to avoid significant material overlap, to be placed along surfaces.

The tetrahedral mesh defining the geometry of a THOR problem, though, presents challenges with regards to compatibility with programs which do not share this mesh format. While Hammer made direct use of THOR's *.thrm* mesh files, MCNP cannot tally fluxes on the THOR mesh. A mapping from MCNP's *FMESH* Cartesian grid to THOR's unstructured tetrahedral grid must then be formulated in order to utilize the GCCSM. The two grids must additionally be similar enough in refinement to avoid losing important data during the mapping process. The mapping procedure developed for the purposes of this work is explained in Section 3.5.1, and details regarding sufficient mesh similarity are provided in Section 3.5.2. As we will show in Section 5.2, the mapping from MCNP to THOR, while not exact as in the case of Hammer to THOR, requires only minimal tuning to produce accurate results.

3.4.2 Scattering Order

The AHOT-C method, as discussed above and as its name suggests, utilizes arbitrarily high order spatial moments of the angular flux to improve accuracy [26]. The notable effects of higher scattering orders on the coupling scheme warrant some discussion. THOR uses scattering expansion orders to treat anisotropic flux [8]. These take the form of so-called P_N cross sections. Each N has its own scattering matrix in the THOR cross section format [9]. P_0 represents a fully isotropic approximation of scattering, while higher N improves anisotropic modeling at the cost of greater computational expense.

Continuous energy MCNP, on the other hand, models scattering differently. The cosine of the angle before and after a scattering collision is sampled from distribution tables in MCNP's cross section library using pseudorandom numbers [7]. Depending on the neutron's energy, elasticity of the collision, data available, and treatments such as the application of $S(\alpha, \beta)$ data for moderators [29], varying degrees of anisotropy are modeled in MCNP. This is one of the advantages of MC in terms of accuracy of solutions. Since the *FMESH* tally of MCNP can only output scalar flux, however, only the P_0 scattering cross sections from THOR are applied during the conversion to a source term in the coupling process; this is illustrated in Eq. 3.7 in Section 3.5.3. Inherently, then, an isotropic approximation is made for the flux at the n^{th} collision.

The question of continuing the THOR run following the isotropic approximation presents several options. If a higher N scattering order is chosen, then the isotropic approximation at the n^{th} collision represents a single collision of inaccuracy. Alternatively, THOR can carry on using P_0 scattering for computational efficiency, terminating anisotropic treatment at the n^{th} collision. Analyses of results in Chapter 5 include runs using multiple different THOR scattering orders to compare effects and to demonstrate convergence of the different solutions at high MCNP collision counts.

3.5 Coupling Scheme

Having explained the meshes used by MCNP for tallying the MC results and THOR for solving the S_N equations as well as the theory behind the GCCSM, we may now explain how coupling between MCNP and THOR has been performed. First, what we dub the centroidal mapping scheme developed for this project is detailed. The value of agreement — that is, similarity in refinement — between MCNP and THOR meshes is then introduced and will be expanded upon in Section 5.2.2. The process for conversion from MCNP fluxes to THOR sources is explained. This is followed by a brief discussion of induced fission, why it is disabled by default when applying our method, and how it is accounted for. The procedure for recombining fluxes from MCNP and THOR and calculating the total detector response is presented. Finally, some postprocessing steps are described.

3.5.1 Centroidal Mapping

As discussed in Section 3.4.1, the THOR mesh is comprised of tetrahedra of varying sizes in an unstructured grid. All information about a tetrahedral mesh element is given by its nodes, or vertices, that connect in pairs to form edges and in triplets to form triangular faces. In comparison, the flux tallied using MCNP's *FMESH* card is ordered in a structured Cartesian grid, with elements which are not necessarily uniform in size. When representing the same geometry, then, the two grids inherently will not align. Some tetrahedra (hereafter referred to as "THOR elements") may be fully enclosed within a Cartesian grid cell ("MCNP cells"), and vice-versa. A choice must then be made regarding how to select which MCNP cell's flux is used to compute the source for a given THOR element.

A simple method of mapping was developed which relies only on using the centroids of the THOR elements. The coordinates of an element's centroid may be computed by

averaging the coordinates of its four composing nodes:

$$C_x = \frac{N_x^1 + N_x^2 + N_x^3 + N_x^4}{4}, \quad (3.9)$$

$$C_y = \frac{N_y^1 + N_y^2 + N_y^3 + N_y^4}{4}, \quad (3.10)$$

$$C_z = \frac{N_z^1 + N_z^2 + N_z^3 + N_z^4}{4}, \quad (3.11)$$

where a letter subscript indicates the coordinate point (x, y, z) , C denotes the coordinates of an element's centroid, and N denotes the coordinates of each of the four nodes of the THOR element. Once the centroid of an element has been found, a sweep across MCNP's Cartesian grid can locate which MCNP cell "contains" the centroid. Then, the $(n - 1)$ -collided flux can be taken from the MCNP cell and copied to the respective THOR element. This value is then used to compute the n -collision source following application of the appropriate cross sections for the THOR element as described in Section 3.5.3.

Pause should be given to note that the flux which is computed by MCNP and mapped to THOR represents only a heatmap of its shape for the problem used in the validation of this method. This behavior can be enforced or adjusted by altering the particle weight in MCNP, using or omitting the *NONU* card, and/or performing postprocessing steps. There are several multiplicative correction terms used in this validation exercise which are discussed in Section 3.5.6.

We also note that two separate output files are prepared at this stage. Both contain fluxes mapped to the THOR mesh. The first contains the n -collided source to be used in the subsequent, coupled THOR calculation. This process is detailed in Section 3.5.3. The second file contains the uncollided through $(n - 1)$ -collided flux — see ϕ_{MC} in Eq. 3.6. The scalar fluxes in this file are summed together with those obtained following the THOR run in a post-processing step described in Section 3.5.5.

Alternative mapping schemes from Cartesian to tetrahedral cells certainly exist and may provide advantages over this centroidal method. One proposed method has been

to average the flux value found in the MCNP cell containing the centroid of a THOR element with the values found in the cells containing the element's nodes. This could potentially better-resolve areas of rapid change, such as material boundaries. It could also provide benefits for large cells of lower importance, such as the air cells in the model shown in Section 4.2. In this case, smoothness of the solution would be better-preserved by averaging over the large area. Finally, in cases of poorly-tuned Cartesian grids, it could help avoid mapping identical MCNP fluxes to different THOR elements in the event that multiple elements' centroids are contained within the same MCNP cell. However, the search through the Cartesian grid is computationally expensive, and with unknown benefits, the simpler mapping scheme, which shall be shown to be sufficient for accurate calculations, was chosen for use in this work.

3.5.2 A Note on Mesh Agreement

The fundamental assumption of the centroidal mapping scheme presented here is that, on average, the flux found at a THOR element's centroid is approximately representative of the actual flux if THOR's mesh were to be tallied upon directly. Of course, this is only an approximation. If an MCNP *FMESH* grid is too coarse, multiple THOR elements might be assigned the same flux despite, for example, being found on opposite ends of a material interface with a large change in actual flux magnitude. Alternatively, if the *FMESH* grid is too fine, the THOR elements may be assigned a value which is only accurate for a small, localized region of the larger tetrahedron, resulting in a misrepresentation of average behavior in the element.

Therefore, there must be a degree of agreement between the *FMESH* grid and THOR's unstructured mesh despite their inherent differences. As will be demonstrated in Section 5.2.3, it has been found that certain factors related to grid refinement prove beneficial for solution fidelity: having similar sizes of THOR elements and MCNP cells and having similar numbers of both. That is, if for each THOR element there exists approximately

one matching MCNP cell of a roughly similar size at a nearby location, the error introduced by mapping is reduced when compared to significantly different refinement scales between the two grids, an intuitive conclusion. It is therefore of benefit to expend some effort performing variable refinement on the *FMESH* grid in order to more closely match the refinement of the THOR mesh in areas of significance to the response calculation. Likewise, reducing the refinement in areas of little importance, such as in air gaps, can reduce the size of output files and the time expended performing the coupling procedure. As our results in Section 5.2.2 demonstrate, however, even a minimally-optimized *FMESH* grid may yield accurate results using the coupling scheme with particles having low MCNP collision counts.

3.5.3 Cross Section Application

Consider Eq. 3.7, in which the n -collided THOR source Q_n is generated from the $(n - 1)$ -collided MCNP flux. In order to produce the source from MCNP's $(n - 1)$ -collided scalar flux once it has been mapped to THOR's elements, the isotropic P_0 scattering matrix in THOR's multigroup cross sections must be applied on a per-element basis:

$$Q_n = \begin{bmatrix} \Sigma_{s,1 \rightarrow 1} & \Sigma_{s,2 \rightarrow 1} & \dots & \Sigma_{s,G \rightarrow 1} \\ \Sigma_{s,1 \rightarrow 2} & \Sigma_{s,2 \rightarrow 2} & \dots & \Sigma_{s,G \rightarrow 2} \\ \dots & & & \\ \Sigma_{s,1 \rightarrow G} & \Sigma_{s,2 \rightarrow G} & \dots & \Sigma_{s,G \rightarrow G} \end{bmatrix}_{P_0} \begin{bmatrix} \phi_1 \\ \phi_2 \\ \dots \\ \phi_G \end{bmatrix}_{n-1} \quad (3.12)$$

The n -collided source Q_n can be seen to consist of the neutrons having $(n - 1)$ collisions which subsequently undergo a scattering collision in their given THOR elements. The ϕ_{n+} computed by THOR then represents the continuation of these particles' histories in the S_N calculation.

3.5.4 Induced Fission

Ray effects mitigation through collision-based coupling in the case of a heavily localized source can be hampered if that source can undergo induced fission. This is because, if left enabled, a large magnitude of fast flux is produced in the fissionable source. In the S_N code, then, this new, localized source of fast neutrons can overwhelm the benefits provided by the GCCSM, and the overall S_N source will remain heavily localized rather than well-distributed. Herring found that, with induced fission enabled for a model of a spherical plutonium source, using the GCCSM still drastically underpredicts count rates due to ray effects [6]. The plutonium source in question is discussed in greater detail in Section 4.1.1 and is the source for the problems of interest for validation of the work presented in this thesis.

Since many localized sources are indeed composed of fissionable material, a simple procedure is used to combat this limitation. First, the problem can be run in MC with induced fission enabled, as is the default behavior of MCNP. Then, it is run again with induced fission disabled, which is made possible with the *NONU* input card. This card effectively zeros out the ν_f fission neutron output term, resulting in gross zero neutron production from induced fission. The response of interest may be compared between the two runs, and a multiplicative factor may be computed:

$$F = \frac{R_{induced}}{R_{NONU}}. \quad (3.13)$$

Of course, in practical applications where the response of interest is not amenable to being computed via MCNP calculations, an alternative method of computing F may be devised, for example by simplifying the configuration and calculating the magnitude of the flux in the source region alone, once without the *NONU* card and once with it. The ratio of the results would give an indication of the scaling factor for the response between runs with induced fission enabled and disabled. It was found in [6] that, by disabling induced

fission and using this correction factor generated by comparison of the detector responses computed in MCNP, ray effects can be mitigated properly by the GCCSM. The application of this factor is discussed in Section 3.5.6. Of course, when transitioning over to the S_N portion of the computation, induced fission must also be disabled for consistency. THOR therefore also has a *NONU* card available which is capable of preventing fission production [9].

3.5.5 Summing Fluxes & Detector Response Calculation

Having produced n -collided sources in each element of THOR's mesh, THOR may compute the n to fully-collided fluxes (see ϕ_{n+} in Eq. 3.8) using the S_N equations. As discussed in Section 3.5.1, a file is output during the mapping procedure containing the ϕ_{MC} fluxes as assigned to each THOR element. Therefore, the overall group scalar fluxes throughout the mesh can be computed according to Eq. 3.3.

For the purposes of the validation problems investigated in this work, for which the predicted count rate from a detector is the desired response, this addition is performed automatically in THOR's detector response calculator, a postprocessing utility. It can read multiple flux output files and sum them together on a per-element basis [9], allowing for the addition of the ϕ_{MC} and ϕ_{n+} terms within THOR's elements. It then applies user-supplied cross sections in the volume of interest to compute a reaction rate. The detector response in counts per second (CPS) is given:

$$R_D = \sum_{r \in D} V_r \sum_g \Sigma_{r,g} \phi_{r,g} \quad (3.14)$$

where R_D is the total response in the detector, V_r is the volume of a single element r in the detector's active region D , $\Sigma_{r,g}$ is the detector's response function for a single element r and energy group g , and $\phi_{r,g}$ is the scalar flux in element r in group g . We note that MCNP and THOR output scalar fluxes, not angular fluxes, for use in such reaction rate

computations, so no summation over angle is required.

3.5.6 Postprocessing

Once the raw detector response has been computed, it must be adjusted using correction terms to yield the proper magnitude of the response, in this case the detector's count rate. Taking R_{raw} as the raw detector response with no correction terms, using a unit source magnitude in MCNP, with induced fission disabled, we can find the overall detector response as:

$$R = R_{raw} \cdot S \cdot F \cdot D, \quad (3.15)$$

where S is the true source strength, F is the multiplication factor for induced fission discussed in Section 3.5.4, and D is the density factor of the active detector region's material as computed in THOR. S is known and may be omitted if the MCNP source is given a magnitude equal to S . Alternatively, both the particle weight in MCNP and S may be kept at 1 in order to compute a detector response on a per-neutral-particle basis. Since S is a scaling factor that is dependent on experimental results and applied to every single computed response, it need only be consistently maintained during comparisons. D is computed automatically by THOR using a ratio of a user-supplied material region volume, which may be determined analytically, and the total volume of mesh elements representing the same material region in order to account for meshing inaccuracies. The density of each THOR element's material is multiplied by the appropriate D in THOR to equalize the amount of material in the meshed and true geometries. Since the detector response calculator does not innately compute and apply this factor in the same way THOR itself does, it must be supplied by the user during postprocessing.

CHAPTER 4

EXPERIMENT & MODELING

4.1 Experiment

The validation of the proposed extension to the GCCSM is informed by an experiment performed using a small, subcritical plutonium source and a SNAP neutron detector. With an air gap between the two, such a problem is heavily susceptible to severe ray effects in S_N computations due to the localized nature of the source. A description of the source itself, thoroughly studied as a benchmark problem, is provided, followed by a description of the experiment against which the method is validated.

4.1.1 BeRP Ball Description

The subcritical plutonium source is referred to as the Beryllium Reflected Plutonium ball, or BeRP ball, because the original experiments using the metallic source were performed with beryllium reflection [30]. It has been used for numerous benchmark experiments and detector validation measurements since its fabrication in 1980.

Table 4.1: Estimated isotopic composition of the plutonium metal comprising the BeRP ball on 6 January 2009. Copied here from Ref. [31].

Constituent	Mass Fraction	Constituent	Mass Fraction
Pu238	0.016%	Mo98	2.2 ppm
Pu239	93.265%	Mo96	1.5 ppm
Pu240	5.907%	Zn66	1.4 ppm
Pu241	0.068%	Mo92	1.3 ppm
Pu242	0.028%	Cd111	1.3 ppm
Am241	2472.4 ppm	Cd110	1.2 ppm
Ga	335 ppm	Cd113	1.2 ppm
C	230 ppm	Pb206	1.2 ppm
U235	746.9 ppm	Pb207	1.1 ppm
U236	174.1 ppm	Be	1 ppm
Zr	100 ppm	B	1 ppm
Np237	77.8 ppm	Mg	1 ppm
Na	50 ppm	Cu	1 ppm
U234	39.2 ppm	Ag	1 ppm
Fe	10 ppm	Bi	1 ppm
Al	5 ppm	Zn68	0.97 ppm
Si	5 ppm	Mo100	0.90 ppm
Ni	5 ppm	Mo97	0.87 ppm
Sn	5 ppm	Mo94	0.81 ppm
Cr52	4.2 ppm	Cd116	0.77 ppm
Ca40	2.9 ppm	Cr53	0.48 ppm
Cd114	2.9 ppm	Cr50	0.21 ppm
Pb208	2.6 ppm	Zn67	0.21 ppm
Zn64	2.4 ppm	Cr54	0.12 ppm
Cd112	2.4 ppm	Cd106	0.12 ppm
		Balance	2860.3 ppm

Mattingly in [31] performed an extensive characterization of the BeRP ball. The BeRP ball is a 4483.884g sphere of α -phase plutonium clad in stainless steel 304. The plutonium sphere has a volume of 228.72cm³, having a computed density of 19.604g/cm³. The stainless steel cladding is 0.0305cm thick with a density of 7.62g/cm³. The gap between it and the plutonium has a maximum width of 0.0680cm. Several different reflector materials have been used for experiments with the BeRP ball despite its name. In a series of benchmarking experiments performed by Mattingly in 2009, several layers of high density polyethylene (HDPE) were used as reflectors [31]. Mattingly estimated the isotopic composition of the plutonium metal on 6 January 2009, the first day of these

experiments. The table detailing the estimated composition of the metal is included here as Table 4.1.

4.1.2 Experimental Configuration

A series of experiments using the BeRP ball which were performed on August 3, 2017 are described in [6]. The experiments were intended to generate data for the validation of both THOR and the GCCSM, an MC- S_N coupling scheme, presented in that work. A picture of the experimental setup is given in Figure 4.1. Besides the steel cladding, the BeRP ball is bare, with no layers of HDPE or other reflectors, and is placed on an aluminum stand.



Figure 4.1: Experimental configuration of the BeRP ball and detectors on August 3, 2017. Copied here from Ref. [6].

Two neutron detectors, a SNAP detector and an MC-15 multiplicity counter, are placed on opposite sides of the BeRP ball, the former 1m and the latter 0.5m from the center of the ball. The SNAP detector counts (n, p) reactions in an active He-3 tube. Upon reaching

the detector, existing thermal neutrons are filtered out by a thin cadmium absorber. Fast neutrons pass through this absorber and are then thermalized by a layer of HDPE in order to improve the detector sensitivity due to the higher (n, p) cross sections at thermal energies as discussed in Section 4.2.3. This configuration is designed to maximize the detection of fission neutrons, which exit the BeRP ball fast, and to minimize detection of already-thermal neutrons. Both the SNAP and MC-15 detectors are described in greater detail in [30] and [31]. The latter notes the removable HDPE slab on the SNAP detector as having a width of 1in and a density of $0.962\text{g}/\text{cm}^3$. Its inclusion leads to a preference for the SNAP detector to count only the highest-energy neutrons. Moderate-energy neutrons are sufficiently thermalized by the extra HDPE in the slab to be absorbed in the cadmium, such that only neutrons with a very high incident energy make it past the cadmium to go on to be detected in the helium region.

Additional experiments were performed in 2017 using various configurations of shells of HDPE, tungsten, and copper surrounding the BeRP ball as reflectors, the details of which may be found in [6]. However, the case of the bare BeRP ball is used for validation of this implementation of the GCCSM. This decision is justified by that fact that, as will be discussed in Section 4.2, the MCNP and THOR models developed for this validation exercise are already simplifications of the experimental setup presented here. They do not include many details such as the stand, table, or floor. Therefore, the inclusion of the reflector shells proves extraneous for the task of demonstrating a mitigation of ray effects in susceptible problems such as this, as perfectly accurate-to-experiment numerical results are not actually sought. Instead, preference is given to thorough study of ray effects mitigation with cases involving the bare sphere, reducing computational expenses and minimizing additional variables which could unnecessarily impact our analysis.

4.2 Models

Models of the experimental setup presented in Section 4.1.2 were developed to validate the GCCSM presented in [5] and [6]. A series of MCNP models was created, which progressively simplified the problem to eventually contain only the bare BeRP ball and SNAP detector (with and without the HDPE slab) to facilitate development of a discretely-meshed THOR model and to accelerate analysis. These MCNP models were validated against the experimental results in [6]. Hereafter, the configuration without the removable HDPE slab placed on the front of the detector is referred to as the "unshielded" case, while the configuration with the HDPE slab is referred to as the "shielded" case.

The original GCCSM utilized MC tallies from the Hammer code on the same unstructured meshes as THOR, meaning the THOR meshes corresponding to the most simplified problem configurations were created. In the following sections, descriptions of the nine MCNP models representing the various simplifications to the experimental setup will briefly be presented. Then, the simplified THOR tetrahedral meshes of the same problem configurations, both of the unshielded and shielded cases, will be shown. In Chapter 5, the results computed with MCNP-THOR coupling using these models will be compared to pure MC results, considered the reference solutions for the problems, in order to prove the efficacy of the new implementation of the GCCSM at mitigating ray effects in S_N calculations.

4.2.1 MCNP Models

A series of nine MCNP models were created in support of [5] and [6]. The first of these accurately modeled almost every aspect of the experimental configuration shown in Figure 4.1. Full details are given for the SNAP detector, the BeRP ball including its steel shell and ring, and the table, stand, and even the concrete floor. However, the MC-15 detector and the walls and ceiling in the room are excluded.

Successive iterations of the MCNP model removed details such as the tripod holding the SNAP detector, the floor, the steel shell of the BeRP ball, and the table and stand. In the final, fully-simplified model, only a bare plutonium metal sphere representing the BeRP ball and the simplified SNAP detector remain with an air gap between them. The removable HDPE slab is modeled as an inch-long extension of the detector's main HDPE body, which may be toggled between air (unshielded case) and HDPE (shielded case). The model of the simplified configuration as designed in the SolidWorks computer aided design (CAD) program is shown in Figure 4.2. This simplification of the experiment discussed in Section 4.1.2 serves as the basis for the validation work presented in this thesis with an *FMESH* grid overlaid on it in MCNP. When compared to the MCNP analysis performed in [6], higher particle counts and source directional biasing were implemented to improve statistics on the *FMESH* grid, the importance of which is discussed in greater detail in Section 5.2.




Figure 4.2: Simplified CAD model of the experimental setup showing the BeRP ball (left) and SNAP detector (right). The shielded configuration is shown. Copied here from Ref. [6]

Herring [6] found a count rate difference in the unshielded SNAP detector between the full MCNP model and simplified MCNP model of less than 5%, falling from 43.8 CPS to 41.9 CPS. Compared to the 14.7%–17.4% (shielded and unshielded case, respectively) error introduced in the fully-detailed model when compared to experimental results, the simplifications were accepted as having negligible impact while allowing for significantly greater ease in meshing for THOR.

We note that the MCNP results for the computed detector responses were therefore taken as the reference values, rather than the experimentally-determined count rates from the detector. As such, the new implementation of the GCCSM proves beneficial by removing the Hammer step entirely. The results presented in Chapter 5 indicate that, by using MCNP itself to couple with THOR, detector response values which are very close to those predicted by full MCNP runs are found with the coupling scheme, even when using very few particle collisions in MCNP.

4.2.2 THOR Models

The simplified version of the MCNP model was taken as the basis for creating the THOR model. This is primarily due to the difficulties associated with meshing extremely fine details, especially for THOR’s unstructured, tetrahedral grid, which requires elements to connect by face and to be properly oriented, unlike, for example, a Cartesian grid overlaid on CSG. An annotated cross-sectional visualization of the THOR mesh generated from the CAD shown in Figure 4.2 (without the HDPE slab) can be found in Figure 4.3. All visualizations of THOR and MCNP meshes are generated using ParaView [32]. A zoomed-in view of the detector, including mesh boundaries, is provided in Figure 4.4. It displays the high degree of mesh refinement required to resolve the interior regions of the detector, where the computed response is most sensitive to the spatial discretization error inherent to S_N calculations.

The THOR mesh for the unshielded case contains 18,278 nodes which bound 105,166




Figure 4.3: Annotated view of the simplified model of the unshielded case as used in THOR. The view is a slice through the middle of both the BeRP ball and detector. Mesh boundaries are hidden.




Figure 4.4: Enhanced view of a cross section of the unshielded detector model with visible mesh boundaries.

tetrahedral elements. The majority of these elements are used to mesh important details in the BeRP ball and the SNAP detector, evidenced by the relatively large cells composing the inactive detector body (dark blue) and air (dark orange, visible at the bottom) in Figure 4.4. The air far from the detector and BeRP ball has the largest elements.

In contrast, the THOR mesh for the shielded case contains 34,280 nodes which bound 202,950 tetrahedral elements. The geometry of the shielded case is shown in Figure 4.5. An enhanced view of the detector with its mesh boundaries visible is given in Figure 4.6. The THOR mesh for the shielded problem has a drastically greater level of mesh refinement in the detector than that for the unshielded problem, and the majority of the new elements are located inside the detector accordingly.



Figure 4.5: The simplified model of the shielded case as used in THOR. The view is a slice through the middle of both the BeRP ball and detector. Mesh boundaries are hidden.




Figure 4.6: Enhanced view of a cross section of the shielded detector model with visible mesh boundaries.

4.2.3 Helium Modeling

We make note here of the response cross section referenced in Eq. 3.14. For the purposes of this numerical validation, the total interaction cross section in the active helium region is taken as the response function's equivalent cross section. He-3, the primary isotope in the active helium region, is a low-scattering material with a dominant (n, p) reaction cross section, which is the reaction type registered by the SNAP detector. The ENDF/B-VIII.0 (n, tot) total cross sections and (n, p) reaction cross sections are plotted in Figures 4.7 and 4.8, respectively [33]. Especially at thermal energies, the (n, p) reaction clearly dominates. Due to their similarity and the expense of computing multigroup (n, p) cross sections compatible with THOR's detector response calculator, the total interaction cross section is therefore instead used for computing the detector response.




Figure 4.7: ENDF/B-VIII.0 total cross section for He-3 (barns) [33].




Figure 4.8: ENDF/B-VIII.0 (n, p) cross section for He-3 (barns) [33].

CHAPTER 5

RESULTS & VALIDATION

5.1 Introduction

The following sections detail the results of the numerical validation exercises for the new implementation of the GCCSM detailed in Section 3.5. For the remainder of Section 5.1, we provide details necessary to accurately compute the detector responses for the problems of interest. The ray effects inherent to problems with the BeRP ball in air using the S_N method are shown in Section 5.2.1. Because the configurations modeled have large spatial domains containing mostly air, the impact of ray effects shown in this section justifies our use of the MCNP-THOR coupling scheme for these problem configurations. For the unshielded detector case, the effects of poor MC- S_N mesh agreement are noted in Section 5.2.2. Then, the computed detector responses for the unshielded case found through coupling when using a better-optimized *FMESH* grid are presented in Section 5.2.3. The same grid is then used to couple MCNP and THOR for the shielded detector case, the results of which are presented in Section 5.3.

5.1.1 Explanation of Important Model & Running Parameters

In an analog MC simulation, a majority of particles emanating from the BeRP ball in directions not pointing to the detector can exit the domain in straight line trajectories, undergoing no collisions, without ever reaching the detector. This can result in poor MC statistics near and within the detector due to a lack of particles reaching this area and scoring tallies that would properly contribute to the computed solution. Results for these problem configurations can therefore suffer from a lack of optimization in MCNP. Accordingly, source directional biasing is implemented as a simple VR technique. Using the *sb* card with a vector favoring the $+y$ direction, simulated particles favor traveling toward the detector from the BeRP ball. This avoids the waste of over half of all simulated particles on account of leakage out of the domain of the problem. Source directional biasing is favored over other VR techniques due to its simplicity of implementation and obvious practical value for this problem.

10^8 particle simulations in MCNP are used for all coupled cases. All results of coupled problems are compared to continuous energy MCNP reference solutions given in [6]. In all executions of MCNP, the particle weight is kept at 1, meaning the S term in Eq. 3.15 must be applied during postprocessing. Its value is given in the next section. All runs utilize an S_2 quadrature order in THOR to emphasize the presence of the ray effects which are addressed by the coupling scheme. For the modeling of anisotropic scattering, limited memory resources necessitate running THOR with a maximum of a P_3 scattering order, which is considered sufficient to demonstrate the impact of anisotropy on solution accuracy for the problems of interest. Justification for excluding higher-order scattering is found in the figures of results presented in Section 5.2.3 on account of the convergence of the P_0 , P_1 , and P_3 solutions in the unshielded case, which is also found in the shielded case.

5.1.2 Postprocessing

The postprocessing terms in Eq. 3.15 must be determined. The fission correction factor F is computed as approximately 3.95 using MCNP, by running the case first with induced fission enabled and then disabled and computing the ratio of the computed responses. The density factor of the active helium $D = 1.0924$ in the unshielded case and $D = 1.0204$ in the shielded case as automatically computed by THOR. The lower density factor in the shielded case is attributable to the more refined meshing of the detector in this case as observed in Figure 4.6.

The source strength S proves the most troublesome, since the BeRP ball was fabricated in October 1980 with some uncertainties in its exact composition even at that time [31]. In the more than four decades since its fabrication, the composition of the BeRP ball has changed due to radioactive decay and fission during experiments, meaning only estimates can be made regarding the exact source strength. For the purposes of this work, the spontaneous fission source rate S is taken as $2.4989 \times 10^5 \frac{n}{s}$ as computed by Herring using data from the Korea Atomic Energy Research Institute in [6]. Using this source strength maintains consistency with previous work and enables comparison of results. However, it should be noted that subsequent estimations of the BeRP ball's spontaneous fission rate have varied as high as $2.90 \times 10^5 \frac{n}{s}$. For this reason and due to aforementioned model simplifications, direct comparison to experimental data is of less concern than comparison to continuous energy MC calculations in this work.

5.2 Unshielded Case

The geometry of the unshielded case has been shown in Figure 4.3. The removable HDPE slab is not present in this case, with air in its place as visible in Figure 4.4. Without the HDPE slab, only neutrons reaching the detector while already having energies in the thermal range are filtered out by the cadmium sleeve.

5.2.1 Uncoupled S_N Solution

For an adequate presentation of ray effects mitigation, we first demonstrate the severity of the ray effects inherent to this problem by running THOR without the MCNP coupling scheme. THOR's regression test suite contains a 32 group, two source-region implementation of the geometry of Figure 4.3 using the spontaneous fission spectrum in the BeRP ball as the sole nonzero source. The source is highly localized in the BeRP ball with little scattering in the air. The configuration is therefore susceptible to severe ray effects.

This case is run with an S_2 quadrature set. Running the uncoupled test problem produces the flux output found in Figure 5.1. In this view, which is a slice down the mid-plane of the problem cutting through both the center of the BeRP ball and of the detector, four rays of non-physically high flux emanate from the BeRP ball, none of them reaching the detector located at the center of the left edge of the plot. In fact, when rotating the domain, it can be found that none of the angles in the three-dimensional quadrature set starting at the BeRP ball, which are visible via the rays, intersect with the detector.

With no rays reaching the detector, the detector response naturally is under-predicted on account of the ray effects. Due to differences in source magnitudes, direct numerical comparisons between this uncoupled problem's solution and those of the validation cases are not fruitful for observing ray effects mitigation. However, this general conclusion shall inform analyses of the efficacy of the coupling scheme in subsequent sections. That is, ray effects in this specific problem geometry with an S_2 quadrature set lead to an underprediction of the detector response. Since the validation cases use this same S_2 quadrature set, they will be subject to an underprediction of the detector response on account of the ray effects. With increasing MCNP collision counts, then, the overall computed detector response should be expected to rise toward an asymptote near the expected value, the point at which additional MCNP collisions do not contribute to a


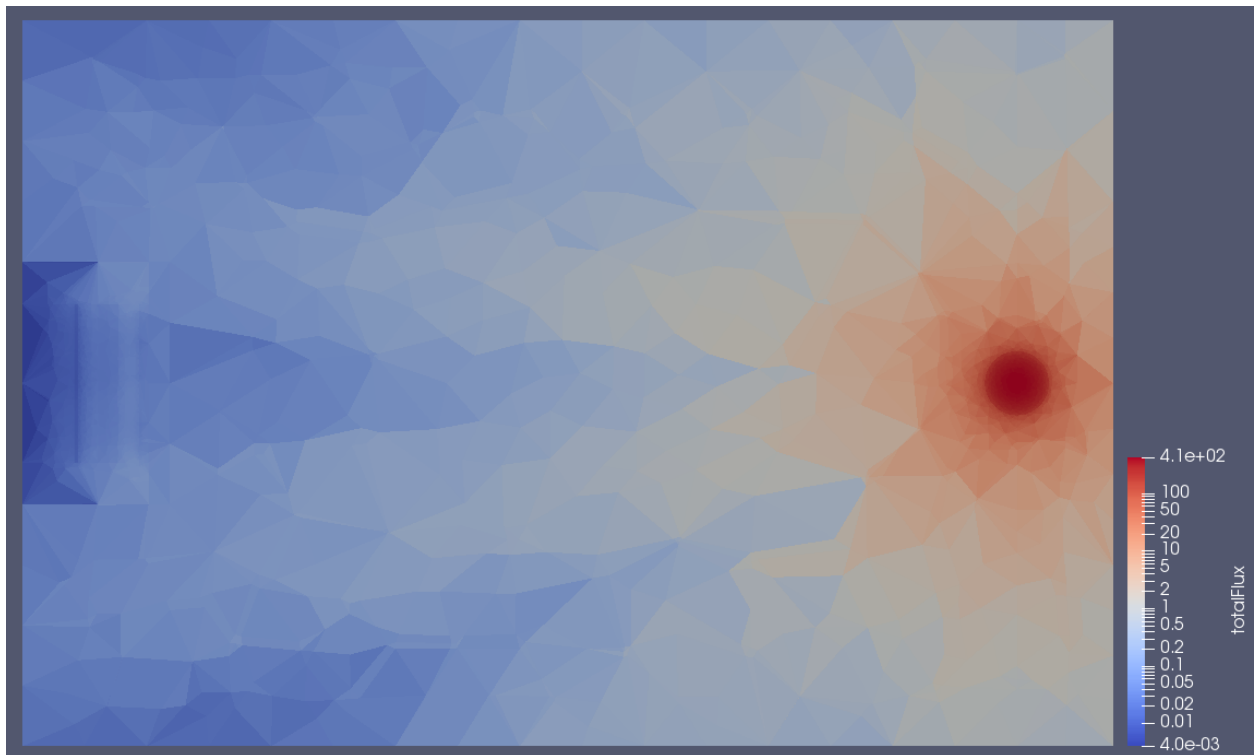
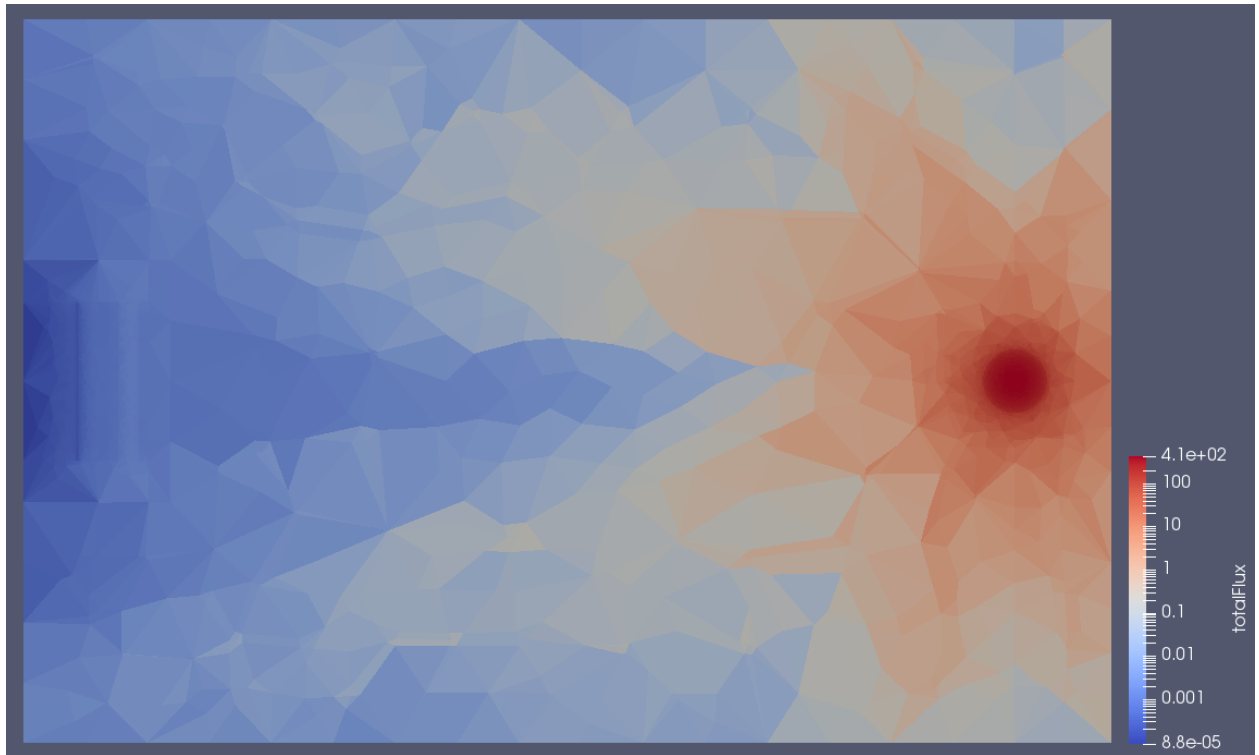


Figure 5.1: Heatmap of the flux output from THOR in the uncoupled test problem using S_2 quadrature, with the plane taken at the center of both the source and detector.

further mitigation of ray effects in a notable fashion.

Furthermore, in Figure 5.2 on the following pages, the results of the same uncoupled test problem using S_4 , S_8 , and S_{16} quadrature sets are shown separately. This source-detector case demonstrates the greatest flaw with attempting to address ray effects by raising the quadrature order of the calculation. Even upon visual inspection, all rays clearly miss the detector even with an S_8 quadrature set. At orders above S_4 , computational efficiency dramatically suffers, with wall clock time more than doubling for a doubling of the quadrature order. Using four processors, for instance, the S_{16} problem takes multiple days to execute.

Figure 5.2: Flux output from THOR in the uncoupled test problem using increasing quadrature orders. The view is down the problem's midplane. The first figure uses an S_4 quadrature set. The second uses S_8 . The third uses S_{16} .






Table 5.1 contains the raw detector responses for the uncoupled solutions, not weighted by the factors from Eq. 3.15. We find that even upon reaching S_{16} , the solution is still not converged with increasing quadrature order, with computational costs growing too great to reasonably find the asymptotic solution. Although less visible, then, the ray effects clearly still dominate the solution. Therefore, raising the quadrature order is not a realistic option when attempting to mitigate ray effects for problems such as this, motivating use of our new coupling scheme.

Table 5.1: Raw detector responses of the uncoupled problem at the various quadrature orders.

Quadrature Rule	S_2	S_4	S_8	S_{16}
Raw Detector Response	$8.5E - 5$	$4.2E - 3$	$1.5E - 1$	$1.3E + 0$

5.2.2 Unoptimized Coupling Case

Model

The first results generated using the new coupling scheme for the GCCSM employ a simplistic, uniform MCNP *FMESH* overlay grid. The grid contains $48 \times 48 \times 48$ cells for a total of 110,592 cells. A heatmap showing the computed total flux (of all collisions and energy groups) in MCNP is found in Figure 5.4. The *FMESH* grid (blue gridlines) and THOR mesh (multicolored lines) are both overlaid on the figure to compare the meshes, while the actual magnitude of the flux is not of particular interest. Note how the uniform Cartesian grid very poorly resolves the BeRP ball and SNAP detector when compared to the THOR mesh. Additionally, the grid excessively resolves the air which has comparably large THOR elements. The air is naturally an area of very little interest when seeking the detector response, making this configuration computationally wasteful.

The simultaneous over-resolution of the air and under-resolution of the detector have compounding effects on the coupling scheme's solution accuracy. The large THOR air elements become susceptible to more dramatic changes in mapped flux, as the centroids of adjacent elements can lie inside *FMESH* cells far from one another. This poor smoothness of the resulting source in the air contributes to a less-accurate collided distributed source. More importantly, large groups of the very fine detector elements are liable to all be assigned the same flux value, which may not be accurate for the region. For instance, there are *FMESH* cells containing cadmium thermal shielding material, HDPE moderator, and active helium at the front of the detector, meaning a variety of competing effects on the flux magnitude are captured in the same *FMESH* cells with a single value of flux for each energy group. Additionally, a mere six cells contain the majority of the active helium in the detector in this view down the problem midplane, even though this is a key area for accurately determining the detector response.

Figure 5.5 displays another heatmap, this of the first collision source generated using




Figure 5.4: Cross-sectional heatmap of MCNP flux using the uniform *FMESH* grid. Both the THOR mesh and *FMESH* grid are overlaid.

uncollided MCNP flux values. It is summed over all 32 energy groups in THOR after applying the coupling scheme. Figure 5.5 displays some of the problematic characteristics predicted for a poorly-optimized *FMESH* grid. The hundreds of THOR elements modeling the BeRP ball are all contained within a few *FMESH* cells, creating visible patches across the surface. A similar phenomenon occurs in the detector, where many fine THOR elements receive identical source values.




Figure 5.5: First collision total source (summed over energy groups) found using the unoptimized grid for coupling to THOR.

Results

Figure 5.6 displays the computed detector responses versus particle collision count in MCNP, ranging from $(n - 1) = 0$ (uncollided) to $(n - 1) = 32$, using this mesh configuration. They are compared to the computed continuous energy MCNP response of 41.9 CPS computed by Herring [6]. P_0 , P_1 , and P_3 scattering cross sections in THOR are applied, with each increasing the degree of anisotropy modeled: P_0 cross sections in THOR indicate fully isotropic scattering in both the n -collided source's generation and the subsequent S_N calculation.

Analysis

We first note that, even for this case with an unoptimized *FMESH* grid, the ray effects inherent to the problem are reduced dramatically when generating THOR sources using




Figure 5.6: Detector response (CPS) computed using P_0 , P_1 , and P_3 scattering cross sections in THOR with the unoptimized *FMESH* grid. MCNP reference solution provided for comparison [6].

neutrons having MCNP collision counts between two and 16. At eight MCNP collisions, the coupling scheme with S_2 quadrature has an error compared to the reference value of -0.64% when using P_0 cross sections and -15.58% when using P_3 . In [6], Herring noted diminishing returns in solution accuracy after about the fifth-collided source, generally terminating analysis at six MC collisions due to convergence of the solution at this point. The results found in this case are therefore consistent with the Hammer-based GCCSM.

Additionally, it is predicted that for this geometric configuration, anisotropic scattering should play a more significant role than in other common applications, such as reactor modeling. This is due to the hard neutron spectrum entering the active helium region in the detector after passing through the thermal cadmium filter. Modeling forward-peaked anisotropic scattering should reduce the count rate when compared to modeling only isotropic scattering due to increased leakage from the detector in the direction normal

to the BeRP ball. Indeed, Figure 5.6 shows a sizeable drop in count rates for lower collision counts when using P_1 and P_3 scattering. Since the P_0 results closely match the reference detector response, significant error is found for the anisotropic cases at low MCNP collision counts. However, the $(n - 1) = 16$ value matches exceptionally well. On account of these results, and due to the atypical moderators chosen, it is predicted that the scattering is mostly modeled isotropically in MCNP, hence the strong agreement between the P_0 coupled result and the pure MCNP result and the convergence of the three scattering orders' solutions.

Despite the apparent successes of the coupling scheme for this configuration, especially when using P_0 cross sections in THOR, it exhibits problematic behavior with very high MCNP collision counts. For instance, when coupled with the 32-collision MCNP neutron flux, an erroneously high detector response of approximately 61–62 counts per second (CPS) is computed. This jump in count rate is predicted to have two primary causes. First, an inadequate resolution of the thermal flux generally is inherent to the MC solution scheme for this problem. Incoming thermal flux is filtered out at the cadmium sleeve in the detector, meaning the majority of counts produced in the detector arise from neutrons which thermalize inside of it. At 32 collisions in MCNP, there should be sufficiently-many thermal neutrons to produce sizeable counts, but finding neutrons having a collision count of exactly 32 inside the detector is a very rare event. Therefore, MC cannot properly resolve the 32-collided thermal flux in order to successfully perform coupling.

This error is compounded by *FMESH* cells which are too large in the detector and overlap multiple regions. The cells containing most of the active helium, for example, also contain sizeable quantities of HDPE, meaning active helium THOR elements are likely assigned an excessive, erroneous thermal flux contributed by the HDPE region at such high collision counts. Since, as seen in Figure 4.7, the response function's equivalent cross section is highest for energies in the thermal range, these two effects which contribute to inaccurate thermal flux values in the detector likely cause the excessively high count rate

in the coupled scheme when using very high MCNP collision counts to generate sources in THOR.

These two erroneous results establish the aforementioned importance of mesh agreement between MCNP and THOR. As will be shown in the following section, the requisite adjustments to the mesh prove simple without detracting from the generality of the coupling scheme, especially since the P_0 results in this case do, in fact, approximate the anticipated detector response with very little error at reasonable MCNP collision counts.

5.2.3 Optimized Coupling Case

Model




Figure 5.7: Optimized *FMESH* grid (white gridlines) overlaid on the geometry of the unshielded case.

Simple changes to the MCNP Cartesian mesh yield results in better agreement with theoretical expectations. The *FMESH* grid is more carefully tuned for this problem, and the new mesh may be seen in Figure 5.7 overlaid on the geometry. Coarse grid points are placed around the BeRP ball and SNAP detector with some air gap, and the degree of refinement is raised within these regions and lowered in the intermediate air. The net effect improves resolution of the points of interest without excessively increasing the cell count and subsequent computational expense.

A view down the midplane of the problem geometry of the first collision source, summed over all energy groups, which is generated from the uncollided MCNP flux and mapped to the THOR mesh, is shown in Figure 5.8. This figure demonstrates an example of information loss during the coupling process. The locations of *FMESH* cells' edges are not actually printed by MCNP. Instead, their center coordinates are printed with only three decimals of precision. When computing the locations of the cell edges, they are found to not necessarily actually be contiguous, especially for *FMESH* grids with very small cells such as in this case. THOR elements' centroids can therefore happen to be found to fall within the spaces between the bounds computed for the *FMESH* cells, meaning no *FMESH* cell effectively contains the THOR element's centroid. This leads to an inability to directly assign a source value. This phenomenon is most visible in the dark air cells at the top of the figure.

In such cases, the search is performed again by rounding the coordinate values of the THOR element's centroid to two decimal places. If a match is not found even then, the source value is simply assigned 0 rather than interpolated. The impact on the accuracy of the solution is minimal as discussed in the analysis for this case, which in fact demonstrates the robustness of the GCCSM. We note that the majority of elements which experience this phenomenon are in fact found on the central plane shown, and compared to the total number of elements, very few (less than 0.5%) are unable to be assigned a source value. Even so, this erroneous phenomenon partially motivates the

proposal in Section 6.2.2 for developing additional mapping techniques which could allow for weighted averages/interpolations of data in nearby *FMESH* cells to make up for any failures to match *FMESH* cells to THOR elements.




Figure 5.8: First collision source found using the better-optimized grid for coupling to THOR.

Results

Figure 5.9 displays the computed detector response using P_0 , P_1 , and P_3 scattering in THOR with the *FMESH* grid's optimizations in place. Results of cases with sources generated using neutrons having MCNP collision counts from 0–8, 16, and 32 are included as before. Herring's computed continuous energy MCNP response of 41.9 CPS [6] is shown for comparison. This configuration allows for the analysis of the point of diminishing returns for the method and of the convergence of the different scattering orders' solutions when considering sources generated from neutrons having very high MCNP collision

counts.

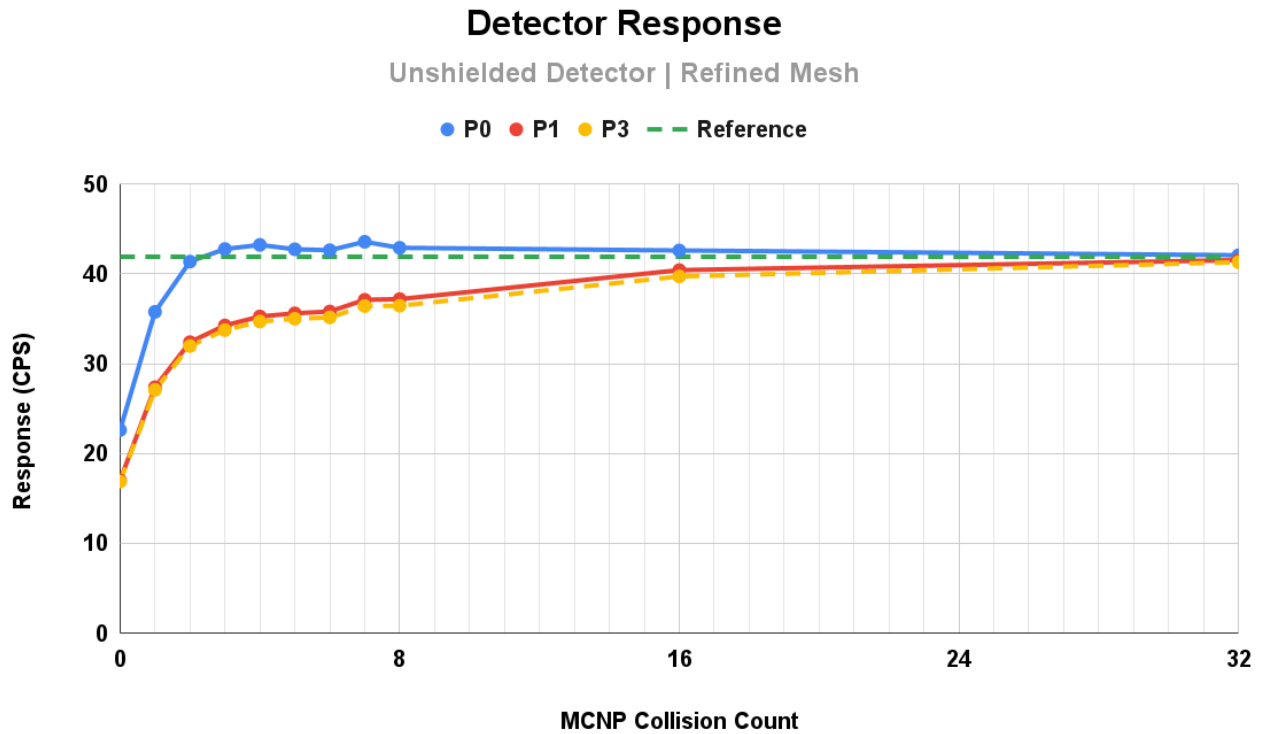


Figure 5.9: Detector response (CPS) computed using P_0 , P_1 , and P_3 scattering cross sections in THOR with the optimized *FMESH* grid. The continuous energy MCNP reference solution is provided for comparison [6].


Analysis

Several promising trends reveal themselves in Figure 5.9. First, the P_0 solution slightly over-predicts the detector response for cases utilizing MCNP collision counts greater than two, but the error when compared to reference remains low. For detector responses computed with third-collided and higher sources, a maximum error of 3.16% is found for the case using a fifth-collided source (generated from fourth-collided MCNP flux). The losses in accuracy from introducing anisotropic scattering in THOR do not introduce as significant an error as in the case with the unoptimized *FMESH* grid (see Figure 5.6). Additionally, most of the loss introduced by anisotropic scattering is incurred during the


application of the P_1 cross sections, meaning increasing the scattering order any higher, resulting in computationally intensive calculations, is not predicted to continue to notably change the computed detector response. So, the trends found in Figure 5.9 point to the quality of the coupling scheme independent of the THOR model.

In cases using very high MCNP collision counts (16–32), the P_0 results fall toward the P_1 and P_3 results, which monotonically approach the reference solution. When using sources generated from MCNP fluxes having 32 collisions, the case with the highest error, the P_3 case, has an error compared to reference of only -1.43% . This behavior is expected and should be contrasted to the behavior found in the case when using the unoptimized *FMESH* grid. At very high collision counts, the thermal flux computed by the continuous energy MCNP run is concentrated in the detector, and it contributes the majority of the response, not the ϕ_{n+} computed by THOR. This is evidenced in Figure 5.10, comparing the fast/epithermal (Figure 5.10a) and thermal (Figure 5.10b) flux magnitudes in the *FMESH* cells along the centerline of the problem. The neutrons considered have any number of collisions ($0-\infty$). The cells containing the outer radius of the cadmium sleeve and all HDPE and helium within it are found between the second set of blue lines on the right of the figures. Therefore, those lines effectively surround the active section of the detector. The magnitude of thermal flux in the detector is only approximately one order of magnitude or less below that of the fast flux, while the thermal cross sections are five or more orders of magnitude above the fast cross sections as shown in Figure 4.7. Therefore, the thermal flux computed by MCNP at very high collision counts dominates the computed response, so THOR sources generated from these fluxes are indeed expected to produce results very close to the continuous energy MCNP reference solution.

Introducing the different order scattering cross sections after the coupling step to primarily thermal flux should incur less error than when coupling with low collision count neutrons, since thermal flux tends to scatter with a more isotropic distribution



(a) Fast & Epithermal Flux



(b) Thermal Flux

Figure 5.10: Line plots of the flux of neutrons having any number of collisions down the centerline of the problem. The left set of blue lines encompasses the BeRP ball, while the right set contains the cadmium-to-cadmium width of the detector.

than fast flux, rendering all solutions similar. Finally, it has already been observed that introducing the higher order cross sections drops the computed response, which is predicted to be attributable to higher leakage from the detector. Therefore, the marginal difference in results found between the three scattering orders indicates the elimination of ray effects, since the error from the different cross sections has been diminished at such high MCNP collision counts. The remaining difference is attributable to the few particles in the THOR run exiting the BeRP ball and reaching the detector, which tend to contribute less to the computed response when anisotropic scattering is applied due to forward peaking of fast neutrons undergoing scattering interactions.

An additional measure indicating the gradual elimination of ray effects by considering neutrons of higher MCNP collision counts comes from the magnitude of the flux inside and immediately surrounding the BeRP ball. Since induced fission is disabled in MCNP and is instead corrected for in postprocessing, additional fission neutrons are not generated in the BeRP ball as the MCNP simulation progresses. Therefore, by summing the flux of a given collision count inside the ball, the ratio of this quantity to the total magnitude of uncollided flux is indicative of the fraction of flux exiting the BeRP ball as collisions progress. Since there is very little scattering in the air, as soon as neutrons exit the BeRP ball, they will imminently either leak from the system or reach the detector. In both cases, they will not contribute to ray effects. Therefore, the ratio of n -collision to uncollided flux in and around the BeRP ball (summed from the *FMESH* cells) is expected to correlate with the impact of ray effects on the problem. Indeed, as found in Figure 5.11, the magnitude of flux in the BeRP ball is slightly more than halved with each additional collision in MCNP. By three collisions, the anticipated error from ray effects therefore falls to less than 10%, indicating that ray effects are mostly eliminated within 3–4 MCNP collisions. The detector response results corroborate this finding, particularly for the P_0 THOR cross sections, with very little change in detector count between 3–8 collisions.




Figure 5.11: Comparison of the n -collided to uncollided MCNP flux remaining in the BeRP ball, indicating a mitigation of ray effects with increasing collision counts.

Finally, an analysis of the error of the MCNP flux at various collision counts is warranted, as all data presented so far indicates that raising the MCNP collision count consistently improves the results toward the continuous energy reference solution as seen in Figure 5.9, albeit perhaps with diminishing returns. The average relative error of the thermal flux (having energy less than 0.625eV), inside the detector, in a two-group run of the simulation, is given in Figure 5.12. As should be expected, the average error generally rises with a higher collision count due to the lower likelihood of encountering exactly n -collided neutrons with increasing n . This indicates that, when performing coupling with neutrons having high collision counts — greater than about six in this case — the ϕ_{n+} computed by THOR uses a very unreliable fixed source with average relative errors in the detector above 10%. This finding further justifies using lower collision count MCNP neutrons for coupling. Which MCNP collision count actually proves optimal remains a topic for study and is discussed in Section 6.2.3 as one proposal for future work.

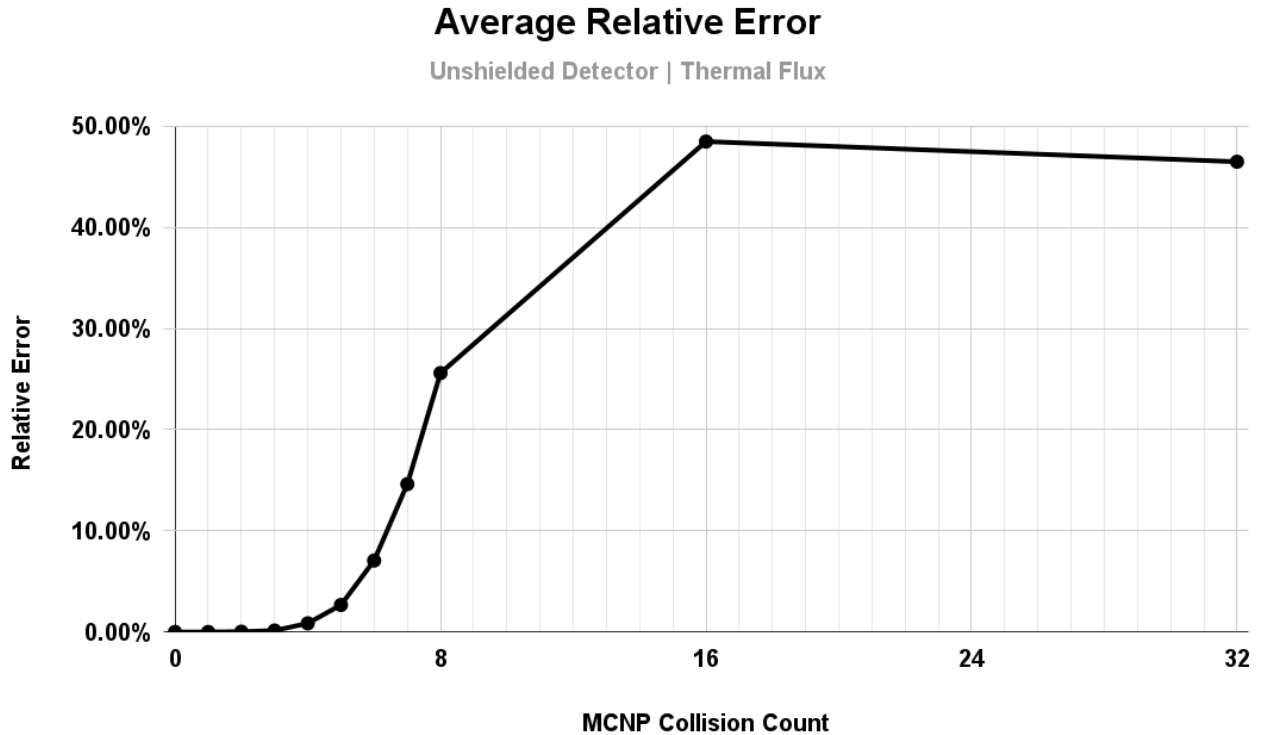


Figure 5.12: Average relative error of thermal flux in the *FMESH* cells making up the detector.

5.3 Shielded Case

Model

The geometry of the shielded case is found in Figure 4.5. The one-inch HDPE slab ("shield") can be seen as an extension of the detector body with an air gap between the body and cadmium. The shield filters out slower neutrons which are thermalized within it and then absorbed in the cadmium sleeve. Therefore, only the fastest neutrons coming from the BeRP ball are counted after additional thermalization in the inner HDPE layer. For that reason, the shielded case should have generally lower computed detector responses than the unshielded case.

The same *FMESH* grid utilized for the optimized unshielded case can be applied to the shielded case. It is shown in Figure 5.13 overlaid on the new geometry. The even

greater refinement of the SNAP detector in the shielded case once again demonstrates the importance of at least a sensible agreement in mesh refinement level between MCNP and THOR.




Figure 5.13: *FMESH* grid overlaid on the geometry of the shielded problem.

A view of the first collision source generated from the uncollided MCNP flux in the midplane of the geometric configuration is shown in Figure 5.14. The figure once again contains tetrahedra without assigned source values. They comprise roughly 0.1% of all elements in the problem and are primarily found on the plane shown. The shield becomes a significant first collision source while reducing the magnitude of the source in the inner HDPE layer when compared to the unshielded problem in Figure 5.8.




Figure 5.14: First collision total source found for the shielded case.

Results

Figure 5.15 depicts the computed detector responses for the shielded case. P_0 , P_1 , and P_3 THOR cross sections are applied. The continuous energy MCNP reference result computed by Herring [6] is 33.5 CPS and is included for comparison.

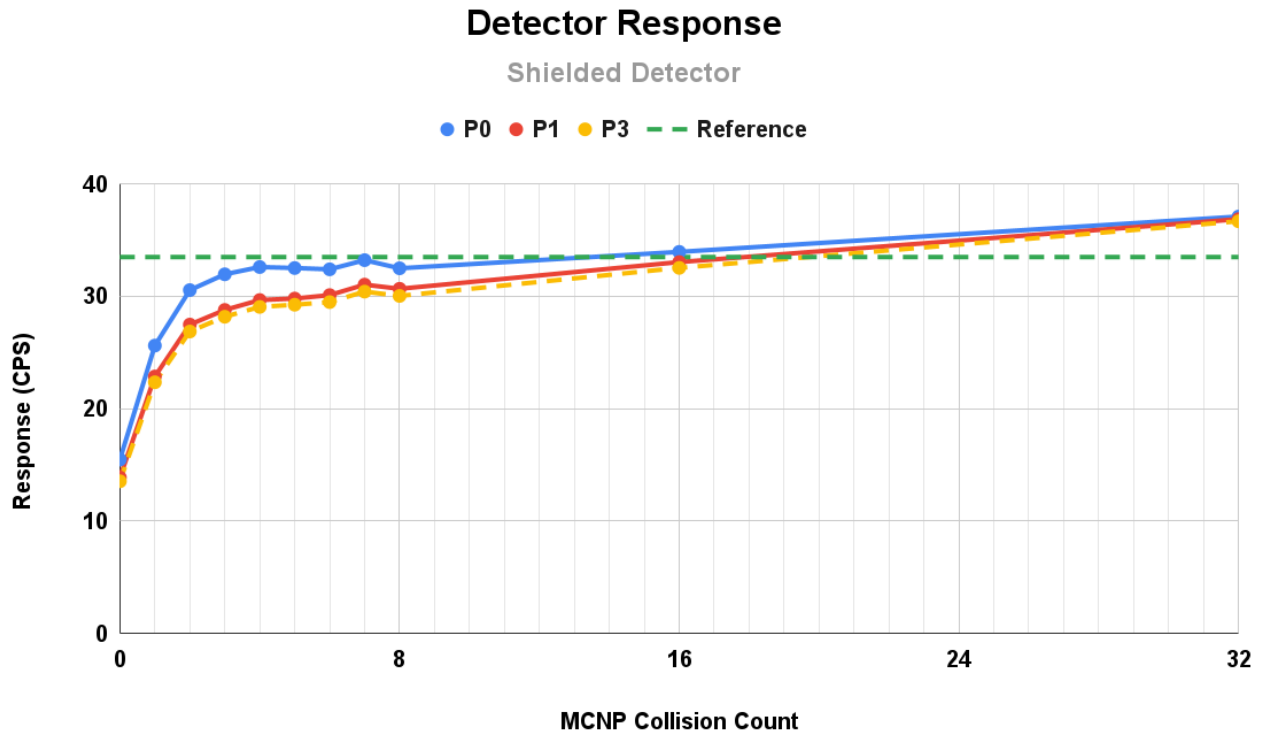


Figure 5.15: Detector response (CPS) computed using P_0 , P_1 , and P_3 scattering cross sections in THOR for the shielded case. MCNP reference solution provided for comparison [6].

Analysis

As in the unshielded case, the response stabilizes when using neutrons having between 3–8 MCNP collisions to generate the THOR sources. When using sources generated from neutrons having higher MCNP collision counts, the three solutions converge as in the unshielded case. The P_0 result using sources generated from neutrons having 32 MCNP collisions has only a 5.4% error compared to reference. However, with greater error than in the cases with lower MCNP collision counts, it is clear that ray effects mitigation is sufficient without taking the MCNP run to very high collision counts.

Comparison of the results of the shielded case to the unshielded case provides insights to the accuracy of the method as a whole. For both problems, using neutrons from MCNP having collision counts of 0–2 to generate THOR sources, the computed detector response

rapidly rises for all THOR scattering orders. Between 3–8 MCNP collisions, the response stabilizes, with the P_0 THOR cross sections typically yielding a solution very close to the MCNP reference solution. Higher scattering order solutions tend to slowly continue approaching the reference solution with higher MCNP collision counts, but they do not match P_0 results until unreasonably high collision counts. This is attributable to the fact that, at such high MCNP collision counts, the vast majority of the problem has been simulated in MC, with very little contribution from THOR. In neither problem does taking the MCNP run to more than four particle collisions improve the response result for cases using P_0 scattering cross sections in THOR.

Therefore, we may conclude that, despite containing nearly twice as many THOR elements while using the same *FMESH* grid, ray effects are successfully mitigated in the shielded case, and additional refinement in the *FMESH* grid is not found to be necessary despite once again observing the rise in count rates when using 32-collision MCNP neutrons. With evidence from both the tested problems and the results found by Herring [6], we conclude that MC particles may be terminated after 3–4 collisions, with using higher collision counts to generate sources in THOR not substantially improving the computed detector response values.

Minimizing the number of collisions performed in the MC simulation can save wall-clock time when compared to completing a full simulation, as particle histories can terminate soon after reaching high-scattering media, such as the HDPE in the detector, in which the particles will interact frequently without traveling far. Killing particles at this point prevents the MC code from unnecessarily using computational resources to simulate particles which fail to undergo absorption or leakage in a reasonable number of collisions in such media. The numerous collisions experienced by these neutrons do not contribute much to the accuracy of the solution, as evidenced by the increasing relative error in MCNP with particle collision count observed in Figure 5.12 and by the sufficiently accurate results found for low MCNP collision counts (3–8) in the coupled scheme. By

saving on the cost of simulating numerous particle collisions per history, more total histories can instead be run in order to improve the accuracy of the flux used to generate the source for the S_N portion of the calculation, or total computational time may instead be reduced if the error is already within acceptable limits.

The ability to cut off expensive MC particle simulations at very low collision counts while obtaining accurate results as shown in this section therefore demonstrates the value of the MC- S_N coupling scheme. Where finding an accurate MC solution for many problems requires numerous particle history simulations and use of VR techniques, coupling the results of even a relatively cheap MCNP run with THOR can yield computed responses in excellent agreement with an expensive, full MCNP calculation. Additionally, the coupling scheme is far more broadly applicable and requires less tuning by the user than VR techniques in MC, which are either configuration-specific (as for, e.g., the *sb* card) or already require an additional calculation (as for, e.g., CADIS [24]).

CHAPTER 6

CONCLUSIONS

6.1 General Remarks

The ray effects inherent to the S_N method hamper its ability to provide accurate solutions to the neutron transport equation. Past attempts at coupling MC and S_N codes have found the approach capable of effectively reducing the impact of ray effects in problems prone to them [4][6]. This thesis has focused on a new implementation of the GCCSM, which takes MC-computed fluxes of neutrons having a given collision count of $(n - 1)$ to generate sources of neutrons having n collisions in the S_N code. Using MCNP as the MC code and THOR as the S_N code, the method has proven to mitigate ray effects and to bring the coupled solution to within a reasonable margin of error of the fully-detailed, continuous energy MCNP solution.

The coupling scheme presented herein has several notable advantages to those that have preceded it. Most importantly, the method does not rely on tallying MC fluxes on identical grids to those used by the S_N code, as was the case in the Hammer-THOR implementation [6]. Furthermore, the ability to use continuous energy MCNP calculations

before binning into a desired energy group structure allows the final, coupled solution to closely match continuous energy MC results. In contrast, Hammer employed the same multigroup cross sections as THOR, limiting accuracy of the MC calculation. Therefore, the goal of merging the benefits of MC and S_N — namely, accuracy and computational efficiency, respectively — is successfully achieved with this implementation of the GCCSM for ray effects mitigation.

Two problems modeling experiments of a localized plutonium source with a one-meter air gap to a neutron detector have been analyzed. Subject to severe ray effects, these configurations prove wholly unsuited for standard S_N calculations without a special quadrature set or a method of ray effects mitigation. Using the GCCSM, neutron sources generated from neutrons having as few as two to three collisions in MCNP proves sufficient to reduce ray effects to acceptable levels and allow for accurate S_N calculations. Despite a significant change in the level of refinement in the detector model between the two problems, the same Cartesian mesh tally generated by MCNP may be used for both, pointing to a robust method which demands minimal problem-dependent tuning.

Therefore, the GCCSM presented in [6] has been improved upon using a novel coupling scheme between MCNP and THOR. The method produces results in better agreement with computationally expensive, accurate, continuous energy MC calculations. Due to the accessibility of MCNP and other MC codes, the method can be implemented easily when seeking to mitigate ray effects in other S_N codes having alternative spatial discretization schemes.

6.2 Proposals for Future Work

6.2.1 Parallelization & Runtime Improvements

This novel implementation of the GCCSM currently does not employ methods of parallelization, and the bounds of every cell in the entire *FMESH* grid must be searched

over for each THOR element until the cell containing the element’s centroid is found. Once a map — a list of which *FMESH* cell contains the centroid of each THOR element — has been generated, it may be read in for future runs, and the coupling scheme then takes less than four minutes to generate a THOR source from an MCNP output file in all cases considered.

However, both the initial mapping process and the generation of the THOR sources through application of the multigroup cross sections could be dramatically accelerated. Both steps are embarrassingly parallel in theory. Splitting the THOR mesh into groups of tetrahedral elements for multiple processors, the wall-clock time of the entire process could roughly be divided by the number of participating processors, not accounting for communication overhead. Furthermore, in the mapping process, more robust search algorithms could rapidly parse through the list of *FMESH* cells without checking the bounds of each cell until a match is discovered, which is the current, naïve searching approach. Finally, implementing the method in Fortran [34] or Julia [35] could improve speed when compared to the current Python [36] implementation.

6.2.2 Alternative Mapping Techniques

As implemented, the GCCSM has been found to successfully mitigate ray effects. However, the method of mapping between MCNP’s Cartesian grid and THOR’s tetrahedral mesh, described in Section 3.5.1, is simplistic. Noting that THOR elements may overlap multiple *FMESH* cells and vice-versa, some information is inherently lost during the mapping procedure. One proposal to help reduce the possible impact of mesh differences would involve finding not only the *FMESH* cell containing the centroid of each THOR element, but also finding the *FMESH* cells containing the vertices of each THOR element. Numerical experiments could be performed to determine the optimal weighting factors to find weighted averages of the fluxes found in the five (not necessarily unique) *FMESH* cells containing the THOR element’s centroid and vertices. It is predicted that such a

mapping scheme would smooth differences in neutron flux found at material interfaces, potentially improving the accuracy of the subsequent S_N simulation.

6.2.3 Optimization of MC Collision Count

The results found in this work corroborate those found by Herring [6] indicating that ray effects are mostly eliminated in S_N when using fixed sources generated from neutrons having three or more MC collisions. However, it has been noted in cases using sources generated from very high-collision count neutrons that increasing the MC collision count does not necessarily continuously improve the accuracy of the S_N solution. This is due to the significant statistical error incurred when attempting to tally neutrons of a given, very high collision count, which are rare when compared to neutrons having undergone only a few collisions. A study of the optimal MC collision count, thoroughly analyzing statistical error in MC and marginal error with increasing MC collision count in the final S_N solution, would optimize usage of the coupled method and minimize the expensive tallying of neutrons having various collision counts.

Additionally, such a study would enable the user to terminate neutron histories at the optimal collision count in MC, allowing for more of the problem to be modeled with the computationally efficient S_N method. Due to this work's focusing on broadly demonstrating mitigation of ray effects, MCNP simulations were executed to completion. This enabled the generation of tallies of neutron fluxes having arbitrary collision counts for study. However, in practical implementations, to save on computational costs, the MC run should actually terminate particles after they reach their optimal number of collisions.

6.2.4 Full Validation

The two problems considered in this work are the result of numerous simplifications to models of the experiment discussed in Section 4.1. While fully-detailed MCNP models

of the experiment exist and were used to generate these simplified configurations, no such mesh has been generated for use by THOR on account of the intensive mesh refinement necessary to resolve the finer details. Using a fully-detailed mesh would require extensive computational resources due to the number of spatially-discretized THOR elements. However, it would enable a full validation of the MCNP-THOR coupling scheme, improving upon the purely numerical validation against the MCNP reference solutions for the simplified cases as presented in this work.

6.2.5 Investigation of Higher Quadrature Orders

In order to demonstrate a mitigation of ray effects entirely attributable to the GCCSM, only S_2 quadrature sets have been employed in THOR for the analysis of the problems considered in this thesis. For the uncoupled problem, it was shown that raising the quadrature order is of minimal benefit to physical accuracy. Whether, however, benefits might be gained by a minor refinement of the angular variable — e.g. using an S_4 quadrature set — in combination with this coupling scheme is currently unknown. A cost-benefit analysis for raising the quadrature order while also employing the GCCSM is therefore warranted.

6.2.6 Analysis of Mesh Similarity Requirements

When finding detector responses for the two problems of interest, a degree of mesh refinement suited for the problem geometry helped to achieve reasonable results. However, when almost doubling the THOR element count, further refinement did not prove necessary. A sensitivity study would help influence cost-benefit decision making between the added computational cost incurred and the accuracy gained with increased *FMESH* grid refinement.

6.2.7 Implementation Using OpenMC

A primary goal for the novel implementation of the GCCSM is to demonstrate its broad applicability and ease of implementation for various MC- S_N code pairs. Although MCNP is widely available, its distribution is controlled by RSICC [7], and not all developers may have ready access to MCNP.

OpenMC is a collaborative, open source, object-oriented Monte Carlo code project that is available for any researcher. It has both Python and C/C++ APIs available [37]. An implementation of the GCCSM in OpenMC-THOR, applying the insights gained from this work, could be made using the *openmc.CollisionFilter* tally filter, meaning that no modification to the underlying code would be necessary, even though its open source nature would make such modification possible for an advanced user. Such a coupling would provide a fully open-source coupled code platform employing the benefits of MC in the mitigation of ray effects as documented in this work.

REFERENCES

- [1] J. Kleijnen, A. Ridder, and R. Rubinstein, "Variance Reduction Techniques in Monte Carlo Methods," *CentER Discussion Paper Series No. 2010-117*, Nov. 2010.
- [2] K. D. Lathrop and B. G. Carlson, "Transport Theory: The Method of Discrete Ordinates," Los Alamos Scientific Laboratory of the University of California, Los Alamos, New Mexico, Tech. Rep., Apr. 1965.
- [3] K. D. Lathrop, "Ray Effects in Discrete Ordinates Equations," *Nuclear Science and Engineering*, vol. 32, no. 3, pp. 357–369, 1968. doi: 10.13182/NSE68-4.
- [4] R. E. Alcouffe, "A First Collision Source Method for Coupling Monte Carlo and Discrete Ordinates for Localized Source Problems," in *Monte-Carlo Methods and Applications in Neutronics, Photonics and Statistical Physics*, Provence, France, Apr. 1985, pp. 352–366.
- [5] N. F. Herring, R. A. Yessayan, K. A. Beyer, *et al.*, "Ray Effects Mitigation Through Monte Carlo Coupling for Detector Problems," in *Proceedings of the International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering M&C 2019*, Portland, OR, Aug. 2019, pp. 2228–2237.
- [6] N. F. Herring, "Ray Effects Mitigation in S_N Problems through General Collision Monte Carlo Coupling and Numerical Validation of the THOR S_N Code for Nuclear Nonproliferation Applications," M.S. thesis, North Carolina State University, 2018.
- [7] C. J. Werner *et al.*, "MCNP6.2 Release Notes," Los Alamos National Laboratory, Tech. Rep. LA-UR-18-20808, 2018.
- [8] S. Schunert, R. Ferrer, and Y. Azmy, "Verification of the Three-Dimensional Tetrahedral Grid S_N Code THOR," in *Proceedings of M&C 2013*, Sun Valley, ID, May 2013, pp. 1135–1150.
- [9] N. Herring, R. Yessayan, S. Schunert, R. Ferrer, and Y. Azmy, *THOR User's Manual*, <https://github.com/NCSU-NCSG/THOR/raw/v1.0.2/docs/usermanual/CurrentVersion/usermanual.pdf>, Version 1.0.2, Sep. 2022.
- [10] R. C. Runkle, A. Bernstein, and P. E. Vanier, "Securing Special Nuclear Material: Recent Advances in Neutron Detection and Their Role in Nonproliferation," *Journal of Applied Physics*, vol. 108, 2010. doi: 10.1063/1.3503495.
- [11] A. Haghghat, *Monte Carlo Methods for Particle Transport*. Boca Raton, Florida: CRC Press, 2021.
- [12] J. J. Duderstadt and L. J. Hamilton, *Nuclear Reactor Analysis*. New York, NY: John Wiley & Sons, Inc., 1976.
- [13] E. E. Lewis and W. F. Miller, Jr., *Computational Methods of Neutron Transport*. New York, New York: John Wiley & Sons, Inc., 1984.
- [14] K. Tamura and Y. Shikano, "Quantum Random Numbers Generated by a Cloud Superconducting Quantum Computer," in *Proceedings of MQC 2019*, ser. Mathematics for Industry, vol. 33, 2019, pp. 17–37. doi: 10.1007/978-981-15-5191-8_6.

- [15] F. James, "A Review of Pseudorandom Number Generators," *Computer Physics Communications*, vol. 60, pp. 329–344, 1990. DOI: 10.1016/0010-4655(90)90032-V.
- [16] J. C. Wagner and A. Haghghat, "Automated Variance Reduction of Monte Carlo Shielding Calculations Using the Discrete Ordinates Adjoint Function," *Nuclear Science and Engineering*, vol. 128, no. 2, pp. 186–208, 1998. DOI: 10.13182/NSE98-2.
- [17] E. W. Larsen, "Diffusion-Synthetic Acceleration Methods for Discrete-Ordinates Problems," *Transport Theory and Statistical Physics*, vol. 13, no. 1-2, pp. 107–126, 1984. DOI: 10.1080/00411458408211656.
- [18] K. D. Lathrop and B. G. Carlson, "Discrete Ordinates Angular Quadrature of the Neutron Transport Equation," Los Alamos Scientific Laboratory of the University of California, Los Alamos, New Mexico, Tech. Rep., Feb. 1965.
- [19] R. E. Peterson and G. A. Newby, "Lady Godiva: An Unreflected Uranium-235 Critical Assembly," Los Alamos Scientific Laboratory of the University of California, Tech. Rep., Sep. 1953.
- [20] J. E. Morel, T. A. Wareing, R. B. Lowrie, and D. K. Parsons, "Analysis of Ray-Effect Mitigation Techniques," *Nuclear Science and Engineering*, vol. 144, no. 1, pp. 1–22, 2003. DOI: 10.13182/NSE01-48.
- [21] M. Frank, J. Kusch, T. Camminady, and C. D. Hauck, "Ray Effect Mitigation for the Discrete Ordinates Method Using Artificial Scattering," *Nuclear Science and Engineering*, vol. 194, no. 11, pp. 971–988, 2020. DOI: 10.1080/00295639.2020.1730665.
- [22] T. Camminady, M. Frank, K. Küpper, and J. Kusch, "Ray Effect Mitigation for the Discrete Ordinates Method Through Quadrature Rotation," *Journal of Computational Physics*, vol. 392, pp. 105–123, 2019. DOI: 10.1016/j.jcp.2019.01.016.
- [23] M. B. Emmett, C. E. Burgart, and T. J. Hoffman, "DOMINO, a General Purpose Code for Coupling Discrete Ordinates and Monte Carlo Radiation Transport Calculations," Oak Ridge National Laboratory, Tech. Rep., Jul. 1973.
- [24] 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," *Progress in Nuclear Science and Technology*, vol. 2, pp. 808–814, 2011.
- [25] E. S. Gonzalez, K. A. Beyer, R. J. Fonti, *et al.*, "Hammer: A Monte Carlo Particle Transport Solver to Support Nonproliferation Applications of the THOR Deterministic SN Code," in *Proceedings of ANTPC 2018*, Orlando, FL, Nov. 2018, pp. 110–112.
- [26] Y. Y. Azmy, "Arbitrarily High Order Characteristic Methods for Solving the Neutron Transport Equation," *Annals of Nuclear Energy*, vol. 19, no. 10, pp. 593–606, 1992. DOI: 10.1016/0306-4549(92)90004-U.
- [27] R. M. Ferrer and Y. Y. Azmy, "A Robust Arbitrarily High-Order Transport Method of the Characteristic Type for Unstructured Grids," *Nuclear Science and Engineering*, vol. 172, pp. 33–51, 2012. DOI: 10.13182/NSE10-106.

- [28] R. A. Yessayan, "Improvements to the THOR Neutral Particle Transport Code on High-Performance Computing Systems via Acceleration, Parallelization, and Performance Analysis," M.S. thesis, North Carolina State University, 2018.
- [29] NEA, *International Evaluation Co-operation Volume 42*, Online, 2020.
- [30] J. Hutchinson, J. Bounds, T. Cutler, *et al.*, "A New Era of Nuclear Criticality Experiments: The First 10 Years of Radiation Test Object Operations at NCERC," *Nuclear Science and Engineering*, vol. 195:sup1, S80–S98, 2021. doi: 10.1080/00295639.2021.1918938.
- [31] J. Mattingly, "Polyethylene-Reflected Plutonium Metal Sphere: Subcritical Neutron and Gamma Measurements," Sandia National Laboratories, Tech. Rep., Sep. 2009.
- [32] ParaView Developers, *ParaView Documentation*, Online, <https://docs.paraview.org/en/latest/index.html>, 2020.
- [33] D. Brown, M. Chadwick, R. Capote, *et al.*, "ENDF/B-VIII.0: The 8th Major Release of the Nuclear Reaction Data Library with CIELO-project Cross Sections, New Standards and Thermal Scattering Data," *Nuclear Data Sheets*, vol. 148, pp. 1–142, 2018, Special Issue on Nuclear Reaction Data. Data retrieved from the ENDF database <https://www.nndc.bnl.gov/endl/>. DOI: <https://doi.org/10.1016/j.nds.2018.02.001>.
- [34] Free Software Foundation, Inc., *The GNU Fortran Compiler*, Online, <https://gcc.gnu.org/onlinedocs/gfortran/>.
- [35] The Julia Project, *Julia 1.8 Documentation*, Online, <https://docs.julialang.org/en/v1/>, 2023.
- [36] Python Software Foundation, *Python 3.11.2 Documentation*, Online, <https://docs.python.org/3/>, 2023.
- [37] P. K. Romano, N. E. Horelik, B. R. Herman, A. G. Nelson, B. Forget, and K. Smith, "OpenMC: A State-of-the-art Monte Carlo Code for Research and Development," *Annals of Nuclear Energy*, vol. 82, pp. 90–97, 2015, Joint International Conference on Supercomputing in Nuclear Applications and Monte Carlo 2013, SNA + MC 2013. Pluri- and Trans-disciplinarity, Towards New Modeling and Numerical Simulation Paradigms. DOI: 10.1016/j.anucene.2014.07.048.