

**A COARSE MESH RADIATION TRANSPORT METHOD FOR
REACTOR ANALYSIS IN THREE DIMENSIONAL HEXAGONAL
GEOMETRY**

A Dissertation
Presented to
The Academic Faculty

by

Kevin John Connolly

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

Nuclear and Radiological Engineering and Medical Physics Programs
George W. Woodruff School of Mechanical Engineering
Georgia Institute of Technology
December 2012

**A COARSE MESH RADIATION TRANSPORT METHOD FOR
REACTOR ANALYSIS IN THREE DIMENSIONAL HEXAGONAL
GEOMETRY**

Approved by:

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

Dr. Bojan Petrovic
George W. Woodruff School
Georgia Institute of Technology

Dr. Alireza Haghghat
Department of Mechanical Engineering
*Virginia Polytechnic Institute and State
University*

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

Dr. Tom Morley
School of Mathematics
Georgia Institute of Technology

Date Approved: October 9, 2012

ACKNOWLEDGEMENTS

This work was supported under a Nuclear Regulatory Commission Fellowship, the Georgia Tech President's Fellowship, and a Department of Energy Nuclear Energy University Program Graduate Fellowship. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the author and do not necessarily reflect the views of the Department of Energy Office of Nuclear Energy.

I would first like to thank my advisor, Dr. Farzad Rahnema, for his guidance, patience, and high standards. I have learned volumes from our conversations and sometimes seemingly endless revisions of papers. My pursuit of this degree is, in large part, a result of his encouragement even before my beginning undergraduate course work.

My appreciation extends to the rest of my committee. My thanks to Dr. Bojan Petrovic and Dr. Tom Morley, for their willingness to read and critique my work, Dr. Dingkang Zhang, who walked me through a large amount of computational neutron transport, and Dr. Ali Haghghat, for whose willingness to travel from Virginia Tech to serve on my committee I am especially grateful.

While a student at Tech, I have had the opportunity to listen to lectures, work on assignments and projects, and attend conferences with many fellow Yellow Jackets who have become friends. I would especially like to recognize Geoff, Jason, and Julia, whose friendship lightened the burden of studying undergraduate nuclear engineering. The

camaraderie within the Computational Reactor/Medical Physics lab has made my tenure as a graduate student an enjoyable one. I have had the good fortune to have worked in the same office with Rob for four years, and traversed some of the Southeast with Justin, Steven, Jack, Alex, Saam, Ryan, Andrew, and Daniel; their enthusiasm for engineering research and conference-related road trips have enhanced my desire to work diligently and celebrate the successful publication of results accordingly.

I would like to extend my most sincere thanks to my family and friends who have supported me throughout my time at Georgia Tech, including my friends Joel, John, and Greg, and especially my parents, Kim and John Connolly, my sister Joanne, my grandparents, and my wife Denise. Much of their encouragement and approval was sent over long distances via such modern conveniences as telephones and electronic mail.

With that in mind, I am utterly unable to adequately express my gratitude to those members of my family who, many decades ago, left the familiarity of their homes to brave an Atlantic passage in search of what must have only been, to them, a fabled promised land. Could they have dreamt, while leaving Ireland or Poland, that their journey would eventually enable a great-grandson to pursue a doctorate in a realm of science that was then unknown to humanity? May my quest for a brighter future honor theirs.

TABLE OF CONTENTS

	Page
ACKNOWLEDGEMENTS	iii
LIST OF TABLES	vii
LIST OF FIGURES	viii
SUMMARY	x
<u>CHAPTER</u>	
1 Introduction	1
2 Background	2
Reactor Analysis in Hexagonal Cores	2
The Hybrid Coarse Mesh Transport Method	5
3 Method	11
Response Coefficient Generation	11
Response Coefficient Library	16
Deterministic Solution Construction	18
4 Results	20
Benchmark Problems	20
Response Expansion Coefficient Library	34
Mesh Boundary Analysis	36
Convergence and Acceleration of Solution Method	40
Expansion Order Analysis	43
Control Rod Worth Analysis	59
5 Conclusions	80
APPENDIX A	82

REFERENCES	84
VITA	86

LIST OF TABLES

	Page
Table 1: HTTR, some rods in (SRI)	38
Table 2: HTTR, all rods out (ARO)	38
Table 3: Core A	38
Table 4: Convergence and acceleration options	41
Table 5: Accelerated SRI HTTR core	42
Table 6: Accelerated ARO HTTR core	42
Table 7: SRI HTTR	44
Table 8: ARO HTTR	44
Table 9: Core A	50
Table 10: Core B	51
Table 11: Core C	51
Table 12: Eigenvalue results for rod worth exercise	62
Table 13: Core A results	82
Table 14: Core B results	83
Table 15: Core C results	83

LIST OF FIGURES

	Page
Figure 1: Top mesh face coordinates	15
Figure 2: Fuel, control, solid reflector, and cooled reflector block geometry	21
Figure 3: ARO HTTR fission density	24
Figure 4: SRI HTTR fission density	25
Figure 5: Core A, levels 1-5	28
Figure 6: Core B, levels 1-5	28
Figure 7: Core C, levels 1-5	29
Figure 8: Core A pin fission density map	31
Figure 9: Core B pin fission density map	32
Figure 10: Core C pin fission density map	33
Figure 11: SRI HTTR errors at (2,2,2,2)	45
Figure 12: SRI HTTR errors at (4,4,4,4)	46
Figure 13: ARO HTTR errors at (2,2,2,2)	47
Figure 14: ARO HTTR errors at (4,4,4,4)	48
Figure 15: Core A errors at (2,2,2,2)	52
Figure 16: Core A errors at (4,4,4,4)	53
Figure 17: Core B errors at (2,2,2,2)	54
Figure 18: Core B errors at (4,4,4,4)	55
Figure 19: Core C errors at (2,2,2,2)	57
Figure 20: Core C errors at (4,4,4,4)	58
Figure 21: Color map for all error figures	59
Figure 22: Rod worth test core	60

Figure 23: Color key to pin fission density results illustrated in Figures 24-39	63
Figure 24: Rod worth test core, rod position 0	64
Figure 25: Rod worth test core, rod position 1	65
Figure 26: Rod worth test core, rod position 2	66
Figure 27: Rod worth test core, rod position 3	67
Figure 28: Rod worth test core, rod position 4	68
Figure 29: Rod worth test core, rod position 5	69
Figure 30: Rod worth test core, rod position 6	70
Figure 31: Rod worth test core, rod position 7	71
Figure 32: Rod worth test core, rod position 8	72
Figure 33: Rod worth test core, rod position 9	73
Figure 34: Rod worth test core, rod position 10	74
Figure 35: Rod worth test core, rod position 11	75
Figure 36: Rod worth test core, rod position 12	76
Figure 37: Rod worth test core, rod position 13	77
Figure 38: Rod worth test core, rod position 14	78
Figure 39: Rod worth test core, rod position 15	79

SUMMARY

A new whole-core transport method is described for 3-D hexagonal geometry. This is an extension of a stochastic-deterministic hybrid method which has previously been shown highly accurate and efficient for eigenvalue problems. Via Monte Carlo, it determines the solution to the transport equation in sub-regions of reactor cores, such as individual fuel elements or sections thereof, and uses those solutions to compose a library of response expansion coefficients. The information acquired allows the deterministic solution procedure to arrive at the whole core solution for the eigenvalue and the explicit fuel pin fission density distribution more quickly than other transport methods. Because it solves the transport equation stochastically, complicated geometry may be modeled exactly and therefore heterogeneity even at the most detailed level does not challenge the method.

In this dissertation, the method is evaluated using comparisons with full core Monte Carlo reference solutions of benchmark problems based on gas-cooled, graphite-moderated reactor core designs. Solutions are given for core eigenvalue problems, the calculation of fuel pin fission densities throughout the core, and the determination of incremental control rod worth. Using a single processor, results are found in minutes for small cores, and in no more than a few hours for a realistically large core. Typical eigenvalues calculated by the method differ from reference solutions by less than 0.1%, and pin fission density calculations have average accuracy of well within 1%, even for unrealistically challenging core configuration problems. This new method enables the

accurate determination of core eigenvalues and flux shapes in hexagonal cores with efficiency far exceeding that of other transport methods.

1. INTRODUCTION

In an effort to improve fuel utilization and therefore power plant economics, the current trend in reactor core design is one of greater heterogeneity at the core and fuel element level. Furthermore, some of the proposed designs for future reactors—such as the Very High Temperature Reactor (VHTR)—are implemented in hexagonal geometry. With a more heterogeneous core and a flux distribution exhibiting strong angular effects, whole core diffusion methods are inadequate to accurately resolve the pin-level fission density map over the core. A new method which relies wholly on transport theory and is tailored specifically to the hexagonal lattice structure of prismatic block reactors and fast reactors is necessary to construct the solution to reactor criticality problems. One which solves the core problems in a matter of minutes rather than days would be a valuable tool.

The coarse mesh transport (COMET) method has previously been demonstrated in Cartesian geometry in three dimensions, and in hexagonal geometry for 2-D core configurations. In all cases, it has been proven accurate and determines detailed solutions far more quickly than other whole core transport methods. In this paper, the COMET method is extended to 3-D hexagonal geometry. Chapter 2 summarizes hexagonal reactor solution methods and gives background to the COMET method. Chapter 3 details the new features and modifications made for 3-D problems. Verification of the method is conducted using a series of benchmark problems in Chapter 4. Finally, conclusions are drawn and future opportunities using this method are presented and discussed in Chapter 5.

2. BACKGROUND

In this section, some methods used for determining the eigenvalue and pin power distribution of reactors are discussed. First, methods which treat hexagonal geometry in reactor cores are examined. A summary of the existing coarse mesh transport method (Ilas and Rahnema, 2003; Mosher and Rahnema, 2006; Zhang and Rahnema, 2012; Connolly et al., 2012), which is the basis of the new method described in this paper, follows.

2.1. Reactor Analysis in Hexagonal Cores

The development of methods specifically tailored to hexagonal reactor core problems is not strictly a new or recent initiative. Historically, whole core eigenvalue problems have been solved using diffusion theory. By dividing the core into homogenized nodes, few-group cross sections are generated, and the solution to the core problem is found via the diffusion equation. Several methods have been developed to treat three dimensional reactor core configurations in hexagonal lattice structure. One method is the response matrix method, where the eigenvalue and node-averaged flux can be determined through the core in seconds (Gheorghiu, 1989). Another uses conformal mapping of the hexagonal faces and expands the flux in terms of orthogonal polynomials to solve the whole core problem (Zimin and Baturin, 2002). The HEXPEDITE method (Fitzpatrick and Ougouag, 1992) couples nodes by net currents at their interfaces. One drawback to nodal diffusion methods is that while they have calculated the block-averaged flux, they do not explicitly calculate the pin power map of the core. To address this shortcoming,

methods have been developed to reconstruct the pin power profile (Singh and Kumar, 1995). All of these methods have been shown accurate in solving eigenvalue problems when compared to fine mesh diffusion benchmarks.

However, in regions of high flux gradients, strongly absorbing media, a pronounced angular distribution of the flux, and asymmetric fuel and control elements, use of the diffusion equation will not yield a desired level of accuracy. Diffusion requires the homogenization of regions, which introduces errors due to the necessary erasing of fine detail in very heterogeneous systems. As regions with such attributes are present in the prismatic VHTR core, it becomes necessary to look to transport methods to supply a high level of accuracy in complex cores. Nodal transport methods, such as one described by Carrico et al. (1992), were first used to address some of the shortcomings associated with diffusion methods, but even these lacked fully heterogeneous treatment of the hexagonal meshes. Since computational memory capacity and processing speed have improved dramatically in recent years, transport-based methods capable of modeling explicit heterogeneity in high fidelity have accordingly become more feasible.

One such transport method developed for solving the VHTR core is DeCART. It has been shown accurate in eigenvalue calculations to within 300 pcm of a whole core Monte Carlo solution for both homogenized fuel and heterogeneous fuel compacts (Lee et al., 2010; Lee and Yang, 2011). Block-averaged power levels also were found to within 1% in most cases and within 2% with all rods in. Runtime figures were not provided in the reference, although two dimensional calculations using this method have previously been

published with stated runtimes on the order of a few hours (Cho et al., 2007). Similar accuracy was achieved in modeling the heterogeneity at the VHTR pin level using the AGENT neutron transport method (Gert and Jevremovic, 2008), although this study was conducted at the block, rather than core, level. For all of the preceding transport references, a method of characteristics approach was taken.

The accuracy of the solutions determined by the transport methods listed was quantified by comparing the result with a Monte Carlo benchmark calculation. Monte Carlo solutions are readily determined for any geometry; however, a low level of statistical uncertainty requires a great deal of computation time. The benchmarks used to evaluate the method presented in this paper were solved using the Monte Carlo code MCNP5 (X-5 Monte Carlo Team, 2005), and they generally required many hours on a large cluster of parallel computers.

While runtime figures are not often mentioned in the references cited, it is apparent that whole core transport calculations modeling fine detailed heterogeneity require prohibitive computational time and resources when accuracy comparable to Monte Carlo is desired. Clearly, the opportunity exists to make improvements in computational runtime while maintaining the accuracy inherent in transport methods. The development of whole core transport methods for 3-D hexagonal lattice structures with a high level of heterogeneity is a field in its infancy.

2.2. The Hybrid Coarse Mesh Transport Method

The purpose of this paper is an extension of a hybrid method which combines the geometric flexibility of a Monte Carlo simulation with the speed inherent in coarse mesh deterministic processes (Mosher and Rahnema, 2003). This is an incident flux response expansion method, typically referred to as COMET (an acronym for coarse mesh transport), which determines the solution to the transport equation throughout the entire core. It is a novel method, representing a significant advancement on the class of response matrix solution techniques, where the core problem is solved by combining solutions to sections of the core for which solution information is already available. A discussion on traditional high-order response matrix theory, including a survey of additional references, can be found in Stamm'ler and Abbate (1983). COMET advances from older response matrix methods in its reliance solely on transport theory, its high fidelity modeling of exact core and lattice specifications without any use of homogenization techniques, its explicit treatment of pin-level heterogeneity, allowing for the solution of the pin fission density profile throughout the core, its arbitrarily high order accuracy of mesh boundary conditions, and its adaptability to many different reactor types and other physical systems (e.g., Forget, 2006; Blackburn, 2009; Zhang and Rahnema, 2012). It is summarized here in order to give the proper background information.

The steady state neutron transport equation over an arbitrary volume may be written as

$$\begin{aligned} \hat{\Omega} \cdot \nabla \Phi(\vec{r}, \hat{\Omega}, E) + \sigma_t(\vec{r}, E) \Phi(\vec{r}, \hat{\Omega}, E) = & \quad (1a) \\ \int_0^\infty dE' \int_{4\pi} d\hat{\Omega}' \sigma_s(\vec{r}, \hat{\Omega}', E' \rightarrow \hat{\Omega}, E) \Phi(\vec{r}, \hat{\Omega}', E') & \\ + \frac{1}{4\pi} \chi(\vec{r}, E) \int_0^\infty dE' \int_{4\pi} d\hat{\Omega}' \frac{\nu \sigma_f(\vec{r}, E')}{k} \Phi(\vec{r}, \hat{\Omega}', E'), & \end{aligned}$$

with a general boundary condition

$$\Phi(\vec{r}_b, \hat{\Omega}, E) = B \Phi(\vec{r}_b, \hat{\Omega}', E'), \quad \hat{n} \cdot \hat{\Omega} < 0 \quad \text{and} \quad \hat{n} \cdot \hat{\Omega}' > 0. \quad (1b)$$

In the preceding equations, Φ represents the angular flux of neutrons, which are defined for every point \vec{r} within the volume and traveling with energy E in direction $\hat{\Omega}$. The macroscopic neutron reaction cross sections are represented by σ and specified as the total, scattering, and fission cross sections by the subscripts t , s , and f , respectively. Each fission reaction creates an average of ν neutrons from the χ energy spectrum. The core multiplication factor is denoted by k . In Eq. (1b), \vec{r}_b represents a point along the boundary of the volume which has outward normal \hat{n} . The boundary operator B is left arbitrary; Mosher and Rahnema (2006) have provided example boundary conditions.

The COMET method divides the volume of interest, e. g., a reactor core, into a set of discrete sub-volumes, such as fuel elements. These sub-volumes will herein be referred to as coarse meshes. The transport equation within each coarse mesh i can be written:

$$\hat{\Omega} \cdot \nabla \psi(\vec{r}, \hat{\Omega}, E) + \sigma_i(\vec{r}, E) \psi(\vec{r}, \hat{\Omega}, E) = \quad (2a)$$

$$\int_0^\infty dE' \int_{4\pi} d\hat{\Omega}' \sigma_s(\vec{r}, \hat{\Omega}', E' \rightarrow \hat{\Omega}, E) \psi(\vec{r}, \hat{\Omega}', E') \\ + \frac{1}{4\pi} \chi(\vec{r}, E) \int_0^\infty dE' \int_{4\pi} d\hat{\Omega}' \frac{\nu \sigma_f(\vec{r}, E')}{k} \psi(\vec{r}, \hat{\Omega}', E') + Q_i(\vec{r}, \hat{\Omega}, E).$$

Here the solution to the mesh-level transport equation is represented as ψ ; it represents the exact solution to Eq. (1a) within mesh i when it is equal to Φ within the particular mesh domain. The term k appearing in Eq. (2a) is fixed; it represents the whole core multiplication factor and not an eigenvalue of Eq. (2a). It is present to scale the fission term, as an implicit treatment of the fission source allows coarse meshes to be modeled at different values of the core multiplication factor. An external source may be denoted by Q_i , however, in solving the mesh-level transport problem presented here, this term is equal to no volumetric source. Instead, an incoming boundary condition on the mesh surface is given, so that each individual mesh i is coupled to the neighboring meshes j by a particle balance at the mesh interface \vec{r}_{ij} :

$$\psi_i(\vec{r}_{ij}, \hat{\Omega}_i, E) = \psi_j(\vec{r}_{ij}, \hat{\Omega}_j, E), \quad (2b)$$

or equivalently,

$$J_i(\vec{r}_{ij}, \hat{\Omega}_i, E) = J_j(\vec{r}_{ij}, \hat{\Omega}_j, E). \quad (2c)$$

For both boundary conditions, \hat{n}_i refers to the outward normal of mesh i , and $\hat{n}_i = -\hat{n}_j$.

Most recently, Zhang and Rahnama (2012) used the angular flux boundary condition given by Eq. (2b) and Connolly et al. (2012) utilized the angular current boundary condition in (2c).

Solutions to the set of fixed-source problems defined by Eq. (2) are found, given some specification of an incoming source. Since both incident angular flux and angular current have been shown to be valid and accurate boundary conditions, let β represent the boundary condition of choice to maintain the generality of the method description. The incident boundary condition on face f of mesh i may be expanded exactly in terms of a complete set of functions Γ^a which are orthogonal over the half-space defined by the coarse mesh face,

$$\beta_{if}^{\pm}(\vec{r}, \hat{\Omega}, E) = \sum_{a=1}^{\infty} c_{if}^a \Gamma^a(\vec{r}, \hat{\Omega}, E). \quad (3)$$

The COMET method determines the outgoing condition at the mesh boundary as a response to a given incoming boundary specification, which is chosen to be a unit value of the orthogonal distribution of the a^{th} order on one mesh face and vacuum boundary conditions on all other sides. This process is repeated for each face of each unique mesh present in the core for every incoming expansion order a . For the sake of practicality, the sum must be truncated at some order A which may be chosen to optimize the accuracy of the solution and the computational investment necessary. Although any transport method

may be used to solve the set of fixed-source problems presented by Eq. (2), a Monte Carlo method has been shown superior to others which have been implemented, as exact geometric and material information can be modeled with no approximation (Mosher and Rahnema, 2003). The result is a library of response expansion coefficients, where the outgoing response to every incoming boundary source to a desired maximum order is tabulated for every face of every unique mesh in the core. Also determined are reaction rate responses to the incoming sources; these are used in the construction of the core solution. These calculations are conducted at several different values of k , so that a full database of response functions and coefficients can be used to solve any core problem over a range of eigenvalues without any prior knowledge of the core multiplication factor: the response function for a desired value of k can be determined by simply interpolating the existing response functions between the values of k which have been solved.

The whole core problem is solved via an iterative procedure. Starting with an initial guess for the multiplication factor and the flux or current distribution throughout the core, a deterministic sweeping method calculates new outgoing response functions to the assumed incoming sources, using the pre-computed library and linear superposition of the response functions. The mesh interface conditions are updated and the process continues until a desired convergence criterion is reached, at which point the eigenvalue is calculated via the neutron balance method. The inner iterations on the mesh interfaces and outer iterations on the multiplication factor continue until the latter converges, at

which point the problem has been solved, and the full solution Φ to Eq. (1) is known over the entire core domain.

3. METHOD

The coarse mesh transport method for 3-D hexagonal geometry is described in this section. The means of calculating response expansion coefficients is illustrated. The creation of a library of expansion coefficients is then discussed. A brief overview of the deterministic solution process completes the section.

3.1. Response Coefficient Generation

Because the solution to the whole core problem is not initially known, it is necessary to pre-compute the set of response expansion coefficients which describe the neutron distribution at mesh interfaces. The fixed source problem over each unique mesh is solved, where the governing equation is Eq. (2a), vacuum boundary conditions are given on all faces f' of the mesh except for one face $f \neq f'$, and the incoming boundary condition β source is written:

$$\begin{aligned} \beta_{ij}^- (\vec{r}, \hat{\Omega}, E) &= \Gamma^a (\vec{r}, \hat{\Omega}, E) \\ \beta_{ij'}^- (\vec{r}, \hat{\Omega}, E) &= 0 \end{aligned} \quad (4)$$

A fixed-source problem is solved for every unique mesh i and face f , and for all desired expansion orders a . The outgoing response on each of the eight faces f' is determined as β^+ , and the response expansion coefficients can be found

$$c_{ij}^{a'} = \iiint \beta_{ij}^+(\bar{r}, \hat{\Omega}, E) \Gamma^{a'}(\bar{r}, \hat{\Omega}, E) d\bar{r} d\hat{\Omega} dE. \quad (5)$$

Two mathematical definitions for the boundary conditions, orthogonal function set, and expansion coefficients are addressed in the following subsection.

3.1.1. Mesh Boundary Condition

Some method developers, including the author, have previously used the incident angular current as the boundary condition for coarse mesh fixed source problems (Forget, 2006; Connolly et al., 2012). This treatment expanded the angular current as

$$\beta_{ij}^{\pm}(\bar{r}, \hat{\Omega}, E) = J_{ij}^{\pm}(\bar{r}, \hat{\Omega}, E). \quad (6)$$

The orthogonal set used in this case was the product of the shifted Legendre polynomials over position and angle, as given by Eq. (7):

$$\Gamma^a(\bar{r}, \hat{\Omega}, E) = \tilde{P}_l^{[0,U]}(u) \tilde{P}_m^{[0,V]}(v) \tilde{P}_n^{[-1,1]}(\mu) \tilde{P}_p^{[0,\pi]}(\varphi) \delta(E - E_g). \quad (7)$$

From this point onward, the position \bar{r} will be given in terms of the spatial variables u , which is defined along one direction of the mesh face from 0 to U , and v , which represents the position in the other dimension of the mesh face between 0 and V . The direction $\hat{\Omega}$ will be given in terms of the cosine μ of the polar angle θ from the mesh

face, and the azimuthal angle φ . The delta function is used to specify a multigroup (discrete) energy treatment over G groups.

The expansion coefficient for the angular current J^+ exiting face f' of mesh i , in group g' , with spatial and angular expansion orders l' , m' , n' , and p' , as a response to an incoming current on face f in group g being of spatial and angular expansion orders l , m , n , and p , is determined as follows:

$$c_{i,f',g',l',m',n',p'}^{ifglmnp} = \int \int \int \int J_{if'}^+(u, v, \mu, \varphi, E) \tilde{P}_l^{[0,U]}(u) \tilde{P}_{m'}^{[0,V]}(v) \tilde{P}_{n'}^{[-1,1]}(\mu) \tilde{P}_{p'}^{[0,\pi]}(\varphi) \delta(E - E_{g'}) du dv d\mu d\varphi dE \quad (8)$$

However, coincident research in water-moderated reactors has suggested that expanding the incident angular flux in a combination of the Legendre polynomials over the spatial domain, the Legendre expansion of the cosine of the azimuthal angle, and the Chebyshev polynomial of the second kind defined over the cosine of the polar angle yields better results, as such an expansion preserves not only the particle balance at the mesh interfaces, but also describes an isotropic flux at the boundary when the 0th order is used (Zhang and Rahnema, 2012). As the angular flux has a large isotropic component in highly scattering reactors, it has been suggested that expanding the angular flux requires a lower order of expansion to achieve similarly accurate results when compared to an angular current expansion at the mesh boundary. The use of angular flux as the physical quantity of expansion changes the boundary condition to

$$\beta_{if}^{\pm}(\vec{r}, \hat{\Omega}, E) = \psi_{if}^{\pm}(\vec{r}, \hat{\Omega}, E), \quad (9)$$

and the expansion function is defined

$$\Gamma^a(\vec{r}, \hat{\Omega}, E) = \tilde{P}_l^{[0,U]}(u) \tilde{P}_m^{[0,V]}(v) U_n^{[-1,1]}(\mu) \tilde{P}_p^{[0,\pi]}(\cos \varphi) \delta(E - E_g). \quad (10)$$

The coefficients may be found

$$c_{i,f',g',l',m',n',p'}^{ifglmnp} = \quad (11)$$

$$\int \int \int \int \psi_{if'}^+(u, v, \mu, \varphi, E) \tilde{P}_l^{[0,U]}(u) \tilde{P}_m^{[0,V]}(v) U_n^{[-1,1]}(\mu) \tilde{P}_p^{[0,\pi]}(\cos \varphi) \delta(E - E_g) du dv d\mu d\varphi dE$$

An investigation of these boundary conditions using benchmarks based on the HTTR and on a smaller test core is conducted in Section 4.3. From these results, it may be concluded that a response expansion of the angular flux presents the superior boundary condition, and it is used for the remainder of the paper.

3.1.2. Treatment of Hexagonal Geometry

Previously, the COMET method has been used in 2-D hexagonal geometry and in both 2-D and 3-D Cartesian geometry. In these cases, the faces of the coarse meshes are represented as either one dimensional line segments (in 2-D problems) or as two dimensional rectangular faces (in 3-D Cartesian geometry). A hexagonal prismatic mesh

has six rectangular faces, for which the existing method of expansion in the spatial variables u and v is adapted without modification. However, the top and bottom hexagonal faces present a new challenge: a new scheme is necessary in order to describe the mesh faces in two spatial dimensions.




Figure 1. Top mesh face coordinates

The expansion over the two spatial variables is maintained over the hexagonal mesh faces. Figure 1 illustrates the top face of a hexagon, indicating the u and v axes over the mesh face. The u variable continues to indicate the position along the x axis, however, the v variable must be transformed from describing simply the y coordinate to describing the position in that direction with respect to the edges of the mesh face. This transformation means that the length from 0 to V in the v direction is no longer constant, but is a function of position in the u direction; when u is equal to 0.5, V represents a length which is double that when u is equal to 0 or 1. Based on the results presented in Chapter 4, this does not appear to result in inaccuracies within the solutions determined by the method.

3.2. Response Coefficient Library

Upon the solution of the set of fixed-source problems, the response expansion coefficients determined either by Eq. (8) or Eq. (11) are collected and placed in a library. The library contains coefficients representing each of the desired expansion orders at each of the discrete values of the multiplication factor for which fixed-source problems were solved. The number of angular flux response expansion coefficients in the library for each mesh is determined by

$$K \times (F \cdot G \cdot \Xi \cdot Y) \times (F \cdot G \cdot \Xi \cdot Y), \quad (12)$$

where K different values of the core eigenvalue have been simulated. In coarse meshes for which there exists no fission source, K is simply equal to one. For all hexagonal meshes, the quantity F which represents the number of mesh faces is equal to eight.

Equations (8) and (11) suggest that coefficients must be found for all values of l , m , n , and p from the set of values up to and including the value chosen for truncation of the orthogonal set, for example, $m \in [0, M]$. However, coefficients where the combined sum of the spatial or angular expansion orders is greater than the value chosen for truncation of either variable have previously been found not to contribute meaningfully to the solution (Zhang and Rahnema, 2012). Therefore, the terms Ξ and Y have been introduced to represent the number of spatial and angular expansion terms, respectively. The number of terms such that $l \leq L$, $m \leq M$, and $(l + m) \leq \max(L, M)$ is represented by

Ξ ; the corresponding variable for the angular terms $n \leq N$, $p \leq P$, and $(n + p) \leq \max(N, P)$ is denoted Y . The prohibitive additional memory and computation time requirements to compile a library with these extra terms $\xi > \Xi$ and $\nu > Y$ have impeded any additional investigation in this paper, although due to the highly accurate results presented in the following section, no such investigation was considered necessary.

3.2.1. Selection of meshes

It has been stated throughout this work that response expansion coefficients are determined for only unique meshes within the core. The optimal choice of coarse mesh has been found to be a fuel block (or assembly); since few different block designs are generally present within a core, the number of response expansion coefficients necessary for a complete library is kept relatively low. As operating reactor cores typically consist of many blocks which, at beginning of life, are identically constructed and composed, building a library of only unique meshes rather than each mesh individually allows this method to be efficient without a loss of practicality. The uniqueness of a coarse mesh refers to its block design, but not to its position in the core. Thus, for a block of a particular design, response expansion coefficients need be calculated only once; any number of blocks with that same design may then be placed at any position in the core without necessitating additional response expansion coefficient calculations. For example, test cores are presented as benchmark problems in section 4.1.2. The first of these test cores is constructed of 185 blocks. However, each block is identical in design;

the same mesh is simply repeated 185 times within the core. It is foreseen that the method will be used for reactor calculations throughout the life cycle of the core. The library will then become a function of reactor state parameters, such that response expansion coefficients will be stored for a set of discrete temperatures, burnup steps, and so forth, requiring only a simple interpolation scheme to select the proper value between coefficients in the library.

3.3. Deterministic Solution Construction

The final step in the method is the deterministic iterative procedure which composes the solution to the whole core problem. Starting with an initial guess for the eigenvalue and flux distribution, the outgoing response to the hypothesized incoming current is determined for each mesh, proceeding in a given sweep order. It was previously shown (Connolly et al., 2012) that the sweep order did not affect the accuracy of the solution, but that sweeping from the center of the core outward yielded results more efficiently than other sweeping methods which had been proposed. The method retains this outward sweeping scheme, in which new response expansion coefficients are first determined for the angular flux exiting the center mesh, based on an initial guess for the angular flux entering that mesh. The process repeats for each mesh in the core, with the sequence of meshes ordered in an outward sweep. Determination of an optimum sweep order is beyond the scope of this paper.

3.3.1. Convergence of solution

Once the iterative procedure converges in flux distribution and eigenvalue, the problem is considered solved. While the determination of appropriate convergence criteria is a decision for the user of the method, several options are investigated in section 4.4.

Convergence criteria should be tight enough to ensure a solution is reached; however, given the inherent uncertainty in the solutions due to the stochastic nature of response expansion coefficient generation and the fact that due to the truncation of the infinite sum of orthogonal functions, the solution will not be exact, the convergence criteria should be chosen which are not significantly tighter than the expected uncertainty and deviation from the exact solution.

4. RESULTS

In this section, numerical results are presented in order to evaluate the new method. A series of core configurations is presented, and COMET's ability to solve these problems is compared to a whole core Monte Carlo solution.

4.1. Benchmark Problems

In order to verify the accuracy of COMET, it is tested against several whole core Monte Carlo benchmarks. The first is a benchmark based on Japan's High Temperature Test Reactor (HTTR). Following are three small cores designed from components of gas-cooled reactors, but intended to present more difficult problems to a solution method. The benchmark problems are described in detail below.

4.1.1. High Temperature Test Reactor

The first benchmark problem presented is a stylized depiction of Japan's High Temperature Test Reactor (HTTR) (Zhang et al., 2011). It was chosen as a representative example of a gas-cooled thermal reactor. The core exhibits block and core-level heterogeneity, it accurately models neutronic behavior characteristic of a graphite moderated core, and it is built with a hexagonal lattice structure. Furthermore, the 2-D version of this problem has previously been solved using the COMET method (Connolly et al., 2012).

The HTTR problem uses a six-group cross section library. Seven unique fuel enrichments, control blocks with rods either inserted or withdrawn, two compositions of solid reflector blocks, and reflector blocks with coolant channels running through them are present in the core; a total of twelve unique meshes is therefore represented. A cross-section of each major block design (fuel, control, solid reflector, and cooled reflector) is presented in Figure 2. Gray represents graphite, yellow represents fuel material, and green is boron carbide absorber. Helium coolant, treated simply as vacuum for the sake of neutronic modeling, flows through the white spaces in the figure. The blocks are uniform in the axial direction.




Figure 2. Fuel, control, solid reflector, and cooled reflector block geometry.

Solutions to two different core configurations were found. The first problem is the simpler one, that being the all-rods-out configuration. The second core is the partially controlled configuration which is near critical operating condition. As the reference presents only the partially controlled core with a corrugated boundary, and as the method

developed within this paper is designed to treat only entire blocks and not fractions thereof, the benchmark problems were solved again for this paper using corrugated boundaries and specifying no incoming flux as a boundary condition for both cores. This is the only deviation from the core definition specified by Zhang et al. (2011). The new benchmarks were solved using standard MCNP5 in parallel.

The first problem represents the core configuration with all rods removed. The core was modeled in one-sixth symmetry. 4,000 active cycles of 250,000 neutron histories were simulated after an initial 2,000 cycles were run to converge the fission source. The multiplication factor was found to be 1.09877 with an uncertainty of 0.00002. Pin fission densities for all 4950 pin segments in the core were calculated; the uncertainty in the calculations ranged from 0.02-0.06%. The pin fission density distribution is illustrated in Figure 3. This calculation required 75009 minutes of computing time, divided across 32 processors in parallel. This corresponds to 52 days of cpu time, assuming a scaling factor of unity. Computation times are given for comparative purpose within this work, as all computations presented herein used 2 GHz processors.

The second problem represents the core configuration with some rods in, to model the core as closely as possible to a critical operating state. Like the previous core, this configuration was also modeled in one-sixth symmetry. Similarly, the same amount of active and inactive cycles was simulated with an identical number of particle histories per cycle. The multiplication factor was found to be 0.99740 with an uncertainty of 0.00002. The uncertainty in the pin fission density calculations ranged from 0.02-0.08%. The

range of fission density values is presented by Figure 4. This calculation required 79479 minutes of computing time, divided across 32 processors in parallel. This corresponds to 55 days of cpu time, assuming a scaling factor of unity.

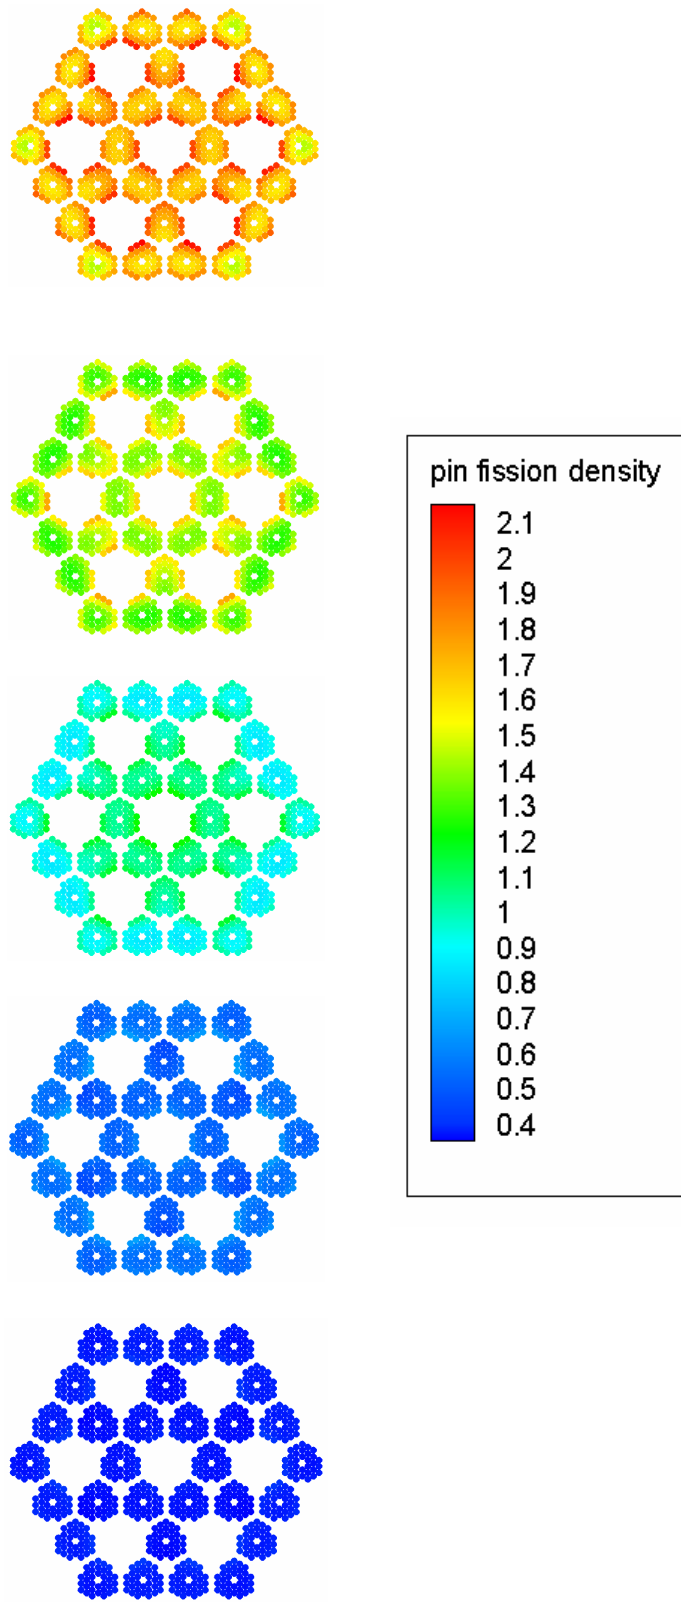


Figure 3. ARO HTTR fission density

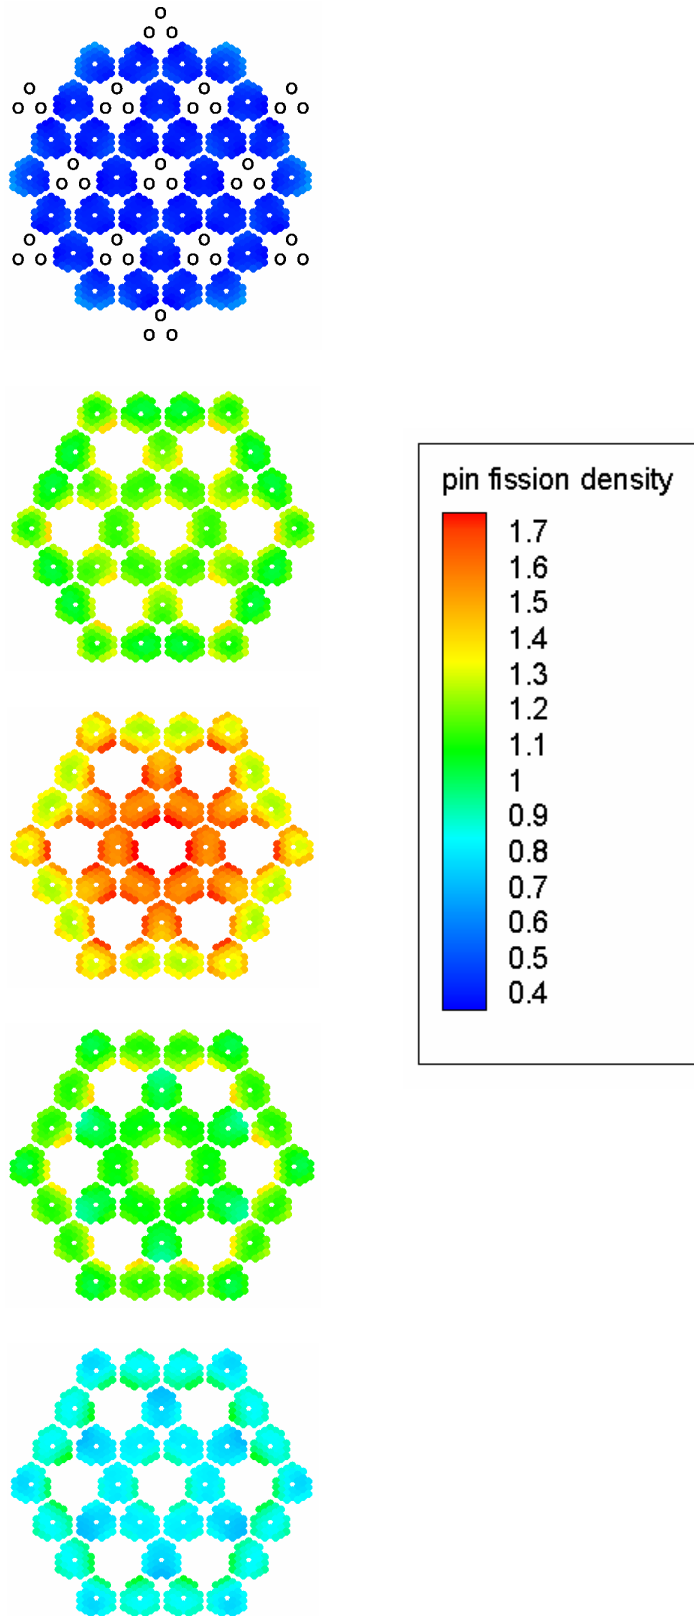


Figure 4. SRI HTTR fission density

4.1.2. Heterogeneous Test Cores

Tests of the method have been conducted on a set of test cores which have been designed to challenge the capabilities of the method while minimizing the need for a large response expansion coefficient library. Although these smaller cores reduce computation time when compared to a realistic operating reactor core, their design is intended to provide a greater challenge to the method by introducing unrealistically dramatic heterogeneity at the core level, sharp flux peaking and gradients, and high leakage.

The test problems have been inspired by the HTTR benchmark, and have previously appeared in an earlier work (Connolly and Rahnema, 2012c). All three cores use the block specifications described in the HTTR benchmark, however, the core geometry is quite different. Only three unique block types appear in the following core specifications: the fuel block specified as enrichment type 5 in the reference (Zhang et al., 2011), the control block with rods inserted, and the permanent reflector block. All three test cores are composed of axial levels consisting of 37 blocks: three rings of blocks around a central block; each core is five axial levels in height, for a total of 185 blocks. The cores are depicted in Figures 5-7, where white hexagonal lattice elements represent fuel blocks, gray is used to symbolize reflector blocks, and blocks featuring control material are colored black. The core illustrations previously appeared in an earlier work (Connolly and Rahnema, 2012c).

As was done for the HTTR solutions, reference solutions for the test cores have been determined using MCNP5. The solution for Core B required 2,000 active neutron history cycles after an initial 500 cycles were skipped to create a converged fission source. 250,000 neutron histories were simulated in each cycle, for a total of 125 million inactive neutron histories and 500 million active histories. Cores A and C modeled twice as many active neutron histories; this was necessary to reduce uncertainty in the pin fission density calculations at the core periphery.

Core A is composed solely of fuel blocks; it is presented first as the simplest core design to model. It is, however, a challenge to solve; the configuration leads to high leakage and steep flux gradients within the core, a global peaking factor of 3.59, and a minimum pin fission density of 0.0285. The core contains 185 fuel blocks, and 6105 pin regions. The multiplication factor of the core was calculated at 0.95333 ± 0.00002 . The pin fission density calculations were determined to an average statistical uncertainty of 0.121%. Figure 8 depicts the power map of the core. The reference solution was calculated over nearly eight days using 64 processors in parallel, corresponding to 508 days of total computing time assuming a scaling factor of one.




Figure 5. Core A, levels 1-5




Figure 6a. Core B, levels 1 and 5




Figure 6b. Core B, levels 2 and 4




Figure 6c. Core B, level 3

Figures 5-7: core diagrams. Fuel blocks are white, reflectors gray, and controls black.




Figure 7a. Core C, level 5




Figure 7b. Core C, level 4




Figure 7c. Core C, level 3




Figure 7d. Core C, level 2




Figure 7e. Core C, level 1

Core B introduces new block types. The core is now fully enveloped in 128 reflector blocks, with 56 fuel blocks surrounding a solitary control block in the center of the core. The 1848 fuel pin regions have a flatter flux profile than the previous core, with fission density ranging from a peak value of 1.23 to a minimum value of 0.708. The fission density map appears in Figure 9, where the pins are colored accordingly and the black circles represent the position of the annular control rods present. The average uncertainty in the pin fission density calculations was 0.074%. Core B has a multiplication factor of 0.90757 ± 0.00003 . Using 64 processors running in parallel, the calculations took nearly 3.5 days; this corresponds to approximately 220 total days of computing time.

Core C is intended to rigorously challenge the method, as it includes strong heterogeneity at the core level, extreme asymmetry, and a dramatically irregular flux profile. The core includes 85 fuel blocks, which contain a total of 2805 fuel pin regions, whose fission densities have been determined to an average uncertainty of 0.254%, but a maximum of 4.000% in some very low-power pins. An additional 76 reflector blocks and 24 control blocks are distributed throughout the core in a haphazard and unrealistic manner. The fission density distribution is illustrated in Figure 10; the fission density values vary from a peak of 2.75 to a low value of 0.00273. Computing the pin fission density profile and the core multiplication factor of 0.85405 ± 0.00002 required over 8.5 days on 64 processors in parallel, or 545 days of total computing time.

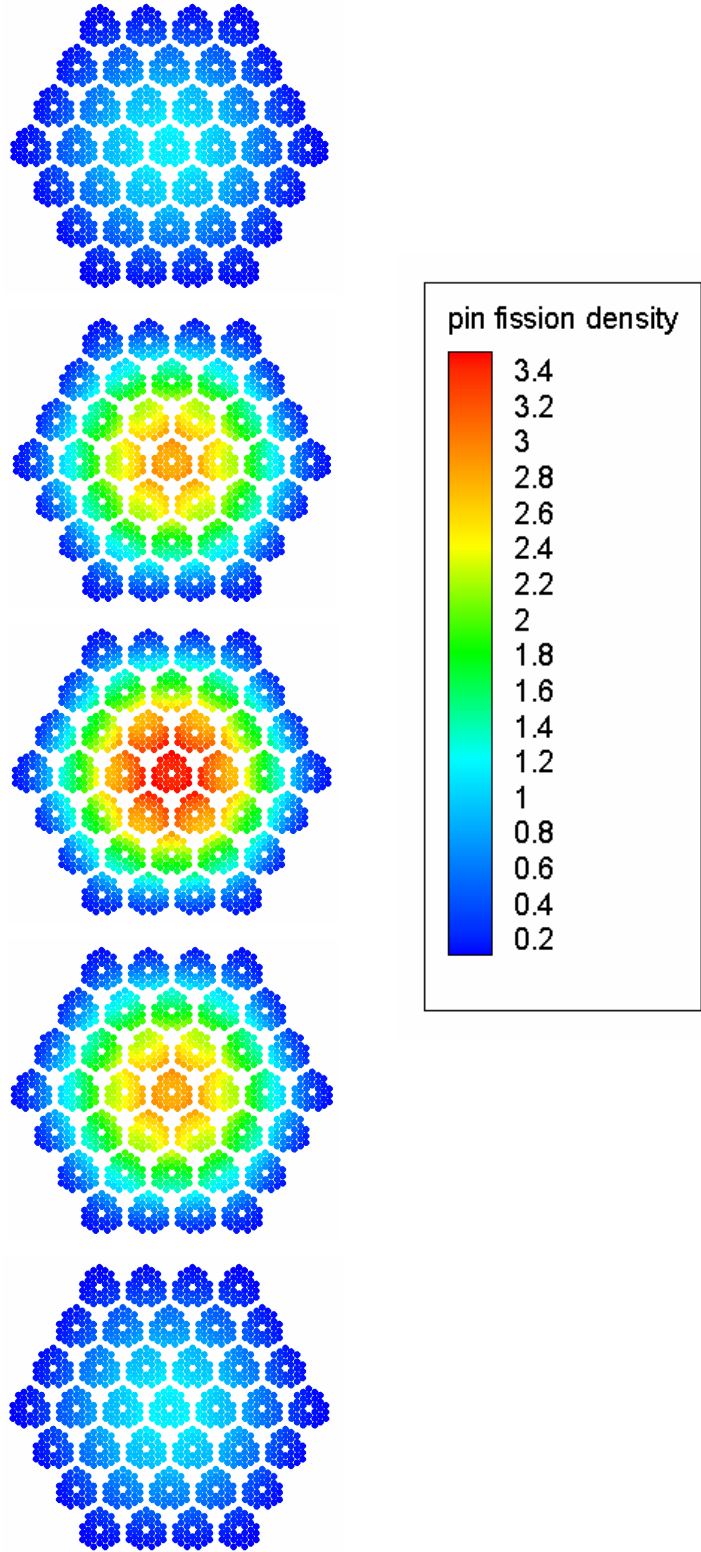


Figure 8. Core A pin fission density map.

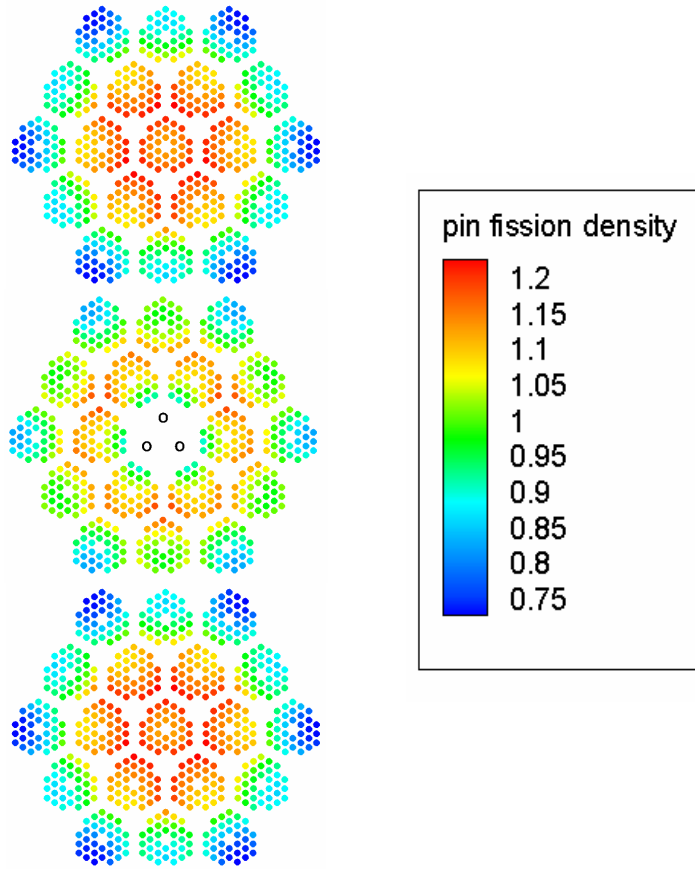


Figure 9. Core B pin fission density map

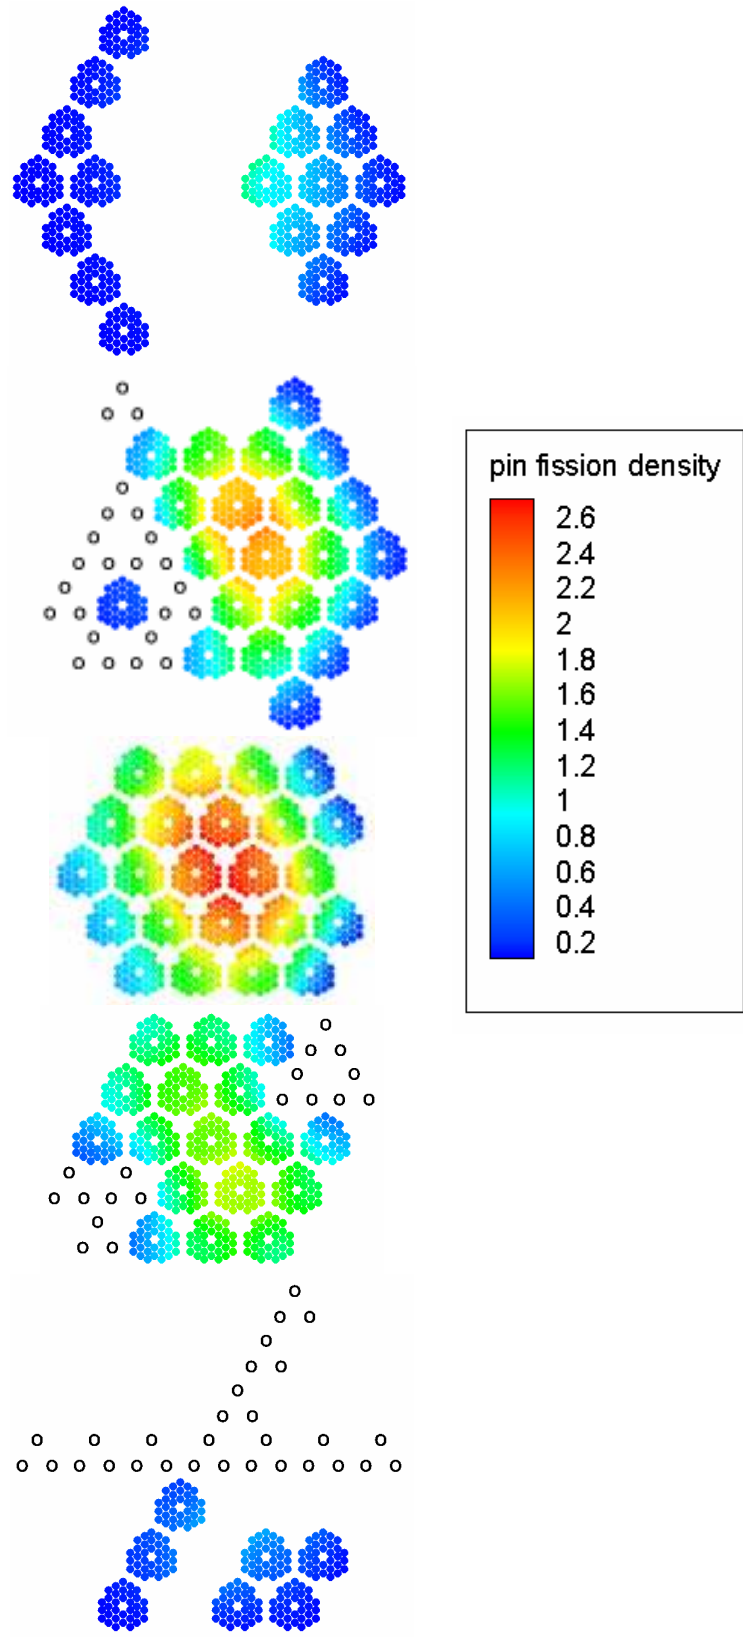


Figure 10. Core C pin fission density map

4.2. Response Expansion Coefficient Library

For the analysis presented herein, three new response expansion coefficient libraries were created. Two libraries were built for the full set of HTTR blocks, one expanding the angular flux, and one expanding the angular current. These libraries used 0.95, 1.05, and 1.15 as the discrete values of the core multiplication factor to model, and simulated 250 million incoming surface source neutrons per unique coarse mesh face and expansion order. Each of these libraries represents 22.7 GB in disk space. As these libraries were adequate to demonstrate the superiority of the angular flux expansion over that of the angular current in section 4.4, a third library was generated using an angular flux expansion only of those unique blocks which appear in the test cores at core multiplication factor values 0.8, 0.9, and 1.0. This final library also simulated 250 million incoming surface source neutrons, and occupies 7.0 GB of disk space. A twin library based on an angular current expansion was not considered necessary.

Earlier works (Connolly and Rahnema, 2012a, 2012c) had created libraries which had modeled only 10 million surface source particles per mesh face and expansion order; these used the same amount of disk space, but represented a much lesser time investment. At the time, the uncertainty in the eigenvalue calculation was determined using the procedure described by Forget et al. (2004), which underestimated the uncertainty, but the magnitude of this was not known. Zhang and Rahnema (2012) have since derived a more accurate means to determine the uncertainty, this has subsequently been implemented. Results using these earlier libraries are included in the appendix to

demonstrate the need for additional source sampling and therefore justify the corresponding time investment.

It is left to the user of the method to determine an optimum balance of time invested, the maximum expansion order investigated, and the inherent statistical uncertainty of the stochastically generated library. The results presented in the remainder of this chapter may suggest that the size of the library necessary for a desired high level of accuracy is significantly smaller than the sizes of those created for this study.

4.2.1. High order cross terms

Compilation of a library which included high order cross terms to the fourth order in both space and angle was attempted, however, as the size of the response expansion coefficient library is directly proportional to the square of the number of expansion orders of interest present in the problem, space limitations became an issue. Without high order cross terms, 225 response coefficients are determined for each face of each mesh at each energy group; when these terms are included, 625 coefficients are calculated. The synthesis of a response library with nearly 8 times as much data for no gain in accuracy is obviously an unnecessary exercise, especially considering the high accuracy of the solutions at fourth or even second order without these extra terms.

4.3. Mesh boundary analysis

The reference solutions used for comparative purposes have been determined by MCNP5. Relative differences between solutions generated by the new method and the reference solutions will be presented in tables. The relative error of any quantity of interest q will be calculated as

$$RE_q = \frac{q - q_{ref}}{q_{ref}}. \quad (13)$$

The relative error in the eigenvalue determination will be given with the stochastic uncertainty in per cent mille (pcm). As the method calculates the fission density for every pin segment in each core explicitly, several average error definitions will be presented to establish the accuracy of the method in a concise manner, thus eliminating the need to present thousands of data points for pin-level calculations. The tables will display the average absolute relative error of the pin fission density, which will be designated AVG and defined by (14), the mean relative error (MRE) given by (15), and the maximum absolute error (MAX). All pin fission density errors will be given in per cent. In Eq. (14) and (15), individual pins or pin segments are designated ζ , they have fission density values of fd_ζ , and the total number of pin segments in the core is given as Z .

$$AVG = \frac{\sum_{\zeta} |RE_{\zeta}|}{Z} \times 100\% \quad (14)$$

$$MRE = \frac{\sum_{\zeta} |RE_{\zeta}| \times fd_{\zeta}}{Z} \times 100\% \quad (15)$$

In order to determine the optimal mesh boundary condition, three cores—the partially controlled HTTR core, the uncontrolled HTTR core, and the smaller Core A—have been solved using both angular current and angular flux as quantities of expansion at the mesh boundaries. Here “tight” convergence (as defined in the following sub-section) has been utilized throughout in order to clearly see the effects of the physical quantity of expansion on the accuracy of the solution. Tables 1-3 present the results for these three cores, where the first line of the table is the order to which the expansion of the mesh interface angular current or flux is carried out: (L, M, N, P) represent the maximum expansion order in the two spatial variables, the polar angle, and the azimuthal angle, respectively. In these three tables, a column designates the physical quantity of expansion (angular current or angular flux) at the mesh interface; the rest of the results tables presented show results for which only the interface angular flux is expanded. The final line on the tables is for the computational runtime in minutes required for the calculations.

Table 1. HTTR, some rods in (SRI)

	(2,2,0,0)		(2,2,2,2)		(4,4,4,4)	
	J	ψ	J	ψ	J	ψ
RE k (pcm)	1306	879	138	73	94	80
σ_k (pcm)	63	63	63	63	63	63
AVG %	1.62	0.88	0.42	0.43	0.39	0.42
MRE %	1.59	0.85	0.39	0.40	0.34	0.36
MAX %	5.34	2.53	2.26	2.50	0.95	1.06
runtime (min)	24	24	229	245	924	793

Table 2. HTTR, all-rods-out (ARO)

	(2,2,0,0)		(2,2,2,2)		(4,4,4,4)	
	J	ψ	J	ψ	J	ψ
RE k (pcm)	1224	751	97	36	42	24
σ_k (pcm)	72	73	72	73	72	73
AVG %	4.06	2.55	0.76	0.67	0.59	0.56
MRE %	3.00	1.87	0.59	0.54	0.43	0.41
MAX %	11.30	6.47	2.44	2.21	1.78	1.66
runtime (min)	44	37	340	336	630	505

Table 3. Core A

	(2,2,0,0)		(2,2,2,2)		(4,4,4,4)	
	J	ψ	J	ψ	J	ψ
RE k (pcm)	2126	1316	50	-37	-18	-32
σ_k (pcm)	12	13	12	13	12	13
AVG %	2.52	0.85	0.48	0.49	0.23	0.23
MRE %	1.58	0.55	0.34	0.33	0.15	0.14
MAX %	10.33	4.37	3.81	4.16	3.28	3.28
runtime (min)	2	2	12	13	42	34

The results for Core A previously appeared in (Connolly and Rahnema, 2012b).

From these results, the expansion of the angular flux is found superior to the expansion of the angular current at mesh interfaces. Although by a fourth-level expansion, both quantities produce results with high accuracy, the lower order expansion of the flux yields results which are, for both the eigenvalue and fission density profile, significantly more accurate than results determined from the current expansion. As the initial guess for higher order calculations is synthesized from a low-order calculation, more accurate low order calculations lead directly to faster runtimes for higher order calculations. This alone is an adequate reason to select the angular flux expansion.

Since the zeroth expansion moment of the angular flux represents an isotropic flux distribution at the mesh boundary (Zhang and Rahnema 2012), the angular flux expansion has been expected to yield superior results to the angular current expansion of the same order when the core of concern exhibits flux with a strongly isotropic component. Based on that, the core environment of a strongly scattering thermal reactor would be expected conducive to the angular flux expansion. That at least a second order expansion in all variables is necessary for accurate results highlights the inadequacy of reliance solely on diffusion or low order methods to determine the core flux shape and eigenvalue of these heterogeneous reactors. The improved accuracy inherent in the angular flux expansion may also result in part from removing any numerical difficulties (near singularities) resulting from an expansion of the angular current at grazing angles.

4.4. Convergence and acceleration of solution method

It is clear that the time required to calculate the flux shape and eigenvalue is related to the convergence criteria of the iterative procedure and the expansion order specified.

Computational runtimes will be reduced with greater accuracy of the initial guesses for the core eigenvalue and flux distribution. In order to supply an initial guess which will enable faster runtimes, a low order acceleration method has been used in previous works; a two-step process using first a (0,0,0) order followed by a (2,2,2) order calculation at looser convergence criteria was found superior to others attempted for two-dimensional hexagonal problems (Connolly et al., 2012). This is adapted for the “tight” convergence criteria as described in Table 4. The “tight” convergence/acceleration option used for the 3-D method differs from the one used in 2-D calculations; due to the relatively poor accuracy at the (0,0,0,0) order, an eigenvalue calculation is not performed once an estimated flux shape is determined. The calculation then proceeds using the flux shape determined from the low order calculation and the initial eigenvalue guess. For calculations in which the final expansion order is equal to or lower than the intermediate (2,2,2,2) step, the expansion order of that step is simply the lesser of (2,2,2,2) and the final expansion of interest. In addition to the “tight” convergence and acceleration options, two new acceleration and convergence schemes have been proposed, with criteria given in Table 4.

Table 4. Convergence and acceleration options

Order	Flux convergence	Fission density	Eigenvalue
“Tight”			
(0,0,0,0)	1×10^{-4}	5×10^{-4}	n/a
(2,2,2,2)	1×10^{-4}	5×10^{-4}	10 pcm
Final	5×10^{-5}	1×10^{-4}	5 pcm
“Moderate”			
(2,2,0,0)	1×10^{-3}	5×10^{-3}	1.0 %
(2,2,2,2)	1×10^{-3}	5×10^{-3}	25 pcm
Final	2.5×10^{-4}	5×10^{-4}	10 pcm
“Loose”			
(2,2,0,0)	5×10^{-3}	5×10^{-3}	1.0 %
(2,2,2,2)	5×10^{-3}	5×10^{-3}	50 pcm
Final	5×10^{-4}	1×10^{-3}	25 pcm

Clearly, the HTTR core solutions have the most to gain in efficiency by accelerating the iterative solution method. Furthermore, the uncertainty in the eigenvalue calculations is far greater than the 5 pcm convergence criterion used in previous calculations, and the magnitude of the pin fission density discrepancies between the method solution and the reference solution suggest that the 0.01% figure used to establish convergence of the fission density profile in the initial “tight convergence” case is unnecessarily strict.

Tables 5 and 6 contain results from the three acceleration options; Table 5 for the partially controlled HTTR core, and Table 6 for the HTTR with all rods out. Each case uses (4,4,4,4) as the final expansion order. In all cases, the eigenvalue results are well within 2σ , and while the calculated pin fission density core maps do fluctuate, the average errors are within 1% of the reference solution in all cases. Clear reductions in the runtimes are obvious with loosening convergence criteria in all cases except that for the HTTR with all rods out, the “tight” convergence and the “moderate” convergence options solve the core problem in the same amount of time. The reason for this is unclear, although a well-converged intermediate acceleration step with high accuracy is

expected to result in a quicker final step, as long as the intermediate step yields a solution with accuracy comparable to the final step. As the (2,2,2,2) calculation for the HTTR all-rods-out configuration is adequate for a highly accurate solution (according to the results presented in Table 8 in the following sub-section), this may be the cause of the low runtime for the “tight” convergence case in the uncontrolled core configuration problem.

Table 5. Accelerated SRI HTTR core

	“Tight”	“Moderate”	“Loose”
RE k (pcm)	80	66	21
σ_k (pcm)	63	63	63
AVG %	0.42	0.17	0.34
MRE %	0.36	0.16	0.33
MAX %	1.06	0.61	1.28
runtime (min)	793	312	278

Table 6. Accelerated ARO HTTR core

	“Tight”	“Moderate”	“Loose”
RE k (pcm)	24	45	-18
σ_k (pcm)	73	73	73
AVG %	0.56	0.94	0.93
MRE %	0.41	0.67	0.69
MAX %	1.66	2.61	2.41
runtime (min)	505	505	340

4.5. Expansion order analysis

The effects of the expansion order are presented in this section, beginning with results for the HTTR core. For this core, the “loose” convergence and acceleration package is used, as it has been shown to best reduce runtimes while maintaining accuracy. As the (2,2,0,0) expansion was given in section 4.3, Tables 7 and 8 begin with the (2,2,2,2) order. It is anticipated that the higher the order at which the infinite sum of orthogonal functions is truncated, the greater will be the accuracy of the solution; theoretically, an infinite sum will produce the exact solution. The results from the HTTR suggest that a second order expansion in both space and angle is adequate for accurate results in eigenvalue and pin fission density distribution. For this large, well-reflected core, the only improvement seen by raising the expansion order beyond this is a small reduction in the maximum error of pins in the partially controlled core; the difference between average pin fission density figures from one expansion order to another in the uncontrolled HTTR core is on the order of the 0.1% fission density convergence criterion. All eigenvalue calculations were within one standard deviation of the reference solution. Runtimes were approximately 500 times faster than the MCNP reference solution, however, the fairness of the runtime comparison can be debated as the method has significantly higher uncertainty in the eigenvalue results, but at an average uncertainty of 0.01%, a far lower expected standard deviation in the pin fission density results.

Table 7. SRI HTTR

	(2,2,2,2)	(2,2,4,4)	(4,4,2,2)	(4,4,4,4)
RE k (pcm)	56	45	22	21
σ_k (pcm)	63	63	63	63
AVG %	0.34	0.37	0.34	0.34
MRE %	0.32	0.34	0.33	0.33
MAX %	1.98	1.90	1.28	1.28
runtime (min)	114	187	147	278

Table 8. ARO HTTR

	(2,2,2,2)	(2,2,4,4)	(4,4,2,2)	(4,4,4,4)
RE k (pcm)	-9	-6	-17	-18
σ_k (pcm)	72	73	73	73
AVG %	0.74	0.59	0.94	0.93
MRE %	0.55	0.46	0.69	0.69
MAX %	2.74	2.41	2.44	2.41
runtime (min)	113	212	170	340

For the partially controlled configuration, pin fission density errors were quite low. At the (2,2,2,2) order, only 74 pins of the 4950 had values differing more than 1% from the reference solution. Only 2 of those 74 pins were at power levels above the average in the core. This means that the method determined the fission density of 98.51% of the pins in the core to within 1% of the reference solution. While the method still performed admirably with the all-rods-out configuration, the steeper flux gradient created by the removal of the flux-shaping rods may have contributed to somewhat higher pin fission density errors when compared to the reference solution. At the (2,2,2,2) order, 97.43% of pins had errors below 2%, but 71.25% of the total were within 1%. None of the pins which deviated by more than 2% from the reference solution had power levels at or

above the core average; most of them were in regions where the fission density was below 0.5 relative to the core average.




Figure 11. SRI HTTR errors at (2,2,2)




Figure 12. SRI HTTR errors at (4,4,4,4)




Figure 13. ARO HTTR errors at (2,2,2)




Figure 14. ARO HTTR errors at (4,4,4)

Figures 11 and 12 depict the errors in the pin fission density calculations throughout the core for the partially controlled HTTR configuration. Figures 13 and 14 show the same, but for the all-rods-out configuration. These four figures, and all subsequent graphical

error map representations, use the color map given as Figure 21; this enables an easy comparison between different core configurations and expansion orders so that the positions within the core for which the method gives the strongest results are apparent. In both HTTR cases, expanding the angular flux to the fourth order yields a flatter error map than that produced by the second order expansion. In both cases, maximum pin fission density errors are reduced, but average errors are not improved. For the ARO (4,4,4,4) calculation, the lower two levels are underestimated by the method when compared to the reference solution, and the upper three levels are overestimated.

Results for cores A, B, and C are presented in Tables 9, 10, and 11, respectively. In solving these cores, the “moderate” convergence and acceleration package was selected, as the “loose” convergence would not be sufficient given the lower uncertainty inherent in the heterogeneous test core simulations. These results conclusively demonstrate the need for at least a second order expansion in angle for accurate eigenvalue results, although the pin fission density maps even assuming an isotropic flux at mesh boundaries is only slightly worse than higher angular order calculations in cores devoid of much control material. The simpler cores A and B exhibit high accuracy in the eigenvalue calculation for expansions of at least second order in both space and angle, with results within 3σ plus the 10 pcm convergence criterion of the reference multiplication factor. However, especially in the unreflected Core A with strong flux gradients, the step from second to fourth order in space reduces pin fission density errors. The spatial expansion of higher order more accurately resolves pin-level details in this core where flux is sharply peaked. Although creating a response expansion coefficient library of higher

order may be necessary to solve the flux map more exactly, the results herein are of such high accuracy that a larger library would only be an unnecessary use of computing time and resources for marginal gains.

The more challenging Core C suggests a need for modeling a fourth order expansion in the spatial distribution of the flux at mesh interface for accurate calculation of the eigenvalue; when this expansion is used, the core multiplication factor is found within 3σ plus the convergence criterion. The fourth order spatial expansion again further enhances the accurate resolution of the pin fission density map over a second order expansion. As was the case for the other cores, a fourth order expansion of the angular distribution of the flux at mesh faces produces no difference in the solution. That the method has produced solutions of such accuracy even in the dramatically unrealistic Core C is a true testament to its value as a fast and accurate transport method.

Table 9. Core A

	(2,2,0,0)	(2,2,2,2)	(2,2,4,4)	(4,4,0,0)	(4,4,2,2)	(4,4,4,4)
RE k (pcm)	1294	-43	-45	1303	-36	-39
σ_k (pcm)	13	13	13	13	13	13
AVG %	0.80	0.50	0.50	0.78	0.24	0.24
MRE %	0.52	0.34	0.34	0.42	0.15	0.15
MAX %	4.09	4.33	4.33	5.74	3.28	3.28
runtime (min)	2	8	11	2	13	25

Table 10. Core B

	(2,2,0,0)	(2,2,2,2)	(2,2,4,4)	(4,4,0,0)	(4,4,2,2)	(4,4,4,4)
RE k (pcm)	1506	-9	-8	1467	-29	-33
σ_k (pcm)	19	19	19	19	19	19
AVG %	0.76	0.72	0.72	0.71	0.70	0.70
MRE %	0.76	0.71	0.71	0.71	0.69	0.69
MAX %	1.94	1.88	1.88	1.52	1.42	1.42
runtime (min)	3	8	13	3	15	30

Table 11. Core C

	(2,2,0,0)	(2,2,2,2)	(2,2,4,4)	(4,4,0,0)	(4,4,2,2)	(4,4,4,4)
RE k (pcm)	1734	-99	-101	1776	-55	-59
σ_k (pcm)	19	19	19	19	19	19
AVG %	1.40	0.64	0.64	1.31	0.34	0.34
MRE %	0.69	0.41	0.41	0.55	0.21	0.21
MAX %	11.13	6.56	6.56	10.78	5.87	5.87
runtime (min)	3	12	17	3	21	42

In Core A, the maximum pin fission density errors were located in very low-power regions of the core. Again, using the (2,2,2,2) order, the method determines 89.07% of the 6105 pins in the core to within 1% of the reference solution. Of the 36 pins (or 0.59% of the pins) with errors over 3%, that with the highest fission density was operating at a power level of 0.19 times the core average. Two pins had errors over 4%; these had fission densities of 0.04 and 0.06. As seen in Figures 15 and 16, core calculations at the (4,4,4,4) order improve the pin fission density calculation accuracy, reducing the errors in the pins at the periphery while also reducing the errors in the center of the core.




Figure 15. Core A errors at (2,2,2)




Figure 16. Core A errors at (4,4,4,4)

At the (2,2,2,2) expansion order, the method determined 75.92% of the pins within Core B to within 1% of the reference solution. Unlike the other cores solved, this core does not have large flux gradients, and so the pin fission density results with errors over 1% are not distributed preferentially among low-power regions. However, since the average

error in pin fission density calculations is below 1% and the maximum error is below 2%, the method performance for eight minutes of runtime compared to the reference solution's runtime requirement of over four orders of magnitude greater is nothing short of outstanding. As shown in Figures 17 and 18, raising the expansion order from (2,2,2,2) to (4,4,4,4) flattens the error map throughout the core in a similar manner to that seen in the ARO HTTR configuration.




Figure 17. Core B errors at (2,2,2,2)




Figure 18. Core B errors at (4,4,4,4)

Core C, due to its unrealistic core configuration of fuel blocks interspersed with control material and reflector blocks specifically designed to cause a drastically irregular flux shape, is expected to challenge a new method beyond what is necessary to establish reliable accuracy in cores resembling operating reactors. At the (2,2,2,2) expansion order, 82.89% of the pins are still found to within 1% of the reference solution. There are 60 pins at greater than 3% error, with the highest-powered pin at 0.4 times the core average, but the rest at below 0.2. Five pins were found with errors over 6% from the reference solution, but all of these had fission density levels of 0.05. These results are improved further with the use of the (4,4,4,4) expansion order—specifically, the number of pins over 3% error is reduced from 60 to 9—but the (2,2,2,2) order has been chosen here to describe the method’s performance at the lowest expansion order and therefore

runtime adequate to accurately resolve both the core eigenvalue and the flux shape throughout. However, the results for both of these orders are illustrated in Figures 19 and 20; it is clear that the largest errors are seen at low-flux periphery regions which are removed from much of the central part of the core and not adjacent to other fuel pins.




Figure 19. Core C errors at (2,2,2)




Figure 20. Core C errors at (4,4,4,4)




Figure 21. Color map for all error figures

The method results presented in this section have been chosen to show the strengths and limitations of this new method. It has been shown that eigenvalues determined by the method have accuracy comparable to their inherent statistical uncertainty, and that pin fission density values are generally highly accurate as well; the limitations of the method in determining fission density figures are typically in those regions which contribute least to the overall power of the core.

4.6. Control rod worth analysis

As a final test, the method will be used to demonstrate its ability to determine the incremental reactivity worth of control rods. The core selected will be a modified version of Core B; the central column of this core will be changed into a single control column as

shown in Fig. 22. It is assumed that the three control rods may move independently of each other for the sake of this exercise. Thus, there will be 16 different configurations of the core modeled: the core with all rods out followed by the first rod being inserted incrementally. Once the first rod is all the way in, the second rod will be inserted one block level at a time (assuming each block level represents a single notch), and finally the third. At each core position, the core eigenvalue is determined and the pin fission density map is constructed. From the eigenvalue calculations, the incremental rod worth ρ_{iw} is determined via Eq. (16), as given by Lamarsh and Baratta (2001).

$$\rho_{iw} = \frac{\Delta k}{k}. \quad (16)$$




Figure 22. Rod worth test core: (a) levels 1 and 5; (b) levels 2-4.

This test of the method first appeared in a preliminary work (Connolly and Rahnema, 2012a), which expanded the angular current rather than the angular flux at the mesh boundary, and also estimated the uncertainty inadequately. These earlier tests had used

the response library which had sampled 10 million surface source particles per unique mesh face, and the uncertainty had been estimated at or around 6 pcm for each eigenvalue calculation. Using the improved method, the true uncertainty was found to be closer to approximately 60 pcm. Clearly, additional precision was necessary for these calculations; the library which used 250 million surface source particles was therefore compiled. The results using this library are presented in Table 12, with runtime and uncertainty figures. These calculations were performed at the (4,4,2,2) order and used the “moderate” convergence criteria.

At first glance, the calculated incremental worth of each rod segment appears as expected: as each rod is inserted, the change in reactivity behaves approximately sinusoidally, with the first rod having the greatest worth. As each additional rod is inserted, it is found to have lower worth, as the flux has already been depressed by the insertion of the first rod. However, COMET’s multiplication factor is lower by more than three standard deviations from the reference value in some cases clustered near $k_{\text{eff}} = 0.9$. Additional investigation is being undertaken at present to mitigate these errors.

Table 12. Eigenvalue results for rod worth exercise

Rod pos.	COMET k_{eff} (σ)	COMET runtime (min)	MCNP k_{eff} (σ)	MCNP runtime* (min)	COMET ρ_{iw} in pcm (σ)	k_{eff} error in pcm (σ)
0	0.95741 (0.00030)	17	0.95745 (0.00003)	12739		-4 (30)
1	0.95332 (0.00029)	17	0.95324 (0.00003)	12799	429 (42)	8 (29)
2	0.93659 (0.00026)	17	0.93677 (0.00003)	12643	1786 (39)	-18 (26)
3	0.91035 (0.00025)	16	0.91119 (0.00003)	12307	2882 (36)	-84 (25)
4	0.89395 (0.00025)	16	0.89549 (0.00003)	12137	1835 (35)	-154 (25)
5	0.89222 (0.00019)	16	0.89383 (0.00003)	12280	194 (31)	-161 (19)
6	0.89054 (0.00024)	17	0.89201 (0.00003)	12169	189 (31)	-147 (24)
7	0.88218 (0.00024)	19	0.88335 (0.00003)	11915	948 (34)	-117 (24)
8	0.86927 (0.00023)	19	0.87016 (0.00003)	11871	1485 (33)	-89 (23)
9	0.86173 (0.00025)	19	0.86245 (0.00003)	11875	875 (34)	-72 (25)
10	0.86099 (0.00018)	19	0.86179 (0.00003)	11770	86 (31)	-80 (18)
11	0.86011 (0.00024)	18	0.86079 (0.00003)	11782	102 (30)	-68 (24)
12	0.85515 (0.00023)	19	0.85564 (0.00003)	11655	580 (33)	-49 (23)
13	0.84748 (0.00022)	18	0.84780 (0.00003)	11746	905 (32)	-32 (22)
14	0.84314 (0.00022)	17	0.84345 (0.00003)	11582	515 (31)	-31 (22)
15	0.84273 (0.00017)	16	0.84308 (0.00003)	11565	49 (28)	-35 (17)

*MCNP runtimes are total cpu time for only eigenvalue calculation. COMET runtimes are for both eigenvalue and pin fission density map determination.

The runtimes stated in Table 12 for the MCNP reference solution were for an eigenvalue calculation without the construction of the pin fission density profile. However, in the listed time, COMET completed not only the eigenvalue calculation, but also a complete

pin fission density profile. These fission density maps are illustrated in Figures 24-39, with Figure 23 giving the color key for all of the following core depictions. The presence of control rods is indicated by a black sphere. The flux depression in pins near control material is clear; the utility of a method which is quickly able to determine the behavior of every fuel pin segment in a core in response to major or minor changes in the core environment is obvious.




Figure 23. Color key to pin fission density results illustrated in Figures 24-39.




Figure 24. Rod worth test core, rod position 0.




Figure 25. Rod worth test core, rod position 1.




Figure 26. Rod worth test core, rod position 2.




Figure 27. Rod worth test core, rod position 3.




Figure 28. Rod worth test core, rod position 4.




Figure 29. Rod worth test core, rod position 5.




Figure 30. Rod worth test core, rod position 6.




Figure 31. Rod worth test core, rod position 7.




Figure 32. Rod worth test core, rod position 8.




Figure 33. Rod worth test core, rod position 9.




Figure 34. Rod worth test core, rod position 10.




Figure 35. Rod worth test core, rod position 11.




Figure 36. Rod worth test core, rod position 12.




Figure 37. Rod worth test core, rod position 13.




Figure 38. Rod worth test core, rod position 14.




Figure 39. Rod worth test core, rod position 15.

5. CONCLUSIONS

A new transport method has been developed for solving whole reactor core eigenvalue and fission density distribution problems in three-dimensional hexagonal geometry. It has accuracy comparable to whole core Monte Carlo calculations, but requires mere thousandths of the computational runtime to produce the results. It is well suited for solving problems in new gas-cooled reactor core designs, as it is not subject to inaccuracies resulting from homogenization of regions within the core or the use of low-order approximations such as diffusion. The problems solved within this paper have illustrated that the method is robust with respect to the core geometry; the cores solved here have ranged from realistic core configurations in the style of an operating gas-cooled reactor to an unrealistic core with reflectors, fuel, and control material arranged seemingly at random throughout. In all cases, the pin fission densities were found on average to be within 1% of the reference solution, and eigenvalue solutions were within 0.1%.

The method's speed advantages over other whole-core transport methods are apparent. Its rapid determination of the explicit pin fission density profile of whole cores makes it extremely well-suited as a neutronics method to couple with thermal hydraulics calculations; this would enable a more thorough understanding of the neutronic and thermal fluid behavior of these reactors. Time-dependent calculations may also be foreseen, allowing future capability to handle transient analysis. Further work is needed to establish the feasibility of coupling the method to a burnup module for fuel cycle

calculations. Additional acceleration is possible in the near future by adapting it for parallel computer architecture. Due to the method's generality and ability to handle arbitrary material and geometry specifications within hexagonal lattice elements, it is expected that it will perform to equally high standards when tested with other reactor designs, such as fast breeder reactors. The method as currently presented balances a desire for the accuracy inherent in transport calculations with runtimes in minutes rather than days; as such, it represents a valuable contribution to advancing the performance and safety of a new generation of nuclear power plants.

APPENDIX A

These results for the three heterogeneous test cores were published using the eigenvalue uncertainty definition of Forget et al. (2004). A line has been added to the tables below [new $\sigma_{k\text{-eff}}$ (pcm)] to indicate the uncertainty as determined by the method of Zhang and Rahnema (2012).

All solutions used the same “tight” convergence criteria as defined in Section 4.4.

Table A.1 presents the COMET results for Core A. Analysis of the results which was presented at the time of their initial publication has not been given here. Results for Core B are given in Table A.2, and for Core C in Table A.3.

Table 13. Core A results

	2,2,2,2	4,4,2,2	2,2,4,4	4,4,4,4
k-eff (pcm)	10	16	47	38
old $\sigma_{k\text{-eff}}$ (pcm)	4	4	4	4
new $\sigma_{k\text{-eff}}$ (pcm)	32	33	33	35
AVG RE (%)	0.646	0.499	0.633	0.471
RMS RE (%)	0.839	0.643	0.844	0.609
MRE (%)	0.475	0.370	0.457	0.362
MAX RE (%)	4.342	3.672	4.561	3.672
AVG σ_{pin} (%)	0.108	0.111	0.111	0.116
Time (min)	34	46	46	69

Table 14. Core B results

	2,2,2,2	4,4,2,2	2,2,4,4	4,4,4,4
k-eff (pcm)	77	57	16	10
old $\sigma_{k\text{-eff}}$ (pcm)	6	6	6	6
new $\sigma_{k\text{-eff}}$ (pcm)	41	42	42	44
AVG RE (%)	0.778	0.731	0.768	0.727
RMS RE (%)	0.933	0.853	0.923	0.845
MRE (%)	0.765	0.722	0.753	0.717
MAX RE (%)	2.712	2.080	2.759	2.092
AVG σ_{pin} (%)	0.057	0.058	0.058	0.062
Time (min)	36	44	50	80

Table 15. Core C results

	2,2,2,2	4,4,2,2	2,2,4,4	4,4,4,4
k-eff (pcm)	42	0	106	67
old $\sigma_{k\text{-eff}}$ (pcm)	5	6	6	6
new $\sigma_{k\text{-eff}}$ (pcm)	45	46	47	49
AVG RE (%)	0.888	0.602	0.871	0.580
RMS RE (%)	1.177	0.781	1.181	0.761
MRE (%)	0.649	0.474	0.611	0.437
MAX RE (%)	5.928	5.142	6.193	5.142
AVG σ_{pin} (%)	0.102	0.106	0.105	0.112
Time (min)	50	71	69	107

In all three cases, improvements made to the method between the publication of these results and the writing of this dissertation led to improved accuracy of the results in a shorter runtime.

REFERENCES

- Blackburn, M. S., 2009. Numerical Benchmarking of a Coarse Mesh Transport (COMET) Method for Medical Physics Applications. Ph.D. Dissertation, Georgia Institute of Technology.
- Carrico, C. B., Lewis, E. E., Palmiott, G., 1992. Three-Dimensional Variational Nodal Transport Methods for Cartesian, Triangular, and Hexagonal Criticality Calculations. *Nuclear Science and Engineering*. 111, 168-79.
- Cho, J.-Y., Kim, K.-S., Lee, C.-C., Joo, H.-G., 2007. Whole Core Transport Calculation for the VHTR Hexagonal Core. *13th International Conference on Emerging Nuclear Energy Systems*, Istanbul, Turkey, June 3-8, 2007.
- Connolly, K. J., Rahnema, F., 2012. Calculation of Incremental Control Rod Reactivity Worth using COMET. *Transactions of the American Nuclear Society*. 106, 747-49.
- Connolly, K. J., Rahnema, F., 2012a. Response Expansion of Incident Angular Flux and Current for Transport Calculations. *Transactions of the American Nuclear Society*. 107 (accepted).
- Connolly, K. J., Rahnema, F., 2012b. The COMET Method in 3-D Hexagonal Geometry. *PHYSOR 2012*, Knoxville, TN, USA, April 15-20, 2012, American Nuclear Society.
- Connolly, K. J., Rahnema, F., Zhang, D., 2012. A Coarse Mesh Radiation Transport Method for 2-D Hexagonal Geometry. *Annals of Nuclear Energy*. 42, 1-10.
- Fitzpatrick, W. E., Ougouag, A. M., 1992. HEXPEDITE: A Net Current Multigroup Nodal Diffusion Method for Hexagonal-z Geometry. *Transactions of the American Nuclear Society*. 66, 268-70.
- Forget, B., 2006. A Three Dimensional Heterogeneous Coarse Mesh Transport Method for Reactor Calculations. Ph.D. Dissertation, Georgia Institute of Technology.
- Forget, B., Rahnema, F., Mosher, S. W., 2004. A Heterogeneous Coarse Mesh Solution for the 2-D NEA C5G7 MOX Benchmark Problem. *Progress in Nuclear Energy*. 45, 233-54.
- Gert, G., Jevremovic, T., 2008. Neutron Transport in Double Heterogeneous Media of High Temperature Reactors. *Transactions of the American Nuclear Society*. 99, 689-90.

- Gheorghiu, H. N. M., 1989. Group Theory Applied to a 3-D Hexagonal Response Matrix Method. *Annals of Nuclear Energy*. 16, 477-82.
- Ilas, D., Rahnema, F., 2003. A Heterogeneous Coarse Mesh Transport Method. *Transport Theory and Statistical Physics*. 32, 445-71.
- Lamarsh, J. R., Baratta, A. J., 2001. *Introduction to Nuclear Engineering, 3rd Edition*. Prentice Hall, Upper Saddle River, NJ.
- Lee, C., Yang, W. S., 2011. Validation Tests of Whole-Core Transport Code DeCART against HTTR and VHTRC Experiments. *Transactions of the American Nuclear Society*. 104, 754-6.
- Lee, C., Yang, W. S., Jung, Y. S., 2010. Improved Capabilities of Whole-Core Transport Code DeCART for Prismatic VHTR Analysis. *Transactions of the American Nuclear Society*. 102, 513-4.
- Mosher, S. W., Rahnema, F., 2003. Monte Carlo Adaptation of a Heterogeneous Coarse Mesh Transport Method. *Transactions of the American Nuclear Society*. 89, 310-2.
- Mosher, S. W., Rahnema, F., 2006. The Incident Flux Response Expansion Method for Heterogeneous Coarse Mesh Transport Problems. *Transport Theory and Statistical Physics*. 35, 55-86.
- Singh, K., Kumar, V., 1995. Reconstruction of Pin Power in Fuel Assemblies from Nodal Calculations in Hexagonal Geometry. *Annals of Nuclear Energy*. 22, 629-47.
- Stamm'ler, R. J. J., Abbate, M. J., 1983. *Methods of Steady-State Reactor Physics in Nuclear Design*. Academic Press, London.
- X-5 Monte Carlo Team, 2005. MCNP—A General Monte Carlo N-Particle Transport Code, Version 5. Los Alamos National Laboratory.
- Zhang, D., Rahnema, F., 2012. An Efficient Hybrid Stochastic/Deterministic Coarse Mesh Neutron Transport Method. *Annals of Nuclear Energy*. 41, 1-11.
- Zhang, Z., Rahnema, F., Zhang, D., Pounders, J. M., Ougouag, A., 2011. Simplified Two and Three Dimensional HTTR Benchmark Problems. *Annals of Nuclear Energy*. 38, 1172-85.
- Zimin, V. G., Baturin, D. M., 2002. Polynomial nodal method for solving neutron diffusion equations in hexagonal-z geometry. *Annals of Nuclear Energy*. 29, 1105-17.

VITA

Kevin John Connolly was born in 1986 to Kim and John Connolly. He grew up in Framingham, Massachusetts with his younger sister Joanne, with whom he shared a love of reading and, later, the privilege to drive a cherished 1993 Caravan. As his parents had, he attended Framingham Public Schools. His long strides and tendency to run in sunglasses during cross country races earned him a letter, and his political commentary for the high school television station earned him an Emmy Award. He has hiked hundreds of miles in New Mexico's Rocky Mountains and New Hampshire's White Mountains, and has paddled over a hundred miles through the lakes of Minnesota and Ontario in a canoe, encountering friendly loons, bald eagles, moose, and the rock paintings of civilizations long past. An Eagle Scout, his enjoyment of clean air, clear water, and physics influenced his decision to study nuclear engineering.

After enduring nineteen New England winters, Connolly was admitted to the Georgia Institute of Technology, where he graduated first in the George W. Woodruff School in 2008. He continued at Tech for his graduate studies, earning a Master of Science in Nuclear Engineering in 2011. While a student, he maintained such extravagant hobbies as playing the harmonica and dabbling in amateur astronomy under light-polluted skies, and indulged in such excesses as having several shelves full of engineering reference books. He has presented his work at several professional conferences, and laments the scarcity of impromptu discussions of Homer's *Odyssey* in between engineering research sessions.

Connolly lives in Atlanta with his wife, Denise. The couple enjoys the city's many conveniences, such as easy access to many jammed roadways, an omnipresent cacophonous symphony of sirens, and night skies with a lurid pink hue. The Connollys hope to live the American Dream of raising a family, owning a house with a porch swing, and parking a Plymouth or two in the garage.