ITERATIVE HOMOGENIZATION METHOD FOR IMPROVING COMPUTATIONAL EFFICIENCY IN SOLVING EIGENVALUE PROBLEMS IN NEUTRON TRANSPORT

A Dissertation
Presented to
The Academic Faculty

by

Gabriel Kooreman

In Partial Fulfillment
of the Requirements for the Degree
Nuclear Engineering in the
George W. Woodruff School of Mechanical Engineering

Georgia Institute of Technology
December 2016

ITERATIVE HOMOGENIZATION METHOD FOR IMPROVING COMPUTATIONAL EFFICIENCY IN SOLVING EIGENVALUE PROBLEMS IN NEUTRON TRANSPORT

Approved by:

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

Dr. Tom Morley
School of Mathematics
Georgia Institute of Technology

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

Dr. Jeffery Densmore
Reactor Physics Methods Development
Naval Nuclear Laboratory

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

Date Approved: October 17, 2016

To Peierls.

ACKNOWLEDGEMENTS

To say that this thesis has been a learning experience would be a massive understatement. I have truly marvelous gratitude to all of the people who have helped me on this journey, which this acknowledgements section is too narrow to contain. I will instead do my best to acknowledge some of the most exceptional people who have supported me, inspired me, and helped me to complete this work.

First, I would like to thank my advisor, Dr. Farzad Rahnema, for his support, guidance, and encouragement during my tenure as a graduate student. Further, I would like to thank him for his ability to keep me on task and to direct a useful and meaningful timeline for my research and helpful criticism of the work and results. I would like to further extend this gratitude to my committee members, Dr. Bojan Petrovic, Dr. Dingkang Zhang, Dr. Tom Morley, and Dr. Jeffery Densmore, for their participation in the completion and defense of this dissertation. Without academic guidance of this manner, I could never have completed such a work.

Second, I would like to thank my colleagues in the CRMP lab; you were the best research companions I could ever have asked for. While most of you have graduated by now, I would be remiss not to mention you by name. Thank you Ryan Hon, Stefano Terlizzi, Drew Johnson, Kyle Remley, Daniel Lago,

Alex Huning, and Chaps Chapman. Without you, this could never have happened so late.

Thirdly, I would like to thank my support structure at home, my parents, Alan Kooreman and Dawn DiMuro, my 'partner in crime' Puvithel Rajan, and anyone else who helped me through hard times.

And finally, this work was performed with the support of the DOE Office of Naval Reactors Rickover Fellowship program. Any opinions, findings or conclusions expressed in this publication are those of the authors and do not necessarily reflect the views of the DOE Office of Naval Reactors.

Thank you to everyone mentioned in this section and to anyone relevant who has not been mentioned. Believe me when I say that I am extremely grateful.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS iv

LIST OF TABLES viii

LIST OF FIGURES x

LIST OF SYMBOLS AND ABBREVIATIONS xiii

SUMMARY xvi

CHAPTER 1. Introduction 1
1.1 Consistent Spatial Homogenization 2
1.2 Motivation and Goal 3

CHAPTER 2. Background 5
2.1 Neutron Transport for Eigenvalue Problems 5
2.2 Homogenization Methods 8
2.3 Current Technologies 9
2.3.1 Other homogenization methods 10
2.3.2 CSH method 11
2.3.3 DTH method 14

CHAPTER 3. Consistent Spatial homogenization 16
3.1 The CSH Method 16
3.1.1 Theory 16
3.1.2 Solution procedure 21
3.2 The DTH Method 24
3.2.1 DTH Theory 25
3.2.2 Solution method 27
3.3 Implementation in 2-D 29
3.3.1 Selection of spatial basis functions for auxiliary source 30
3.3.2 2-D discretization details 36
3.3.3 Progressively tightened homogeneous mesh and homogeneous convergence 37

CHAPTER 4. BENCHMARK PROBLEMS 42
4.1 2-D C5G7 42
4.1.1 Benchmark geometry and solution 43
4.2 Full scale 2-D BWR benchmark problem 48
4.2.1 Cartesian mesh 2-D BWR benchmark description 49
4.2.2 Reference solution for 2-D BWR benchmark 53
4.3 2-D S_N Transport Package 54
4.4 Figures of Merit 56

CHAPTER 5. CSH results 60
5.1 Choice of Re-homogenization Convergence Criteria 61
5.2 Parameter Variations 65
5.2.1 Spatial basis functions 66
5.2.2 Re-homogenization process 68
5.2.3 Homogeneous problem meshing 77
5.3 CSH with Progressively Tightened Convergence and Meshing 80

CHAPTER 6. DTH results 89
6.1 DTH with Progressively Tightened Convergence and Meshing 90
6.2 Other Parameter Variations 96
6.2.1 Spatial basis functions 97
6.2.2 Re-homogenization method 98
6.2.3 Homogeneous problem meshing 101

CHAPTER 7. Discussion and Conclusions 103
7.1 Conclusions 103
7.2 Discussion 104
7.3 Future Work 106

APPENDIX A. Details of 2-D BWR Discretization generation 108
A.1 Pin Cell Meshing Utility 108
A.2 Other Simplifications 110

APPENDIX B. Choice of reference Convergence Criteria 112
B.1 Definition of Convergence Criteria 112
B.2 Choice of Convergence Values 114

REFERENCES 117

LIST OF TABLES

TABLE 1. PROPOSED VOLUME-CONSERVING C5G7 PIN CELL DISCRETIZATION MESH BOUNDARIES IN CENTIMETERS 45

TABLE 2. VERIFICATION OF 2-D SN CODE AND STUDY INTO REASONABLE MESHING FOR CARTESIAN MESH C5G7. 46

TABLE 3. OTHER CARTESIAN MESH C5G7 SOLUTION ERRORS, FROM [16] 47

TABLE 4. CARTESIAN MESH BOUNDARIES FOR 2-D BWR PIN CELL IN CENTIMETERS, AS DISTANCE FROM THE CENTER OF THE PIN. 51

TABLE 5. EFFECT OF SPATIAL BASIS FUNCTION CHOICE ON CSH SOLUTION TO C5G7 BENCHMARK. CALCULATION PERFORMED WITH FOUR CORE SWEEPS PER RE-HOMOGENIZATION AND CLASSIC HOMOGENEOUS MESH. 66

TABLE 6. EFFECT OF SPATIAL BASIS FUNCTION CHOICE ON CSH SOLUTION TO BWR BENCHMARK. CALCULATION PERFORMED WITH FOUR CORE SWEEPS PER RE-HOMOGENIZATION AND CLASSIC HOMOGENEOUS MESH. 67

TABLE 7. EFFECT OF DIFFERENT RE-HOMOGENIZATION METHODS FOR THE C5G7 CORE. AFS REFERS TO ASSEMBLY FIXED-SOURCE RE-HOMOGENIZATION. 69

TABLE 8. EFFECT OF DIFFERENT RE-HOMOGENIZATION METHODS FOR THE BWR CORE. AFS REFERS TO ASSEMBLY FIXED-SOURCE RE-HOMOGENIZATION. 72

TABLE 9. EFFECT OF HOMOGENEOUS PROBLEM MESH SIZE FOR CSH ON THE C5G7 BENCHMARK SOLUTION. 77

TABLE 10. EFFECT OF HOMOGENEOUS PROBLEM MESH SIZE FOR CSH ON THE BWR BENCHMARK SOLUTION. 79

TABLE 11. RESULTS OF THREE SEPARATE PROGRESSIVE MESHING SCHEMES APPLIED TO THE CSH METHOD FOR THE C5G7 BENCHMARK. 81
TABLE 12. RESULTS OF THREE SEPARATE PROGRESSIVE MESHING SCHEMES APPLIED TO THE CSH METHOD FOR THE BWR BENCHMARK. 85
TABLE 13. RESULTS OF THREE SEPARATE PROGRESSIVE MESHING SCHEMES APPLIED TO THE DTH METHOD FOR THE C5G7 BENCHMARK. 91
TABLE 14. RESULTS OF THREE SEPARATE PROGRESSIVE MESHING SCHEMES APPLIED TO THE DTH METHOD FOR THE BWR BENCHMARK. 91
TABLE 15. THE EFFECT OF DIFFERENT SPATIAL BASIS FUNCTIONS ON DTH CALCULATIONS OF THE C5G7 BENCHMARK. 97
TABLE 16. THE EFFECT OF DIFFERENT SPATIAL BASIS FUNCTIONS ON DTH CALCULATIONS OF THE BWR BENCHMARK. 97
TABLE 17. EFFECT OF DIFFERENT RE-HOMOGENIZATION METHODS ON THE DTH METHOD FOR THE C5G7 CORE. 99
TABLE 18. EFFECT OF DIFFERENT RE-HOMOGENIZATION METHODS ON THE DTH METHOD FOR THE BWR CORE. 99
TABLE 19. EFFECT OF HOMOGENEOUS PROBLEM MESH SIZE ON DTH METHOD CALCULATION OF THE C5G7 BENCHMARK. 101
TABLE 20. EFFECT OF HOMOGENEOUS PROBLEM MESH SIZE ON DTH METHOD CALCULATION OF THE BWR BENCHMARK 101

LIST OF FIGURES

FIGURE 1. A 1-D DEPICTION OF HOW B-SPLINE BASIS FUNCTIONS DO NOT CONSERVE INTEGRALS. SHADED TRIANGLES INDICATE AREAS WHERE THE INTEGRAL ESTIMATE WOULD BE EITHER TOO HIGH OR TOO LOW. 34

FIGURE 2. 1-D GRAPHICAL DEPICTIONS OF 0TH ORDER (LEFT) AND 1ST ORDER (RIGHT) INTEGRAL-CONSERVING BASIS FUNCTIONS 36

FIGURE 3. PROPOSED VOLUME-CONSERVING C5G7 PIN CELL DISCRETIZATIONS. FROM LEFT TO RIGHT, 4X4, 9X9, AND 18X18 MESHES. EXACT GEOMETRY IS DEPICTED AS A RED CIRCLE. 44

FIGURE 4. REFERENCE SCALAR FLUX SOLUTION TO C5G7 PROBLEM. LEFT: FAST GROUPS, RIGHT: THERMAL GROUPS. FLUX PRESENTED IN ARBITRARY UNITS WITH MAXIMUM 1. 48

FIGURE 5. 10X10 VOLUME-CONSERVING CARTESIAN MESH OF CLADDED 2-D BWR PIN CELL. RED LINES INDICATE EXACT GEOMETRY. THE COLORS BLUE, GRAY, AND YELLOW INDICATE MODERATOR, CLADDING, AND FUEL, RESPECTIVELY. 51

FIGURE 6. 2-D BWR CARTESIAN DISCRETIZED ASSEMBLY GEOMETRY. EACH COLOR IS A DIFFERENT MATERIAL. NOTE THE SMEARED CONTROL MATERIAL (IN BROWN) AND THE ABSENT CENTRAL WATER CHANNEL CLADDING. 52

FIGURE 7. QUARTER-CORE CORE LAYOUT OF 2-D BWR CORE WITH CARTESIAN MESH. EACH COLOR REPRESENTS A DIFFERENT MATERIAL. CONTROL BLADES ARE INSERTED BETWEEN ASSEMBLIES WITH GREEN AND BLACK MODERATOR 53

FIGURE 8. REFERENCE SCALAR FLUX SOLUTION TO 2-D BWR PROBLEM. LEFT: FAST GROUP, RIGHT: THERMAL GROUP. FLUX PRESENTED IN ARBITRARY UNITS WITH MAXIMUM 1.....54

FIGURE 9. RE-HOMOGENIZATION CONVERGENCE FOR THE C5G7 CORE WITH S_2. RED LINES INDICATE ERROR COMPARED TO THE REFERENCE AT EACH STEP, AND BLACK LINES INDICATE THE CURRENT VALUE OF THE WITHIN-CALCULATION CRITERIA...63

FIGURE 10. RE-HOMOGENIZATION CONVERGENCE FOR THE BWR CORE WITH S_2. RED LINES INDICATE ERROR COMPARED TO THE REFERENCE AT EACH STEP, AND BLACK LINES INDICATE THE CURRENT VALUE OF THE WITHIN-CALCULATION CRITERIA...64

FIGURE 11.WITHIN-CORE FAST AND THERMAL RELATIVE ERROR PROFILES FOR THE C5G7 CORE WITH AFS RE-HOMOGENIZATION.....70

FIGURE 12. WITHIN-CORE FAST AND THERMAL RELATIVE ERROR PROFILES FOR THE C5G7 CORE WITH ONE CORE SWEEP RE-HOMOGENIZATION 71

FIGURE 13. WITHIN-CORE FAST AND THERMAL RELATIVE ERROR PROFILES FOR THE C5G7 CORE WITH 20 CORE SWEEP RE-HOMOGENIZATION 71

FIGURE 14. WITHIN-CORE THERMAL RELATIVE THERMAL ERROR PROFILE FOR THE BWR CORE WITH ASSEMBLY FIXED-SOURCE RE-HOMOGENIZATION 74

FIGURE 15. WITHIN-CORE THERMAL RELATIVE THERMAL ERROR PROFILE FOR THE BWR CORE WITH ONE CORE SWEEP RE-HOMOGENIZATION 75

FIGURE 16. WITHIN-CORE THERMAL RELATIVE THERMAL ERROR PROFILE FOR THE BWR CORE WITH 20 CORE SWEEP RE-HOMOGENIZATION 76

FIGURE 17. FAST SPECTRUM RELATIVE FLUX ERROR PROFILE FOR C5G7 WHEN SOLVED WITH A PROGRESSIVE MESH ENDING WITH AN EXACT HETEROGENEOUS MESH OVERLAY.....84

FIGURE 18. THERMAL SPECTRUM RELATIVE FLUX ERROR PROFILE FOR C5G7 WHEN SOLVED WITH A PROGRESSIVE MESH ENDING WITH AN EXACT HETEROGENEOUS MESH OVERLAY. 84
FIGURE 19. FAST SPECTRUM RELATIVE FLUX ERROR PROFILE FOR BWR BENCHMARK WHEN SOLVED WITH A PROGRESSIVE MESH ENDING WITH AN EXACT HETEROGENEOUS MESH OVERLAY. 86
FIGURE 20. THERMAL SPECTRUM RELATIVE FLUX ERROR PROFILE FOR BWR BENCHMARK WHEN SOLVED WITH A PROGRESSIVE MESH ENDING WITH AN EXACT HETEROGENEOUS MESH OVERLAY. 87
FIGURE 21. FAST SPECTRUM RELATIVE ERROR PROFILE FOR THE C5G7 CORE SOLVED VIA DTH WITH PROGRESSIVE MESH TO THE 'CLASSIC' MESH SIZE. 93
FIGURE 22. THERMAL SPECTRUM RELATIVE ERROR PROFILE FOR THE C5G7 CORE SOLVED VIA DTH WITH PROGRESSIVE MESH TO THE 'CLASSIC' MESH SIZE. 94
FIGURE 23. FAST SPECTRUM RELATIVE ERROR PROFILE FOR THE BWR CORE SOLVED VIA DTH WITH PROGRESSIVE MESH TO THE 'CLASSIC' MESH SIZE. 95
FIGURE 24. FAST SPECTRUM RELATIVE ERROR PROFILE FOR THE BWR CORE SOLVED VIA DTH WITH PROGRESSIVE MESH TO THE 'CLASSIC' MESH SIZE. 96
FIGURE 25. RESULT OF CALCULATION OF REQUIRED EIGENVALUE CONVERGENCE CRITERION VALUE 114
FIGURE 26. RESULT OF CALCULATION OF REQUIRED FLUX CONVERGENCE CRITERION VALUE 115

LIST OF SYMBOLS AND ABBREVIATIONS

CSH	CONSISTENT SPATIAL HOMOGENIZATION
1-D	ONE DIMENSIONAL
DTH	DIFFUSION-TRANSPORT HOMOGENIZATION
2-D	TWO DIMENSIONAL
C5G7	7-GROUP REACTOR BENCHMARK
BWR	BOILING WATER REACTOR
3-D	THREE DIMENSIONAL
$\hat{\Omega}$	UNIT VECTOR ON THE UNIT SPHERE
ψ	NEUTRON ANGULAR FLUX
$\vec{r}$	VECTOR IN 2-D SPACE
E	ENERGY
σ	TOTAL CROSS SECTION
σ_s	SCATTERING CROSS SECTION
χ	FISSION SPECTRUM
σ_f	FISSION CROSS SECTION

ECSH	Efficient Consistent Spatial Homogenization
PWR	Pressurized Water Reactor
MOX	Mixed-oxide fuel
V^h	Homogenized volume
μ_0	Scattering angle
S_{AUX}	Auxiliary source
ϕ	Neutron scalar flux
k	Reactor eigenvalue
$\vec{J}$	Neutron current
MCNP	Monte Carlo N-Particle
AFS	Assembly Fixed-Source
PCM	Per cent mille
AVG	Average flux error metric
MRE	Mean relative flux error metric
MAX	Maximum flux error metric
SPEEDUP	Factor of computational performance increase

MFP MEAN FREE PATHS

SUMMARY

The neutron transport equation often is homogenized in order to simplify its solution procedure in some manner or another. There exist many methods for homogenizing the neutron transport equation with different benefits and detriments. One promising method is the Consistent Spatial Homogenization (CSH) method developed and implemented in 1-D by Yasseri and Rahnema. The method, along with its successor, the Diffusion-Transport Homogenization (DTH) method are promising for their ability to reconstruct accurate fine-mesh angular flux profiles as well as reactor eigenvalue after a re-homogenization procedure. This work will explore the extension of both the CSH and DTH methods to higher spatial dimensionality in order to solve large-scale reactor eigenvalue problems.

The CSH and DTH methods are based around iterated re-homogenization of the neutron transport equation with an auxiliary source term which is used to correct for heterogeneity effects of a given problem. The net effect of this is that the effects of heterogeneity are relegated to a source term, and the homogenized neutron transport equation is solved instead of the heterogeneous equation. This allows for implementation of simpler acceleration techniques to improve the speed and accuracy of the homogenized problem and in multiple dimensions helps to avoid the effects of complicated reactor geometries. The re-homogenization procedure brings the flux solution

back to the heterogeneous discretization in order to generate better approximations for the homogenized cross sections, a better approximation of the auxiliary source term, and most importantly to reconstruct the full heterogeneous angular flux profile.

In this work, the CSH and DTH methods are modified for increased spatial dimensionality and implemented using a 2-D S_N discrete ordinates transport solver. This implementation is tested using Cartesian-mesh variants of the 2D-C5G7 benchmark problem and a 2-D full-scale boiling water reactor (BWR) benchmark problem.

CHAPTER 1. INTRODUCTION

The ability to obtain efficient full-core solutions to the neutron transport equation has been a long sought-after goal in the field of reactor physics. As computers and computing power continues to improve, full-core deterministic calculations, which have long been unfeasible due to the sheer number of unknowns involved in neutron transport calculations, are becoming more and more commonplace. However, even with modern computing power, there is always a need for more advanced computational techniques which can improve the efficiency of solution methods for the neutron transport equation, as full-core 3-D deterministic calculations can tax even the largest computing clusters.

One method which has historically been used for reactor physics calculations is cross section homogenization. Through homogenization, complicated problems can be drastically simplified by reducing the need to model complex reactor geometry and to track large numbers of unique materials within a core. Historically, homogenization methods were among the only reasonable tools for whole-core deterministic transport calculations, and with the use of nodal diffusion methods, they are still used heavily today. These historical homogenization methods are not always appropriate computational

tools, however, as nodal methods are not good tools for constructing highly accurate neutron flux profiles within a core.

1.1 Consistent Spatial Homogenization

Consistent Spatial Homogenization (CSH) [1] is a theory which uses cross section homogenization together with an advanced correction term known as the auxiliary source term to leverage the computational benefits of homogenization while still solving the neutron transport equation for a highly accurate flux profile within a whole core. The CSH method works by iteratively homogenizing with the use of an auxiliary source term. The use of an auxiliary source term allows the CSH method to fold the heterogeneous information of a given problem into the source term of a homogeneous problem. Through the use of on-the-fly re-homogenizing of the cross sections and the auxiliary source term, the CSH method can reconstruct heterogeneous angular flux profile of any reactor core with modest computational speedup and minimal loss of accuracy.

A sister method to the CSH method is known as the Diffusion-Transport Homogenization (DTH) method [2]. The DTH method replaces the homogeneous transport solution of the CSH method with a diffusion calculation, but keeps the fully transport-based re-homogenization stage. Through this process, the DTH method can reconstruct angular flux profiles of

a given reactor core with only slightly less accuracy than the CSH method, but with significant computational benefit.

One benefit of CSH theory is that the calculations require only small modifications to the homogeneous transport equation, in the form of the auxiliary source term. This effect means that another acceleration technique can also be applied alongside either method, and the speedup afforded by the CSH and DTH methods will not be wasted. Unfortunately, until now the CSH and DTH methods have been limited in their implementations to one spatial dimension, so the potential speedup gains they can afford have merely been proof of concept. Additionally, while all implementations of the CSH and DTH methods have employed an iterative re-homogenization procedure, the implementations of the re-homogenization procedure have been disparate, [1] [2] [3] [4] with no clear consensus of the best procedure moving forward.

1.2 Motivation and Goal

This thesis is motivated by the success of the CSH and DTH methods in solving for the heterogeneous flux profile and eigenvalue with significant computational speed improvements, and also by the lack any implementation for the methods in higher spatial dimensionality. Higher dimensionality implementations are crucial for a homogenization method, as one dimensional

problems can never fully encapsulate the complex reactor geometries that exist in higher dimensions.

The goal of this thesis is to provide a basis for implementation of the CSH and DTH methods in 2-D, and to test these methods for difficult full-core reactor benchmark problems. In order to achieve this goal, the CSH and DTH methods will need to be improved and the re-homogenization procedures will need to be clarified and consistent. Successful implementation of the methods in 2-D will see even further improved computational efficiency over the 1-D methods, without further loss of accuracy.

In order to properly investigate the computational performance of the methods, it is important to have a fair basis for comparison for reference calculations. As such, an incidental goal of this thesis is to develop a code base in which future homogenization development work can be performed. This should open new paths for future development both of CSH theory and also of other future methods.

CHAPTER 2. BACKGROUND

In this chapter, the neutron transport equation (or Boltzmann equation), the basics of homogenization, and some of the current technologies used to perform homogenization of the neutron transport equation will be discussed. The purpose of this chapter is to frame the later discussion of the Consistent Spatial Homogenization (CSH) method and its sister method the Diffusion-Transport Homogenization (DTH) method so that neutron transport and homogenization methods in general can be contextualized. This chapter will contain broad-strokes definitions and descriptions of the neutron transport equation itself, the notation used in neutron transport theory, the purpose of homogenization methods themselves, and some current technologies in homogenization of the neutron transport equation.

2.1 Neutron Transport for Eigenvalue Problems

The eigenvalue neutron transport equation is

$$\begin{aligned}\widehat{\Omega}\cdot\nabla\psi(\vec{r},E,\widehat{\Omega})+\ \sigma(\vec{r},E)\psi(\vec{r},E,\widehat{\Omega})\\ =\int_{4\pi}\mathrm{d}\widehat{\Omega}'\int\mathrm{d}E'\sigma_s(\vec{r},E'\to E,\widehat{\Omega}'\to\widehat{\Omega})\psi(\vec{r},E',\widehat{\Omega}')\\ +\frac{\chi(E)}{4\pi k}\int_{4\pi}\mathrm{d}\widehat{\Omega}'\int\mathrm{d}E'\nu\sigma_f(\vec{r},E')\psi(\vec{r},E',\widehat{\Omega}')\,.\end{aligned}\tag{1}$$

In this equation, the principle unknown is the angular flux $\psi(\vec{r}, E, \hat{\Omega})$, which is a function of three phase space variables, location $\vec{r}$, energy E, and angle $\hat{\Omega}$. Angular flux is a differential quantity which has units of neutrons per area per unit time; $\psi(\vec{r}, E, \hat{\Omega})$ is the number of neutrons per second passing through a unit area within some differential volume $d\vec{r}$ centered around $\vec{r}$ with an energy in some differential energy bin dE centered around E, traveling in a direction which is in some differential solid angle $d\hat{\Omega}$ centered around angle $\hat{\Omega}$, which is a vector on the unit sphere. The other unknown of Equation (1) is k, the reactor eigenvalue, the largest such value which leads to a nontrivial solution for angular flux. Physically, k represents the neutron multiplication in such a reactor for each generation of neutrons. In other words, if a reactor has an eigenvalue of $k = 1.1$, one could expect each 'generation' of neutrons to be 1.1 times larger than the previous 'generation'. Operational nuclear reactors generally have an eigenvalue of 1, meaning that the neutron population in the reactor is stable over time.

In Equation (1), $\sigma(\vec{r}, E)$ is the total cross section, and the term $\sigma(\vec{r}, E)\psi(\vec{r}, E, \hat{\Omega})$ represents the total rate at which any sort of neutron reaction occurs as a function of space, angle, and energy. The term $\sigma_s(\vec{r}, E' \to E, \hat{\Omega}' \to \hat{\Omega})$ is the scattering kernel, and it represents the probability of a neutron scattering from some energy E' to some other energy E and from some direction $\hat{\Omega}'$ to some other direction $\hat{\Omega}$. The term $\sigma_f(\vec{r}, E')$ is the fission scattering, which

represents the probability of a fission reaction happening involving a neutron of energy E', and ν is the average number of neutrons emitted per fission reaction. Finally $\chi(E)$ represents the fission spectrum, which is the energy distribution of neutrons that are born from a fission reaction.

For a reactor, the neutron transport equation is generally solved using a set of boundary conditions stating that no neutrons enter the reactor from outside it,

$$\psi\left(\vec{r}_{\partial V}, E, \hat{\Omega}\right) = 0, \quad \hat{n} \cdot \hat{\Omega} < 0, \tag{2}$$

where $\vec{r}_{\partial V}$ is the reactor boundary, and $\hat{n}$ is the outward unit normal on the reactor boundary. In some cases, in order to exploit reactor symmetry, calculations may employ specular reflective boundary conditions, which state that any neutrons which leave the reactor boundary are immediately reflected back in,

$$\psi\left(\vec{r}_{\partial V}, E, \hat{\Omega}\right) = \psi(\vec{r}_{\partial V}, E, -\hat{\Omega}) \quad \hat{n} \cdot \hat{\Omega} < 0. \tag{3}$$

Except for extremely simplified cases, there are no known analytic solutions to the neutron transport equation. It must always be solved numerically. The neutron transport equation is a very difficult problem to solve computationally, as the angular flux has six independent variables (when the

vectors $\vec{r}$ and $\widehat{\Omega}$ are unpacked) and the energy variable may easily span many orders of magnitude, with wildly varying behavior for each energy value.

2.2 Homogenization Methods

Spatial homogenization of the neutron transport equation means constructing an approximate version of Equation (1) where the spatial dependence of the three cross section terms has been removed in some manner, as in

$$\begin{aligned}\widehat{\Omega}\cdot\nabla\psi\left(\vec{r},E,\widehat{\Omega}\right)+\ \sigma^{h}(E)\psi\left(\vec{r},E,\widehat{\Omega}\right)\\ =\int_{4\pi}\mathrm{d}\widehat{\Omega}'\int\mathrm{d}E'\sigma_{s}^{h}\left(E'\to E,\widehat{\Omega}'\to\widehat{\Omega}\right)\psi\left(\vec{r},E',\widehat{\Omega}'\right)\\ +\frac{\chi(E)}{4\pi k}\int_{4\pi}\mathrm{d}\widehat{\Omega}'\int\mathrm{d}E'\nu\sigma_{f}^{h}(E')\psi\left(\vec{r},E',\widehat{\Omega}'\right).\end{aligned}\tag{4}$$

In Equation (4), the superscript h indicates that homogenized values have been used. In most cases, the entire reactor core is not homogenized as one process, instead only some spatial variation is removed. Some examples of this are pin-cell homogenization, where each pin of a reactor core is treated as a single material, or assembly homogenization, where each assembly of a reactor core is treated as its own material. Some advantages of homogenization is that it can eliminate the need for complicated geometry specification, effectively enabling more efficient transport solution methods to be employed in a reactor

calculation. A simplistic homogenization example is to volume-weight the cross sections in an area, or more commonly, the homogenized cross sections may be constructed using a flux-weighted approach with an estimate for the flux in each homogenized area. In practice, some form of flux-weighted homogenization is commonly used, intending to preserve as many core parameters as possible, such as reaction rates, leakages or reactor eigenvalue. Generally, when a flux-weighted homogenization is performed, it is with an estimate for the flux, as whole-core flux profiles are unavailable at the beginning of a calculation.

More advanced homogenization methods combine the homogenization process with some other modification to the transport equation or to the solution procedure in order to improve accuracy. Some examples of this practice are modifications to the boundary conditions of certain homogenized regions [5] [6] or an implementation with some form of a correction term added to the equation [1] [7] [8]. More information about the theory of homogenization can be found in references [9] [10].

2.3 Current Technologies

In this section, current technologies in homogenization will be briefly reviewed and discussed, ending with the CSH and DTH methods, as they were

implemented in 1-D. This review should contextualize the base upon which the 2-D implementations of the CSH and DTH methods have been built.

2.3.1 Other homogenization methods

The CSH and DTH methods are both assembly homogenization techniques, meaning that both methods seek to homogenize cross sections on an assembly level. In transport theory, one such technique is black-box homogenization using discontinuity factors as developed by Sanchez [9]. The overall idea of black-box homogenization with discontinuity factors is that a single assembly is homogenized so that the when the modified incoming current and modified source term to a single assembly is applied to the assembly, then that assembly will reproduce the reaction rates and the outgoing current of the heterogeneous reference. This homogenization with discontinuity factors is generalized in the sense that is can be applied in either transport theory or diffusion theory, through the appropriate use of high and low order operators, so long as leakage is calculable. Ultimately, flux- or current-discontinuity factors are applied in order to calculate the best approximation for the homogeneous incoming current to each assembly. Black-box homogenization with flux discontinuity factors are useful for very quickly reproducing some heterogeneous reactor parameters, but they do not lend themselves necessary to iterated approaches with increased accuracy at each iteration.

One method which does allow for arbitrary accuracy of the homogenized problem is the high-order cross-section homogenization method of Rahnema and McKinley [6]. This method mixes the concept of flux discontinuity factors with on-the-fly updates in order to correct an estimate for reactor core environment effects. The method of Rahnema and McKinley was implemented for nodal diffusion, but it can be seen as a simple precursor to the CSH method. The use of on-the-fly updating of homogenized parameters through some re-homogenization in order to account for core environment effects is a key element of the CSH and DTH methods, and the CSH method is in some sense a natural extension of this method to higher accuracy.

Many other advanced homogenization methods exist, such as Anistratov's Quasi-Diffusion formulation [8], the leakage corrected assembly homogenization technique of Rahnema and Nichita [5] or the cell homogenization techniques of Kozlowski et al. [11], each with their own set of advantages and disadvantages. For a general review of some other advanced homogenization techniques, see References [9] [10].

2.3.2 CSH method

The CSH method operates entirely within transport theory, involving no low-order approximations of the homogeneous transport equation, and instead uses a full transport solve for both the heterogeneous and homogeneous mesh.

All implementations of the CSH method involve employing standard flux-weighted homogenized cross sections with an additional correction term added to the homogeneous neutron transport equation known as the auxiliary source term[1]. The CSH method involves calculating successive estimates of both the homogeneous cross sections and of the auxiliary source term through on-the-fly re-homogenization during the transport calculation. The goal of the CSH calculation is a complete reconstruction of the heterogeneous fine-mesh angular flux solution, along with reactor eigenvalue.

Advances in CSH theory have taken place recently. These advances have improved the re-homogenization method from a series of assembly fixed-source calculations with boundary conditions approximated from the whole-core solution to a system which is based on whole-core transport sweeps on the heterogeneous mesh at each iteration. Additional efficiency improvements that have been documented are a switch from Fourier series spatial basis functions for the auxiliary source term to piecewise-defined B-spline basis functions, as well as an expansion of the auxiliary source term in the angular domain. These enhancements to the method have been dubbed the Efficient Consistent

[1] Previous implementations of the CSH and DTH methods have referred to this term as an auxiliary cross section term. The auxiliary cross section term and the auxiliary source term are identical, except that the auxiliary source term has been scaled using the average scalar flux in each homogenized region. The term "auxiliary source term" is used in order to reduce confusion about its purpose in the homogeneous equation.

Spatial Homogenization (ECSH) method [4], although all future references to the method in this thesis will simply use CSH to refer to a conglomerate of all implementations of the CSH theory.

The CSH method was initially developed by Yasseri and Rahnema [1], and has been tested in 1-D implementations for 1-D problems based off of a boiling water reactor (BWR) core [1], gas-cooled thermal core [12], and to a pressurized water reactor (PWR) core with mixed-oxide (MOX) fuel [4]. The results of these tests have shown that the CSH method affords modest speedup to transport calculations, with speedup factors between 1.2 and 2.5, but at very little cost of accuracy. The CSH method in 1-D has been shown to consistently calculate the reactor eigenvalue to within 10 pcm of the reference, with less than 0.5% mean relative flux error[2]. However, implementations in 1-D neglect some important aspects of spatial homogenization, namely the need to homogenize complicated multidimensional geometry aspects, such as circular pin cells with cladding, and the effect of space-angle coupling on the solution of the neutron transport equation.

[2] The mean relative flux error (MRE) metric will be defined in detail in Section 4.4.

2.3.3 DTH method

One advancement of the CSH method is the DTH method [2] [3], in which the inner homogeneous problem is replaced by a calculation using fine-mesh diffusion. The diffusion equations are significantly simpler than the neutron transport equation due to the elimination of the angular phase space variable, making the homogeneous solution much more efficient. By employing a full transport re-homogenization after each homogeneous solution in order to recalculate the full phase space treatment of the auxiliary source term, the DTH method is able to leverage the computational efficiency of diffusion theory while still maintaining the accuracy of full transport theory. One significant advantage of the DTH method over other diffusion-based homogenization methods is that the natural use of a transport theory re-homogenization stage allows for accurate reconstruction of the full fine-mesh heterogeneous angular flux profile for a given core.

Like the CSH method, the DTH method had only ever been implemented in 1-D, and has been tested against 1-D benchmark problems based off of a BWR core [2] [3], a gas-cooled thermal reactor core [13], and a PWR core with MOX fuel [14]. As expected of the difference between diffusion-based and transport-based methods, the DTH method is computationally faster than the CSH method, at some cost to its accuracy. For the 1-D reactor cores which have been studied, the DTH method has speedup factors between 3 and 13 times, with

eigenvalue errors under 50 pcm, and mean relative flux error of less than 1.8%. This is significantly more error than the CSH method, but the errors are still small, indicating that both methods are able to accurately reconstruct the heterogeneous angular flux profile with some significant speedup. As with the CSH method, no previous study has considered the impact of more difficult geometry in 2-D and 3-D, nor the effect of space-angle coupling which is not present in 1-D.

CHAPTER 3. CONSISTENT SPATIAL HOMOGENIZATION

In this chapter, the consistent spatial homogenization (CSH) method and the diffusion-transport homogenization (DTH) method will be derived, improvements to the derivation and to the method will be discussed, and details of the implementation of CSH and DTH in 2-D will be discussed. Particular attention will be paid to the issues that arise when extending the CSH and DTH methods to two dimensions and to the solutions.

3.1 The CSH Method

3.1.1 Theory

In this section, the updated version of the CSH equations will be derived. The heterogeneous angular flux within some region V^h is with isotropic fission and a scattering kernel dependent only on the scattering angle cosine $\mu_0 = \widehat{\Omega} \cdot \widehat{\Omega}'$ is governed by the neutron transport equation,

$$\begin{aligned} &\widehat{\Omega} \cdot \nabla\psi(\vec{r}, E, \widehat{\Omega}) + \sigma(\vec{r}, E)\psi(\vec{r}, E, \widehat{\Omega}) \\ &\qquad = \int_{4\pi} \mathrm{d}\widehat{\Omega}' \int \mathrm{d}E' \sigma_s(\vec{r}, E' \to E, \mu_0)\psi(\vec{r}, E', \widehat{\Omega}') \\ &\qquad + \frac{\chi(E)}{4\pi k} \int_{4\pi} \mathrm{d}\widehat{\Omega}' \int \mathrm{d}E' \nu\sigma_f(\vec{r}, E')\psi(\vec{r}, E', \widehat{\Omega}') \quad \forall \vec{r} \in V^h. \end{aligned} \tag{5}$$

The fundamental assumption of the CSH method is that Equation (5) can be consistently homogenized through the addition of an auxiliary source term. This equation is given by

$$\begin{aligned}\widehat{\Omega}\cdot\nabla\psi^h(\vec{r},E,\widehat{\Omega})+\ \sigma^h(E)\psi^h(\vec{r},E,\widehat{\Omega})&\\ =\ \int_{4\pi}\mathrm{d}\widehat{\Omega}'\int\mathrm{d}E'\sigma_s^h(E'\to E,\mu_0)\psi^h(\vec{r},E',\widehat{\Omega}')&\\ +\frac{\chi(E)}{4\pi k^h}\int_{4\pi}\mathrm{d}\widehat{\Omega}'\int\mathrm{d}E'\nu\sigma_f^h(E')\psi^h(\vec{r},E',\widehat{\Omega}')&\\ +\ S_{aux}(\vec{r},E,\widehat{\Omega})\qquad\qquad\forall\vec{r}\in V^h,&\end{aligned}\tag{6}$$

Where the superscript h indicates homogenized values. All spatial dependencies in cross section values have been removed in the homogenized equation, and the effects of the heterogeneous cross sections are folded into the auxiliary source term, S_{aux}. At this point, no assumption has been made about *how* the cross sections have been homogenized only that they have had their dependence on $\vec{r}$ removed in some manner. An expression for the auxiliary cross section is found by assuming that the angular flux and eigenvalue solutions of Equation (5) are identical to the solutions to Equation (6). The equations can then be subtracted, and the auxiliary source term can be defined

$$S_{aux}(\vec{r},E,\widehat{\Omega}) = \int_{4\pi} \mathrm{d}\widehat{\Omega}' \int \mathrm{d}E'\ \Delta\sigma_s(\vec{r},E'\to E,\mu_0)\psi(\vec{r},E',\widehat{\Omega}')$$
$$+\frac{1}{4\pi k}\int_{4\pi} \mathrm{d}\widehat{\Omega}' \int \mathrm{d}E'\ \chi(E)\Delta\nu\sigma_{\mathrm{f}}(\vec{r},E')\psi(\vec{r},E',\widehat{\Omega}') \tag{7}$$
$$-\ \Delta\sigma(\vec{r},E)\psi(\vec{r},E,\widehat{\Omega}),$$

Where $\Delta\sigma_i \equiv \sigma_i(\vec{r}) - \sigma_i^h$ for $i = t,s,f$. In the traditional derivation of CSH, the auxiliary source term is then expanded into a set of continuous orthogonal basis functions in both space and angle [1], however more recent derivations of the CSH method have taken to using piecewise-defined basis functions [3] [4] in the spatial domain in order to avoid the need of having extreme high order spatial basis functions and also to better capture the detailed spatial distribution of the auxiliary source. Until now, every derivation of CSH has included a spherical harmonics expansion in angle, however the effect of different angular basis functions has never been studied, and in practice the CSH method has only ever been implemented using discrete ordinates codes, where the angular expansion can be entirely ignored in favor of a quadrature-based approach. In light of this fact and of the eventual implementation of the CSH method in a discrete ordinates context, the auxiliary source term will be left unexpanded in angle. However, for the sake of the derivation and for completeness, the auxiliary source will still be decomposed into arbitrary

spatial basis functions, which can be defined in a piecewise manner and which may or may not be complete. From this step on, due to the potential incompleteness of the spatial basis functions, the auxiliary source term can only be approximated. The auxiliary source term is separated into its components and approximated via expansion into spatial basis functions as

$$S_{aux}(\vec{r}, E, \hat{\Omega}) \approx \sum_{a=0}^{A} E_a(\vec{r})\alpha_a(E, \hat{\Omega}) + \sum_{b=0}^{B} F_b(\vec{r})\beta_b(E, \hat{\Omega}) + \sum_{c=0}^{C} G_c(\vec{r})\gamma_c(E, \hat{\Omega}). \quad (8)$$

where $E_a, F_b,$ and G_c are some arbitrary, normalized spatial basis functions, and where $\alpha_a, \beta_b,$ and γ_c are defined as the components of each spatial expansion order for the scattering, fission, and total cross section components of the auxiliary source term. If $E_a, F_b,$ and G_c are orthogonal, they can be defined as,

$$\alpha_a(E, \hat{\Omega}) = \int_{4\pi} \mathrm{d}\hat{\Omega}' \int \mathrm{d}E' \int_{V^h} \mathrm{d}\vec{r}\ E_a(\vec{r})\Delta\sigma_s(\vec{r}, E' \to E, \mu_0)\psi(\vec{r}, E', \hat{\Omega}'), \quad (9)$$

$$\beta_b(E, \hat{\Omega}) = \frac{1}{4\pi k} \int_{4\pi} \mathrm{d}\hat{\Omega}' \int \mathrm{d}E' \int_{V^h} \mathrm{d}\vec{r}\ F_b(\vec{r})\ \chi(E)\Delta\nu\sigma_{\mathrm{f}}(\vec{r}, E')\psi(\vec{r}, E', \hat{\Omega}'), \quad (10)$$

and

$$\gamma_c\left(E,\widehat{\Omega}\right) = \int_{V^h} \mathrm{d}\vec{r}\; G_c(\vec{r})\Delta\sigma(\vec{r},E)\psi(\vec{r},E,\widehat{\Omega}). \tag{11}$$

Equations (9) through (11) have been left in their fully generalized form. Within the framework of CSH there is no reason to necessarily expect that the same basis functions will be used for each component of the auxiliary source. Indeed, there may be good reasons to use differing basis functions for each component, if – for example – one component of the auxiliary source can be expected to vary in some predictable manner. In practice, it is not expected that any accuracy will be gained by using different spatial expansions for each component of the auxiliary cross section, and all three spatial basis function sets can be reliably replaced with G_c without compromising the method. In fact, all previous implementations of CSH and all of its variants have used a single set of spatial basis functions for each component of the auxiliary source term, including the implementations found here.

Since the manner in which the cross sections were homogenized was originally left unspecified, the CSH equations have not yet been fully constrained. The additional constraint that is added to the system is the requirement that each component (total, scattering, and fission) of the auxiliary source integrates to zero. Equivalently, the system is constrained such that the reaction rates determined by taking the 0^{th} angular moment of the auxiliary source term are conserved between the heterogeneous and

homogeneous solutions. This constraint results in definitions for $\sigma^h(E)$, $\sigma_s^h(E' \to E, \mu_0)$, and $\chi(E)\nu\sigma_f^h(E')$ as the traditional flux-weighted homogenizations,

$$\sigma^h(E) = \frac{\int_{V^h} \mathrm{d}\vec{r}\, \sigma(\vec{r}, E)\phi(\vec{r}, E)}{\int_{V^h} \mathrm{d}\vec{r}\, \phi(\vec{r}, E)}, \tag{12}$$

$$\sigma_s^h(E' \to E) = \frac{\int_{V^h} \mathrm{d}\vec{r} \int_{-1}^{1} \mathrm{d}\mu_0\ \sigma_s(\vec{r}, E' \to E, \mu_0)\phi(\vec{r}, E)}{\int_{V^h} \mathrm{d}\vec{r}\, \phi(\vec{r}, E)}, \tag{13}$$

and

$$\chi(E)\nu\sigma_f^h(E') = \frac{\int_{V^h} \mathrm{d}\vec{r}\, \chi(E)\nu\sigma_f^h(\vec{r}, E')\phi(\vec{r}, E')}{\int_{V^h} \mathrm{d}\vec{r}\, \phi(\vec{r}, E')}, \tag{14}$$

where $\phi(\vec{r}, E)$ is the scalar flux profile. Equation (13) is written for isotropic scattering.

3.1.2 Solution procedure

In the previous section, the fundamental equations of the CSH method were derived. However, the definitions found in Equations (9) through (14) require *a priori* knowledge of the heterogeneous angular and scalar flux profiles. To address this, the heterogeneous flux in each homogenized region is approximated through the use of single-assembly calculations with specular reflective boundary conditions, the same process done in standard

homogenization methods and cross section calculations. The advantage of the CSH method is that since the entire calculation is performed in transport theory, the problem can be iterated upon and re-homogenized on-the-fly at each step, resulting in new homogenized cross sections as well as an updated approximation for the auxiliary source term. Following this logic, the overall solution procedure for the CSH method is described in the steps below

1. Perform heterogeneous assembly-level eigenvalue calculations with approximate (specular reflective) boundary conditions to generate initial homogenized cross sections and the initial auxiliary source term for each assembly.
2. Solve the homogeneous full-core transport equation (Equation (6)) using the auxiliary source term in order to generate solutions for the angular flux and eigenvalue of the core.
3. Re-homogenize the problem using the solution from step 2. This can be done in one of two ways.
    a. Evaluate the angular flux at assembly interfaces using some spatial and angular basis functions at the interfaces, then solve the heterogeneous single-assembly problems using these interface fluxes as incoming boundary conditions. These problems are solved with a fixed eigenvalue equal to the core eigenvalue calculated in step 2. The scattering and fission sources are

updated as usual, but the incoming boundary conditions are held fixed. In practice, it is usually appropriate to use the same spatial and angular basis functions as were used to express the auxiliary source term in the homogenized equation. This method is referred to as "assembly-fixed-source" re-homogenization.

b. Expand the whole core angular flux in some set of spatial basis functions, usually the same set as was used to expand the auxiliary source term, and then perform a small number of fixed-eigenvalue 'sweeps' on the heterogeneous problem using the eigenvalue from step 2. This method is referred to as "core-sweep" re-homogenization. In each sweep, the eigenvalue and fission source terms are updated, and a normalization is applied at the end in order to correct for not updating the eigenvalue.

4. Repeat steps 2 and 3 until the flux and eigenvalue meet the convergence criteria,

$$\max_{Assemblies} \max_{E} \frac{\left|\phi_{avg}^{h,n}(E) - \phi_{avg}^{h,n-1}(E)\right|}{\phi_{avg}^{h,n}(E)} < \epsilon_{\phi}, \tag{15}$$

and

$$\frac{|k^{h,n} - k^{h,n-1}|}{k^{h,n}} < \epsilon_{k}, \tag{16}$$

for some user-defined values of ϵ_ϕ and ϵ_k. In Equation (15), $\phi_{avg}^{h,n}$ is the average flux in a given assembly, and in both Equation (15) and Equation (16), n represents the current outer iteration number.

3.2 The DTH Method

The CSH method can be extended to diffusion theory in what is called the diffusion-transport hybrid (DTH) method. While the DTH method no longer has the consistency afforded by being fully implementable in transport theory, it gains significant computational speed at very little cost of accuracy. One key unrealized advantage of the DTH method over the CSH method is that theories for solving homogenized systems in diffusion theory are significantly more developed than for transport theory, including the potential for modified nodal methods. All current implementations of the DTH method employ fine-mesh diffusion, but a discussion of these potential acceleration techniques takes place in CHAPTER 7. The fundamental difference between the DTH and CSH methods is that in DTH, the inner homogenized equation is solved via diffusion theory rather than in transport theory. This difference causes several significant implications for the solution method, as well. In this section, the DTH method will be derived, following the procedure found in [2], and the differences in solution method between the DTH and CSH methods will be discussed.

3.2.1 DTH Theory

Starting from Equation (6), the homogeneous transport equation with an auxiliary source term, assume linearly anisotropic scattering to obtain

$$\begin{aligned}
&\widehat{\Omega}\cdot\nabla\psi^h(\vec{r},E,\widehat{\Omega})+\ \sigma^h(E)\psi^h(\vec{r},E,\widehat{\Omega})\\
&\qquad=\frac{1}{4\pi}\int\limits_{4\pi}\mathrm{d}\widehat{\Omega}'\int\mathrm{d}E'[\sigma_{s0}^h(E'\to E)\\
&\qquad+3\mu_0\sigma_{s1}^h(E'\to E)]\psi^h(\vec{r},E',\widehat{\Omega}')\\
&\qquad+\frac{\chi(E)}{4\pi k^h}\int\limits_{4\pi}\mathrm{d}\widehat{\Omega}'\int\mathrm{d}E'\nu\sigma_f^h(E')\psi^h(\vec{r},E',\widehat{\Omega}')\\
&\qquad+\ S_{aux}(\vec{r},E,\widehat{\Omega})\qquad\qquad\forall\vec{r}\in V^h,
\end{aligned}\tag{17}$$

where σ_{s0}^h and σ_{s1}^h are the homogenized zeroth and first order Legendre polynomial expansions of the scattering kernel in the scattering angle $\mu_0=\ \widehat{\Omega}\cdot\widehat{\Omega}'$. The diffusion equations are derived by taking the 0th and 1st angular moments of Equation (17) resulting in the equations

$$\begin{aligned}
&\nabla\cdot\vec{J}^h(\vec{r},E)+\ \sigma^h(E)\phi^h(\vec{r},E)\\
&\qquad=\int\mathrm{d}E'\ \sigma_{s0}^h(E'\to E)\phi^h(\vec{r},E')+\frac{\chi(E)}{k^h}\int\mathrm{d}E'\,\nu\sigma_f^h(E')\phi^h(\vec{r},E')\\
&\qquad+\ S_{aux}^0(\vec{r},E)
\end{aligned}\tag{18}$$

and

$$\nabla \cdot \Pi^h(\vec{r}, E) + \sigma_{tr}^h \vec{J}^h(\vec{r}, E) = \vec{S}_{aux}^1(\vec{r}, E). \tag{19}$$

In these equations, S_{aux}^0 and S_{aux}^1 are defined as the 0th and 1st angular moments of $S_{aux}(\widehat{\Omega})$; the tensor term, Π^h is defined as the 2nd angular moment of $\psi^h(\widehat{\Omega})$, and $\sigma_{tr}^h(E)$ is defined as the homogenized transport cross section. These definitions are

$$S_{aux}^0(\vec{r}, E) = \int_{4\pi} \mathrm{d}\widehat{\Omega}\, S_{aux}(\vec{r}, E, \widehat{\Omega}), \tag{20}$$

$$\vec{S}_{aux}^1(\vec{r}, E) = \int_{4\pi} \mathrm{d}\widehat{\Omega}\, \widehat{\Omega} S_{aux}(\vec{r}, E, \widehat{\Omega}), \tag{21}$$

$$\Pi^h(\vec{r}, E) = \int_{4\pi} \mathrm{d}\widehat{\Omega}\, (\widehat{\Omega}\widehat{\Omega}) \psi^h(\vec{r}, E, \widehat{\Omega}), \tag{22}$$

and

$$\sigma_{tr}^h(E) = \sigma^h(E) - \sigma_{s1}^h(E). \tag{23}$$

The definition of the transport cross section in Equation (23) implicitly assumes the principle of detailed balance, $\int dE' \sigma_{s1}(E' \to E) \vec{J}(\vec{r}, E') = \int dE' \sigma_{s1}(E \to E') \vec{J}(\vec{r}, E)$, which is common for multi-group or continuous-energy diffusion derivations. It should be noted explicitly that S_{aux}^1 is a vector. The

second order term Π^h is handled using the P_1 approximation, $\nabla \cdot \Pi^h(\vec{r}, E) = \frac{1}{3}\nabla\phi^h(\vec{r}, E)$. Making all of the substitutions implied above and substituting Equation (19) into Equation (18) yields the DTH equation,

$$\begin{aligned} -\nabla \cdot D^h(E)\nabla\phi^h(\vec{r}, E) + \ \sigma^h(E)\phi^h(\vec{r}, E) \\ = \int \mathrm{d}E' \ \sigma_{s0}^h(E' \to E)\phi^h(\vec{r}, E') \\ + \frac{\chi(E)}{k^h}\int \mathrm{d}E' \ \nu\sigma_f^h(E')\phi^h(\vec{r}, E') + \ S_{aux}^0(\vec{r}, E) \\ - \frac{\nabla \cdot \vec{S}_{aux}^1(\vec{r}, E)}{\sigma_{tr}^h(E)} \qquad \forall \vec{r} \in V^h, \end{aligned} \tag{24}$$

where the homogenized diffusion coefficient $D^h(E)$ is defined in the usual manner as $\frac{1}{3\sigma_{tr}^h(E)}$. While D^h is spatially constant within each homogenized region, in implementation a single core may be composed of several homogenized regions, and so the term D^h must be treated as being only piecewise continuous across each homogenized region. This is why the first term is kept as $\nabla \cdot D^h(E)\nabla\phi^h(\vec{r}, E)$ instead of being simplified to $D^h(E)\nabla^2\phi^h(\vec{r}, E)$.

3.2.2 Solution method

The solution method for DTH is similar to the solution method for CSH but different enough to warrant discussion. As with the CSH method, DTH

employs an outer iteration with on-the-fly re-homogenization at each step in order to update the auxiliary source term and homogenized cross sections. Unlike the CSH method, the DTH method has additional possible methods for re-homogenization, especially for reconstructing the angular flux from the solution to the homogenized equations. The DTH method is employed using the following steps, starting with the same initial step as the CSH method,

1. Perform heterogeneous assembly-level calculations with approximate (specular reflective) boundary conditions to generate initial homogenized cross sections and the initial auxiliary source term for each assembly.
2. Solve the homogeneous full-core diffusion equation (Equation (24)) using the auxiliary first and second angular moments of the auxiliary source term as defined in Equations (20) and (21) in order to generate solutions for the scalar flux and eigenvalue of the core.
3. Reconstruct the angular flux from the scalar flux using the P_1 relation,

$$\psi^h\left(\vec{r}, E, \hat{\Omega}\right) \approx \frac{1}{4\pi}\left(\phi^h(\vec{r}, E) + 3\hat{\Omega} \cdot \vec{J}(\vec{r}, E)\right), \tag{25}$$

and then either perform core sweep or assembly fixed-source re-homogenization calculations, as described in step 3 of the CSH solution procedure. Previous work in the DTH method has shown that the P_1

approximation for angular flux reconstruction is sufficient for getting reasonable flux profile accuracy, compared to other reconstruction methods [2] [4]. In particular, it the P_1 approximation has the advantage of being easily integrated into either the assembly fixed-source or core-sweep schemes. Previous iterations of the DTH method employed several other methods of reconstructing the angular flux including interpolations that used weighted values of the previous iteration's guess at the angular flux.

4. Repeat steps 2 and 3 until the flux and eigenvalue meet the convergence criteria defined in Equations (15) and (16) for some user-defined values of ϵ_ϕ and ϵ_k.

Steps 1 and 4 are identical to the solution procedure for the CSH method.

3.3 Implementation in 2-D

Multidimensional implementations of the CSH and DTH methods have never before been attempted. The details of the 2-D implementation of the CSH and DTH methods will be discussed here along with the problems that arise in that implementation. In all cases, the heterogeneous component of the CSH and DTH methods, Equation (5), has been treated with discrete ordinates in angle and the multigroup approximation in energy. This eliminates the need to consider basis functions in angle, as the auxiliary source can simply be

evaluated for each ordinate. The only set of spatial basis functions which warrants further study is the set used to expand the auxiliary source and the angular flux in the spatial variable.

3.3.1 Selection of spatial basis functions for auxiliary source

Ultimately, the purpose of the spatial expansion function for the auxiliary source as well as for the angular flux is to perform an interpolation from the spatial mesh used to solve the heterogeneous problem to the spatial mesh used to solve the homogeneous problem and back for the re-homogenization step. Therefore, discussion of choices for basis functions must first be informed by an understanding of assumptions made by the spatial discretization used to numerically solve the transport and diffusion equations in the CSH and DTH methods. The method used in this thesis for the transport solver for the heterogeneous transport equations is the diamond difference S_N method as derived in reference [15]. This spatial discretization is derived by integrating the transport equation over some rectangular cell (i,j) that extends from $x_{i-1/2} = x_i - \frac{\Delta x_i}{2}$ to $x_{i+1/2} = x_i + \frac{\Delta x_i}{2}$ in the x-direction and from $y_{j-1/2} = y_j - \frac{\Delta y_j}{2}$ to $y_{j+1/2} = y_j + \frac{\Delta y_j}{2}$ in the y-direction. The resultant spatially discretized equation is the balance relation,

$$\frac{\mu_n}{\Delta x_i}\left(\psi^{n,g}_{i+1/2,j} - \psi^{n,g}_{i-1/2,j}\right) + \frac{\eta_n}{\Delta y_j}\left(\psi^{n,g}_{i,j+1/2} - \psi^{n,g}_{i,j-1/2}\right) + \sigma^g_{i,j}\psi^{n,g}_{i,j} = q^{n,g}_{i,j} \quad . \tag{26}$$

In this equation μ represents the x-component of $\widehat{\Omega}$, and η represents the y-component of $\widehat{\Omega}$. The index n is the discrete ordinates quadrature index, the index g represents the group index, and the right-hand side term $q^{n,g}_{i,j}$ encapsulates all source terms, including scattering, fission, and auxiliary. In Equation (26), $\psi^{n,g}_{i,j}$ is defined as the *average* value of $\psi^{n,g}(\vec{r})$ over the cell (i,j). Likewise, the cell-edge flux $\psi^{n,g}_{i+1/2,j}$ is defined to be the average value of $\psi^{n,g}(\vec{r})$ over the right edge of cell (i,j). This definition of the discretized angular flux is used to inform the process of numerical integration on the mesh. For some mesh-defined quantity $f_{i,j}$, its integral over some volume V is numerically expressed as

$$\int_V \mathrm{d}\vec{r}\, f(\vec{r}) = \sum_{(i,j)\in V} f_{i,j}\Delta x_i \Delta y_j \,. \tag{27}$$

The definition of cell-centered quantities as being the average value over a cell, along with the numerical integration given in Equation (27) together define an interpolation scheme for quantities in CSH and DTH theory. Assuming one wants to use $f_{i,j}\ \forall i,j \in V$ defined on one mesh to calculate $f_{a,b}$ on some other mesh, the process for calculating $f_{a,b}$ is

1. Use the spatial basis functions to generate a piecewise continuous representation of $f(\vec{r})$ using the discrete values $f_{i,j}$, $f(\vec{r}) \approx \sum_p f_p G_p(\vec{r})$, for some spatial basis functions $G_p(\vec{r})$.
2. Integrate the resultant continuous function to find its average value over each of the cells of the new mesh (a, b),

$$f_{a,b} = \frac{1}{V_{a,b}} \sum_p \int_{V_{a,b}} \mathrm{d}\vec{r}\ f_p G_p(\vec{r})\,. \tag{28}$$

In previous implementations of the CSH and DTH methods, the auxiliary source term has been expressed either as a Fourier series or a B-spline expansion. Fourier series expansion of the auxiliary source term was not considered at all for use in 2-D due to its very high order requirements. For relatively simple reactor benchmarks in one dimension, Fourier expansion orders of 30 or higher were necessary to achieve full accuracy of the CSH method [1]. In 2-D, and with the complicated geometries that the extension to 2-D affords the method, these expansion orders would need to be significantly higher, well beyond the realm of what is reasonable. Instead, three sets of spatial basis are presented in this section for expanding the auxiliary source, the linear B-spline bases that were used in previous implementations and two new 'integral-conserving' spatial basis functions. As will be explained in the next section and as will become evident in CHAPTER 5, the linear B-spline based expansion functions have poor performance when used with the CSH

and DTH methods in 2-D, while both the 0^{th} and 1^{st} order 'integral-conserving' basis functions have significantly better performance.

3.3.1.1 B-spline basis functions

The use of linear B-spline basis functions is effectively the same as linear interpolation for generating a continuous representation of a mesh-defined quantity. In one dimensional implementations of CSH and DTH, this representation had great performance, accurately representing both flux and auxiliary source. However even in 1-D, the B-spline representation has some problems when using the mesh transfer routine given by Equation (28). Namely, the B-spline expansion does not locally conserve integrals. A simple graphical depiction of this effect can be seen in Figure 1. In this graphic, the black circles indicate function values on some mesh marked by the dotted red lines. The solid black line depicts the B-spline representation of the continuous function, and the shaded triangles show areas which cause inconsistency. In more than one spatial dimension, it is predicted that this effect will have a much more significant impact on the accuracy of both the flux and eigenvalue.

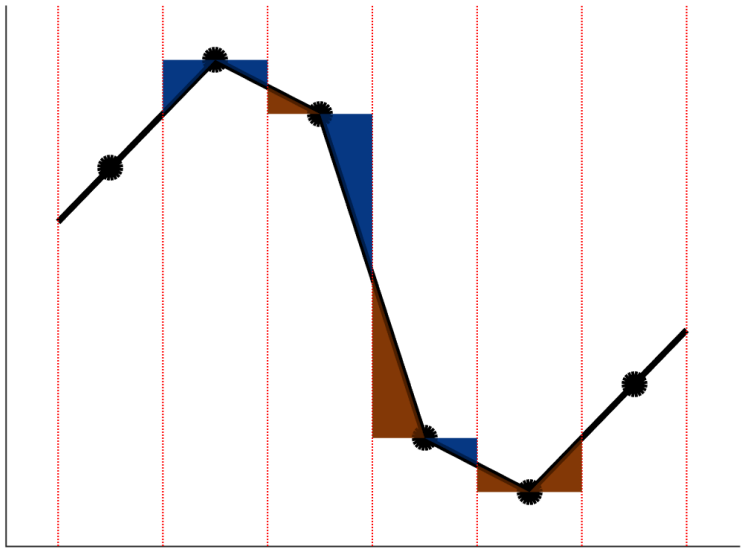

Figure 1. A 1-D depiction of how B-spline basis functions do not conserve integrals. Shaded triangles indicate areas where the integral estimate would be either too high or too low.

3.3.1.2 Integral-conserving basis functions

There are two advantages to the integral-conserving basis functions over linear B-spline expansion. The first advantage is obvious in the name: the integral-conserving spatial basis functions conserve integrals, at least when spatial integrals are computed in the manner given by Equation (27). The second advantage is that the integral-conserving spatial basis functions enable the expanded function to have discontinuities across cell interfaces. At first glance, this effect seems like it may be detrimental to the CSH and DTH

methods; however, the auxiliary source in the CSH and DTH methods is not necessarily continuous across cell boundaries because it can change proportionally to cross section changes, which are necessarily defined as piecewise constant on each cell.

The 0^{th} order integral-conserving basis function is defined simply to be a constant value on each cell, with no other modifications needed. The 1^{st}-order integral-conserving basis function is a slope-limited plane. It is defined on each cell as the plane which passes through the point $(x_i, y_i, f_{i,j})$ that has a slope in the x direction equal either to the forward difference or the backward difference, whichever has smaller magnitude and the same for the y direction. A graphical depiction of the 0^{th} and 1^{st} order integral-conserving basis functions can be found in Figure 2. A benefit of the integral-conserving basis functions is they are less susceptible to affecting reaction rates when used to transfer angular flux between different spatial meshes.

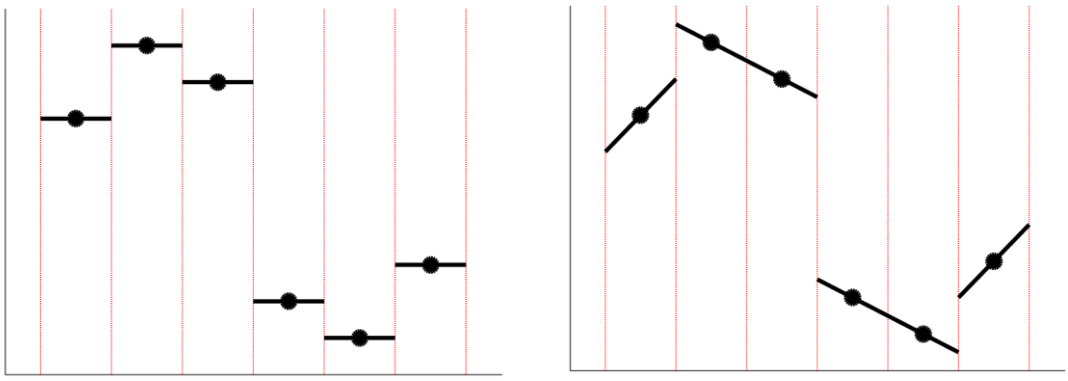

Figure 2. 1-D graphical depictions of 0th order (left) and 1st order (right) integral-conserving basis functions

3.3.2 2-D discretization details

The discretization used for the transport solution component of the CSH and DTH methods is 2-D S_N. The chosen 2-D S_N implementation has been modified to accommodate the auxiliary source term S_{aux}. The heterogeneous and homogeneous transport equation discretization is identical to that of Equation (26) with $q_{i,j}^{n,g}$ defined as the average value of $q^{n,g}(\vec{r})$ over the cell (i,j). When the scattering is isotropic, as is the case for both benchmark problems described in CHAPTER 4, the spatially dependent source term $q_{i,j}^{n,g}$ is defined as

$$q_{i,j}^{n,g} = \sum_{n'} w_{n'} \sum_{g'} \sigma_{s0,i,j}^{g' \to g} \psi_{i,j}^{n',g'} + \frac{\chi_{i,j}^{g}}{k} \sum_{n'} w_{n'} \sum_{g'} \nu \sigma_{f,i,j}^{g'} \psi_{i,j}^{n',g'} + S_{aux,i,j}^{n,g} . \qquad (29)$$

where the quadrature weights w_n are assumed to sum to 1, and where $S_{aux,i,j}^{n,g}$ is defined as the average value of the auxiliary source in cell (i,j) corresponding to ordinate n and group g. For the heterogeneous discretization, the same definition for $q_{i,j}^{n,g}$ is used but without the auxiliary source term. The S_N iteration is performed in the traditional manner, as described in reference [15], by marching from each corner in the direction of each ordinate and updating the cell-wise angular flux estimates.

The reactor eigenvalue is calculated through a power iteration, applied after each single source sweep by applying the power iteration formula

$$k^{\text{new}} = k^{\text{old}} \frac{\sum_{g,i,j} \Delta x_i \, \Delta y_j \left(\text{Fiss}_{i,j}^{g,\text{new}}\right)^2}{\sum_{g,i,j} \Delta x_i \, \Delta y_j \text{Fiss}_{i,j}^{g,\text{new}} \text{Fiss}_{i,j}^{g,\text{old}}}, \tag{30}$$

where $\text{Fiss}_{i,j}^{g,\text{new}}$ is the fission source calculated after the previous iteration has completed, and $\text{Fiss}_{i,j}^{g,\text{old}}$ is the fission source that was used in the previous iteration. The fission source $\text{Fiss}_{i,j}^{g}$ is defined

$$\text{Fiss}_{i,j}^{g} = \frac{\chi_{i,j}^{g}}{k} \sum_{n'} w_{n'} \sum_{g'} \nu\sigma_{f,i,j}^{g'} \psi_{i,j}^{n',g'}. \tag{31}$$

3.3.3 Progressively tightened homogeneous mesh and homogeneous convergence

In previous implementations of the CSH and DTH methods in 1-D, significantly improved convergence rates were obtained through the use of progressively tightened convergence criteria for the homogeneous equations. This behavior can be discussed in the context of a simplified 1-D fixed-source version of the CSH equations. In 1-D, the simplified heterogeneous transport equation can be written

$$\mu \frac{\partial \psi^*(x,\mu)}{\partial x} + \sigma(x)\psi^*(x,\mu) = Q(x,\mu). \tag{32}$$

Its corresponding homogenized CSH equation is

$$\mu \frac{\partial \psi^*(x,\mu)}{\partial x} + \sigma^h \psi^*(x,\mu) = Q(x,\mu) + S^*_{aux}(x,\mu), \tag{33}$$

where $S^*_{aux}(x,\mu)$ is the auxiliary cross section calculated using the exact angular flux solution as

$$S^*_{aux}(x,\mu) = -\Delta\sigma(x)\psi^*(x,\mu), \tag{34}$$

and where $\psi^*(x,\mu)$ is the exact solution. By the definition of the auxiliary source, the solutions to Equations (32) and (33) are identical. In practice, due to the fact that the exact solution cannot be known *a priori*, the auxiliary source can only be approximated using some estimated angular flux profile as

$$\begin{aligned} S_{aux}(x,\mu) \approx -\Delta\sigma(x)\psi^{\text{est}}(x,\mu) = \Delta\sigma(x)(\psi^{\text{est}}(x,\mu) + \psi^*(x,\mu) - \\ \psi^*(x,\mu)) = S^*_{aux}(x,\mu) + e_{aux}, \end{aligned} \tag{35}$$

where e_{aux} is the error of the auxiliary term, defined

$$e_{aux} = \Delta\sigma(x)\,(\psi^{\text{est}}(x,\mu) - \psi^{*}(x,\mu)). \tag{36}$$

In order to numerically solve Equations (32) and (33), they must be discretized in space. In practical implementations of the CSH and DTH methods, the heterogeneous and homogeneous equations are likely to have different spatial meshes. Assume that the operators D_{het} and D_{hom} are numerical derivative operators on the heterogeneous problem grid and the homogeneous problem grid, with associated discretization errors e_{het} and e_{hom} respectively. Restricting Equations (32) and (33) to their respective meshes and applying the operators D_{het} and D_{hom} results in

$$\mu D_{het}\psi_i(\mu) + \sigma_i\psi_i(x,\mu) = Q_i(\mu) + e_{het}, \tag{37}$$

and

$$\mu D_{hom}\psi_a(\mu) + \sigma^h\psi_a(\mu) = Q_a(\mu) + S_{aux}(x_a,\mu) + e_{hom} + e_{aux}, \tag{38}$$

where i indicates a cell index on the heterogeneous grid, and a indicates a cell index on the homogeneous grid. Assuming that $Q(x,\mu)$ can be represented on these grids with no loss of accuracy, and with proper definitions of e_{aux}, e_{het}, and e_{hom}, the solutions to each of these equations are both mesh restricted version of the exact solution ψ^{*}. In Equation (38), the auxiliary source term is not yet written in a manner which can be easily evaluated on the new grid (a,b). In order to correct for this, The mesh transfer procedure given in

Equation (28) is applied to Equation (38) resulting in the fully discretized version of the simplified CSH equation,

$$\mu D_{hom}\psi_a(\mu) + \sigma^h\psi_a(\mu) = Q_a(\mu) + S_{aux,a}(\mu) + e_{hom} + e_{aux} + e_{basis}, \tag{39}$$

where e_{basis} is the error introduced through the mesh transfer process using some set of spatial basis functions. In practice, the three error terms of Equation (39) cannot be evaluated and a solution is obtained by neglecting them. If Equation (39) is solved via an iterative method, then the error of the solution to the discretized CSH equation, e_{CSH} will be

$$e_{\mathrm{CSH}} = O(e_{hom} + e_{aux} + e_{basis} + e_{conv}), \tag{40}$$

where e_{conv} is error introduced by eventually cutting off the iterative method to solve for the angular flux. It can be expected that for any reasonable discretization and choice of basis functions that the magnitude of both e_{hom} and e_{basis} will decrease as the homogeneous problem mesh size decreases, however e_{aux} has no dependence on the homogeneous problem mesh size. It is entirely dependent on the accuracy of ψ^{est}. In CSH, ψ^{est} is taken to be either the initial estimate angular flux taken from the single-assembly problems or the output from the previous iteration. The CSH method in general converges to the heterogeneous solution, so the magnitude of e_{aux} can be assumed to decrease with each re-homogenization iteration. Since the homogeneous mesh and the homogeneous solution method convergence criteria are user-defined,

only e_{aux} is rigidly specified in Equation (40). If a user wishes to balance accuracy with computational speed, it would be prudent to choose the loosest convergence criteria and largest homogeneous mesh size such that the magnitude of $(e_{hom} + e_{basis} + e_{conv})$ is less than the magnitude of e_{aux} at each re-homogenization iteration. Previous implementations of the CSH and DTH methods capitalized on this tradeoff by tightening the convergence criteria of the homogeneous problem solve at each re-homogenization step, however this analysis implies that further gains to the performance of the CSH method can be obtained through some corresponding tightening of the homogeneous problem mesh at each re-homogenization iteration. Specific schemes for progressive mesh tightening will be discussed in CHAPTER 5.

CHAPTER 4. BENCHMARK PROBLEMS

This chapter will describe the benchmark problems used to test the CSH and DTH methods as well as explain the figures of merits that will be used to judge the efficacy of the methods. The methods were tested on 2-D discrete ordinates models with Cartesian mesh spatial discretization. Reference solutions for each model were calculated using S_8 with diamond differencing and level-symmetric quadrature using a home-grown standalone 2-D S_N transport package, written in compiled FORTRAN as a robust cross-platform Python library.

4.1 2-D C5G7

The 2-D C5G7 benchmark is a small neutron transport eigenvalue problem developed in 2003 primarily to be used to test deterministic neutron transport methods [16]. The 2-D C5G7 benchmark is representative of a small sixteen assembly pressurized water reactor (PWR) with mixed-oxide (MOX) fuel. Each assembly is constructed of a lattice of 17x17 pin cells, and each pin cell is heterogeneously modeled with two materials as a circular fuel region surrounded by moderator. For a full description of the C5G7 benchmark geometry, see [16] in the references. The heterogeneity of the C5G7 benchmark as well as the ubiquity of this problem in neutron transport methods development makes it a useful test case for the CSH and DTH methods. For this reason, the 2-D C5G7 benchmark will serve as the primary test case for

the CSH and DTH methods. That is, the CSH and DTH methods will be tested for accuracy against flux profiles generated using the reference transport solution method, which cannot perfectly represent circular fuel elements. This fact motivates a study into the spatial discretization used for the benchmark case before any CSH and DTH cases are attempted.

In this section, the 2-D C5G7 benchmark will be broken down into several increasingly accurate Cartesian mesh discretizations in order to settle on a reference benchmark problem and corresponding solution that effectively captures the physics of the 2-D C5G7 benchmark problem without overly simplifying the problem or adding artificial difficulty to the problem.

4.1.1 Benchmark geometry and solution

4.1.1.1 Verification of underlying S_N solution code and choice of C5G7 spatial mesh

Three different spatial meshes for a C5G7 pin cell are presented and analyzed using the reference 2-D S_N solution method. Since the underlying 2-D S_N code is custom-built, these tests will also serve to verify the accuracy of the underlying reference solver. The reference solutions should show similar results to other discrete ordinates solutions of the 2-D C5G7 benchmark problem when similar meshes are used. The three proposed spatial meshes

are based around 4x4, 9x9, and 18x18 mesh pin cell descriptions depicted below in Figure 3.

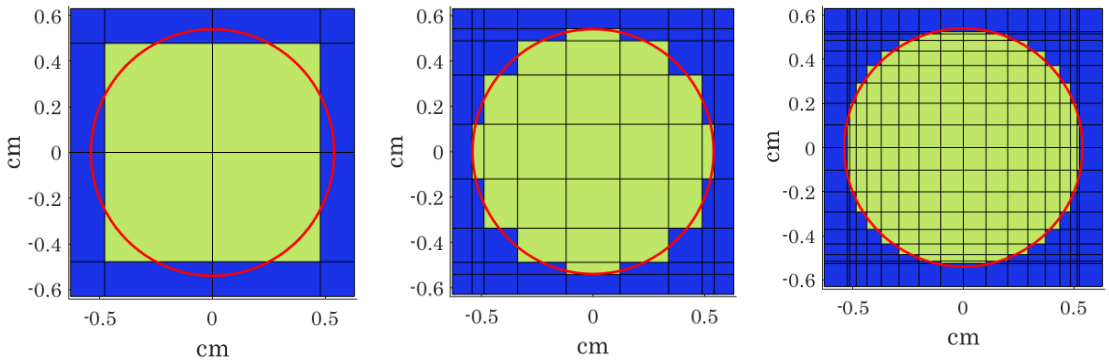

Figure 3. Proposed volume-conserving C5G7 pin cell discretizations. From left to right, 4x4, 9x9, and 18x18 meshes. Exact geometry is depicted as a red circle.

These three depictions were created by first placing mesh boundaries on the fuel periphery with equal-angle spacing, and then by scaling the resulting mesh to achieve exact volume conservation of fuel and moderator. In order to make it possible to reproduce these meshes, the locations of the mesh boundaries relative to the point (0,0) at the center of the pin cell can be found in Table 1, not including the boundary at 0.63 cm. Each cell is rotationally symmetric and has a material identifier as indicated in Figure 3.

Table 1. Proposed volume-conserving C5G7 pin cell discretization mesh boundaries in centimeters

4x4	9x9	18x18
0.0000	0.1207	0.0000
0.4786	0.3381	0.1025
	0.4886	0.2011
	0.5423	0.2919
		0.3715
		0.4369
		0.4854
		0.5153
		0.5254

In order to first verify the accuracy of the 2-D S_N transport solver, these three discretizations were tested along with three increasing level-symmetric quadrature orders against the canonical Monte Carlo reference solution to the C5G7 benchmark problem [16]. Not only does this test verify the efficacy of the 2-D S_N solver, it also provides valuable information about whether or not the proposed pin cell discretizations capture the relevant physics of the C5G7 benchmark problem. The results of this study are presented in Table 2, for eigenvalue error and for the error in pin fission density when calculated for the pin with maximum fission density and for the pin with minimum fission density. This error metric provides a reasonable quantification of overall flux error when compared to the canonical Monte Carlo benchmark results.

Table 2. Verification of 2-D SN code and study into reasonable meshing for Cartesian mesh C5G7.

Pin mesh	Quadrature	Δk (pcm)	Pin fission density error (%)	
			Max. pin	Min. pin
4x4	S_2	62	2.1	3.2
9x9	S_4	45	1.1	3.7
18x18	S_8	111	1.3	0.7

While the eigenvalue convergence in Table 2 is contrary to what one would expect, the pin fission density shows a clear trend, indicating that the eigenvalue results are likely an anomaly due to error cancellation. As expected, the more accurate spatial meshes and more accurate angular quadrature order yield more accurate flux results. Further verification of the underlying S_N solver is done by comparing these results to those generated using other Cartesian mesh discrete ordinates methods in the C5G7 benchmark paper. The results of the four Cartesian mesh discrete ordinates calculations from the C5G7 benchmark specification [16] are included in Table 3. The quadrature orders and specific spatial meshes used by these other codes can be found in [16]. Observation of Table 3 indicates that the magnitude of the errors is similar between the 2-D S_N transport solver used by this thesis and other Cartesian mesh discrete ordinates codes.

With an eye to balancing computational complexity with the ability to capture physics, the 9x9 pin cell representation with S_8 quadrature was chosen for all reference solutions.

Table 3. Other Cartesian mesh C5G7 solution errors, from [16]

Code name	Discretization	Δk (pcm)	Pin fission density error (%)	
			Max. pin	Min. pin
DORT-GRS	$S_{16}, 17x17$	146	0.50	0.02
DORT-ORNL	$S_8, 4x4$	134	0.60	0.22
TWODANT	$S_{16}, 27x27$	11	1.40	0.03
PARTISN	$S_{26}, 15x15$	15	0.18	0.18

4.1.1.2 Reference solutions for the benchmark

The CSH and DTH methods will be evaluated on their ability to accurately calculate reactor eigenvalue and fine-mesh flux profiles. In this subsection, the S_8 reference solution to the 9x9 pin cell problem with half mean free path spacing in the reflector region will be presented. All future references to the "reference solution" of the C5G7 benchmark refer to the solution of the S_8 9x9 volume-conserving pin cell discretization benchmark problem with diamond differencing, power iteration for solution on eigenvalue and converged to 6×10^{-6} in flux and 7×10^{-8} in eigenvalue. The precise meaning of these convergence criteria and the justification for the values chosen is expounded upon in APPENDIX B. These convergence criteria mean that the reference solutions are converged to a maximum error compared to the discretization error of 0.5% in flux and 0.5 pcm in eigenvalue.

The reference flux solution is presented in Figure 4. For the sake of succinct plots, the flux solution has been condensed to two groups whenever it is plotted. The first group (fast) is a sum of groups 1 through 4, and the second group (thermal) is a sum of groups 5, 6, and 7. The reference eigenvalue is 1.185419.

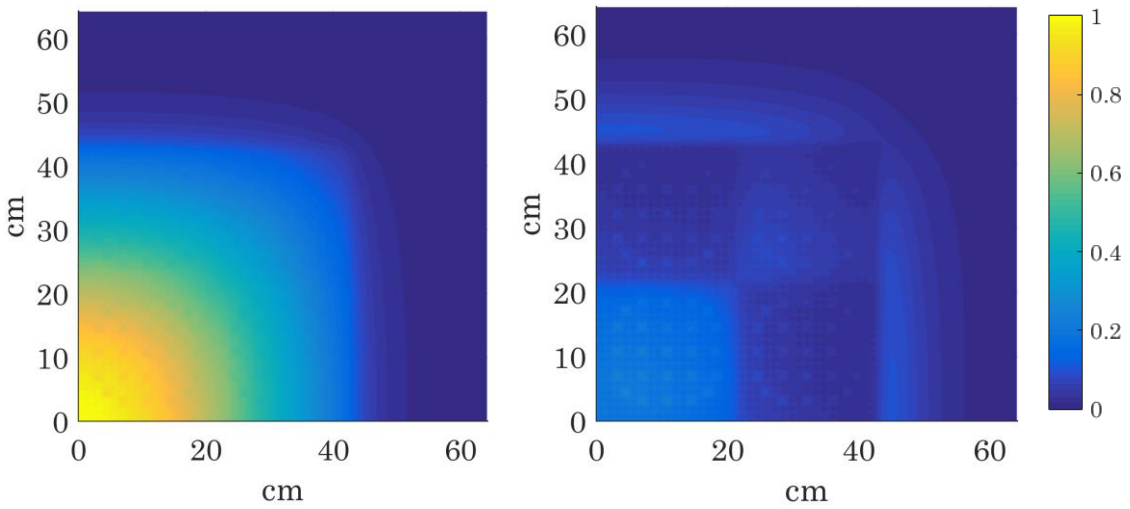

Figure 4. Reference scalar flux solution to C5G7 problem. Left: Fast groups, right: thermal groups. Flux presented in arbitrary units with maximum 1.

4.2 Full scale 2-D BWR benchmark problem

In addition to a Cartesian mesh version of the C5G7 problem, another quarter-core benchmark problem indicative of a full-scale BWR reactor was used to test the efficacy of the CSH and DTH methods. The 2-D model is based off of the model of [17] with minor modifications to work with Cartesian geometry. The benchmark itself is based off of a controlled BWR core with a

HAFAS core layout but with assemblies taken from a GE9-based lattice [18]. The core itself is a checkerboard layout of fresh and depleted assemblies with geometry from the RACER assembly. For added complexity, assemblies near the center of the core are assumed to have 70% or 40% void fraction. Each assembly has a complicated fuel pin layout of ten fuel types, including two gadded pins, resulting in a benchmark problem with a total of 120 unique materials and fully resolved cladding around each pin. For brevity, the following section will only discuss the Cartesian-mesh version of the benchmark problem as tested with the CSH and DTH methods. A full description of the benchmark geometry, materials, and 2-group cross sections can be found in [17].

4.2.1 Cartesian mesh 2-D BWR benchmark description

Very few assumptions were taken in constructing the Cartesian mesh version of the 2-D BWR benchmark problem. The assumptions that were made are documented in this section, along with the information required to reconstruct the Cartesian mesh benchmark from the 2-D benchmark information given in [17].

The major assumption of the benchmark discretization is Cartesian mesh pin cells. The pin cell mesh chosen for this problem is 10x10 and is depicted in Figure 5. This pin cell mesh was calculated by meshing through

points chosen on the fuel and clad surface which are spaced apart with equal angles. These fuel and clad regions were then scaled in size until the pin cell perfectly conserved volume compared to the non-discretized pin cell. The procedure for meshing the cladded pin cell is not straightforward, and the algorithm used to guarantee a representative volume-conserving pin mesh was generated specifically for the purpose of this thesis. For more information about the method used to calculate the appropriate pin cell mesh for the 2-D BWR benchmark, see APPENDIX A. Table 4 includes a list of mesh boundaries for the 2-D Cartesian mesh BWR pin cell. As the pin is symmetric with right-angle rotation, this list of mesh boundaries, along with the information from Figure 5 can be used to accurately reconstruct the pin cell discretization.

The benchmark specification has steel cruciform control structures between assemblies with control rods inserted in the steel. The geometry of the control rods has been neglected in this presentation of the benchmark problem. Instead, the control rods are smeared over the entire section of the cruciform control blades. Figure 6 includes a depiction of a controlled assembly as meshed by this problem, and the rectangular control region can be seen. For a details of the unsimplified assembly geometry, please see the figures of reference [17].

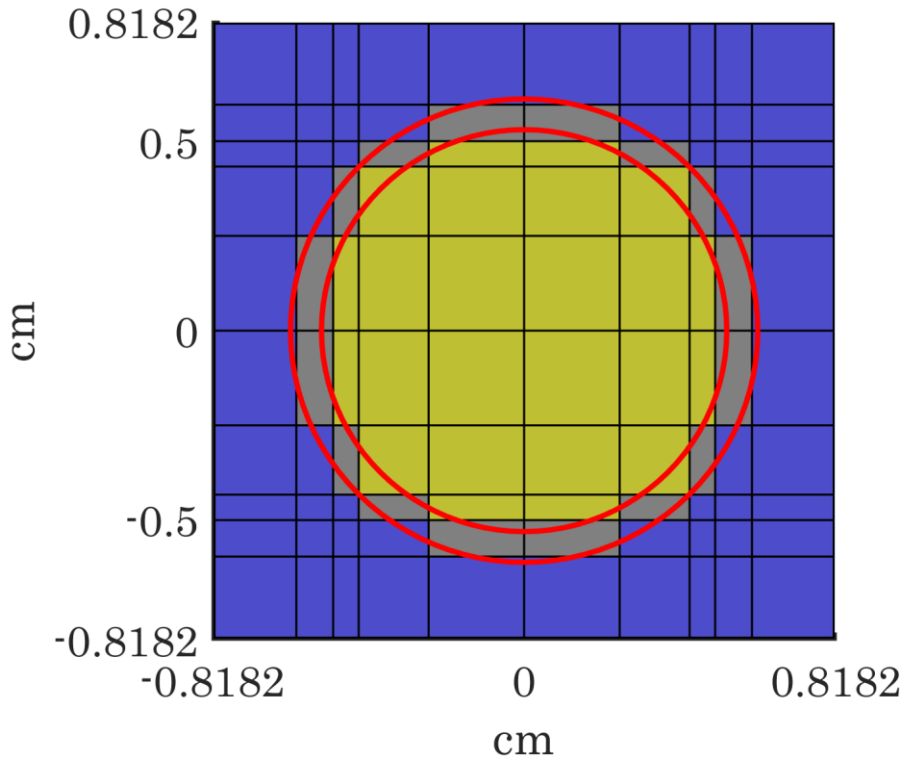

Figure 5. 10x10 volume-conserving Cartesian mesh of cladded 2-D BWR pin cell. Red lines indicate exact geometry. The colors blue, gray, and yellow indicate moderator, cladding, and fuel, respectively.

Table 4. Cartesian mesh boundaries for 2-D BWR pin cell in centimeters, as distance from the center of the pin.

0.0000
0.2508
0.4344
0.5016
0.5982
0.8128

The last simplification made in order to represent the benchmark problem with a Cartesian mesh is to neglect the cladding around the center water channel of each assembly. Instead, the central coolant channel is treated

as if it occupies the entire square section in the center of the assembly with no cladding. Aside from the stated assumptions, the benchmark problem is implemented as described in [17]. The geometry layout of the full core as it is solved for the reference solution and CSH and DTH methods is depicted in Figure 7, with each material in a separate color.

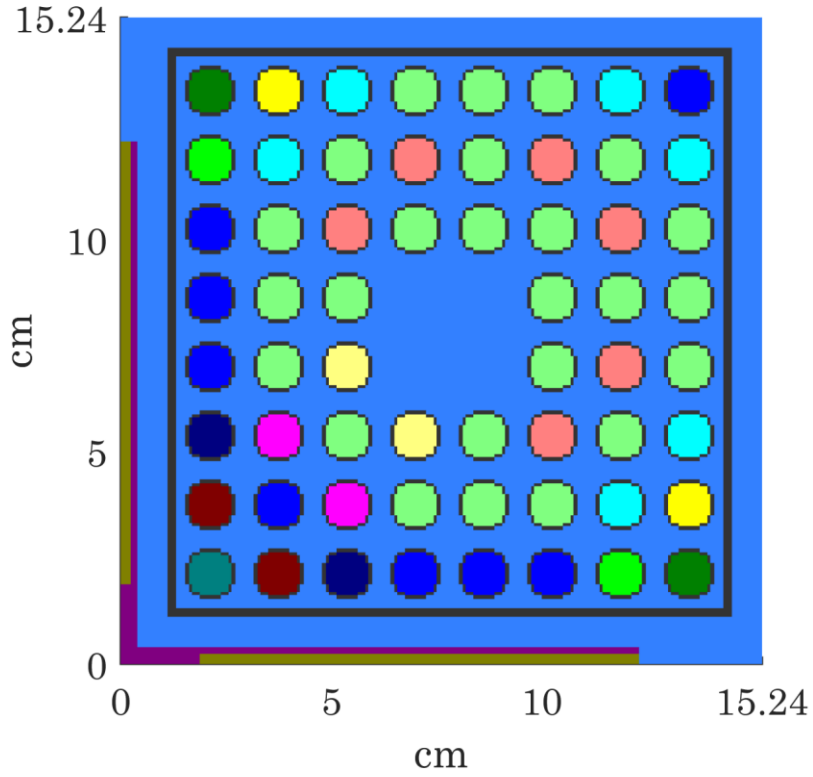

Figure 6. 2-D BWR Cartesian discretized assembly geometry. Each color is a different material. Note the smeared control material (in brown) and the absent central water channel cladding.

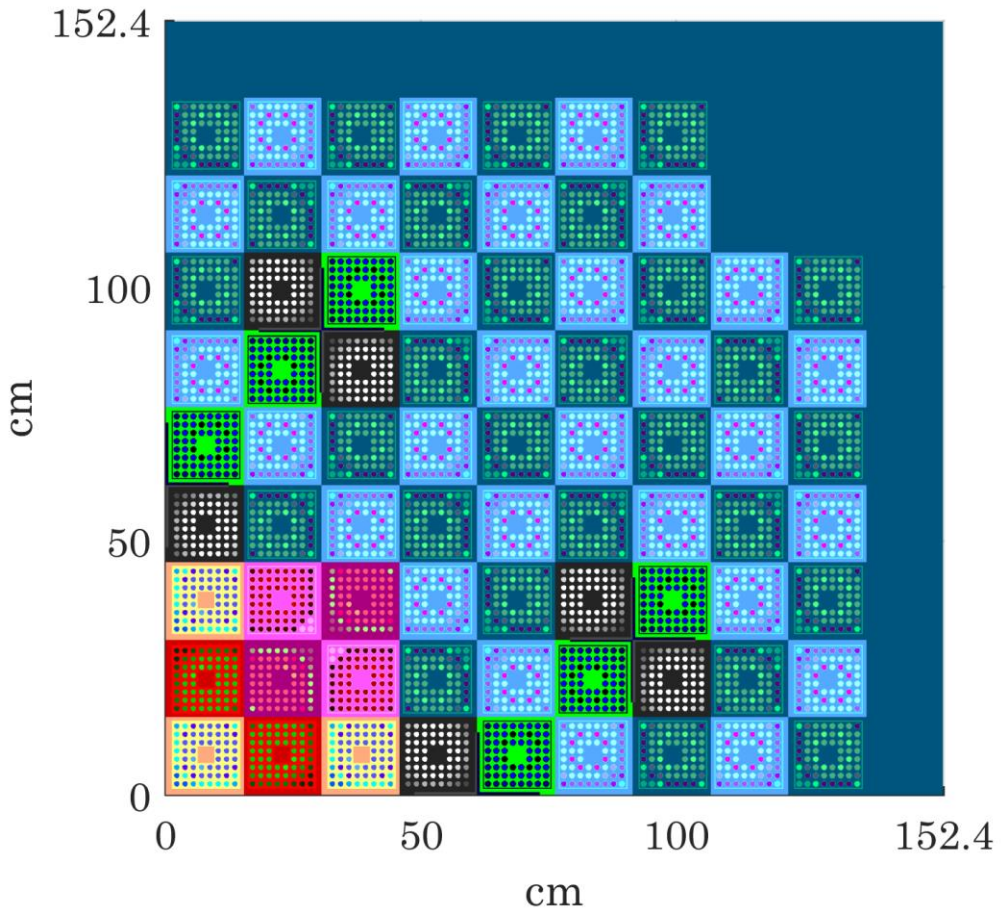

Figure 7. Quarter-core core layout of 2-D BWR core with Cartesian mesh. Each color represents a different material. Control blades are inserted between assemblies with green and black moderator

4.2.2 Reference solution for 2-D BWR benchmark

The accuracy of the 2-D S_N solver itself was verified for the C5G7 benchmark, and so those verifications are not repeated here. Any further references to the 2-D BWR benchmark or to its solution refer to the solution of the Cartesian mesh version of the benchmark problem. This benchmark has only 2-group cross sections, and so no group collapse is necessary in order to

plot the reference fluxes. As before, fine mesh flux is plotted (not pin fission). The reference eigenvalue of the 2-D BWR benchmark is 1.017495. This reference solution was calculated using level-symmetric S_8 quadrature. The reflector region is modelled with half mean free path spacing where appropriate, and thin regions in each assembly, such as the bundle cladding and the control blade cladding are modelled as one cell thick.

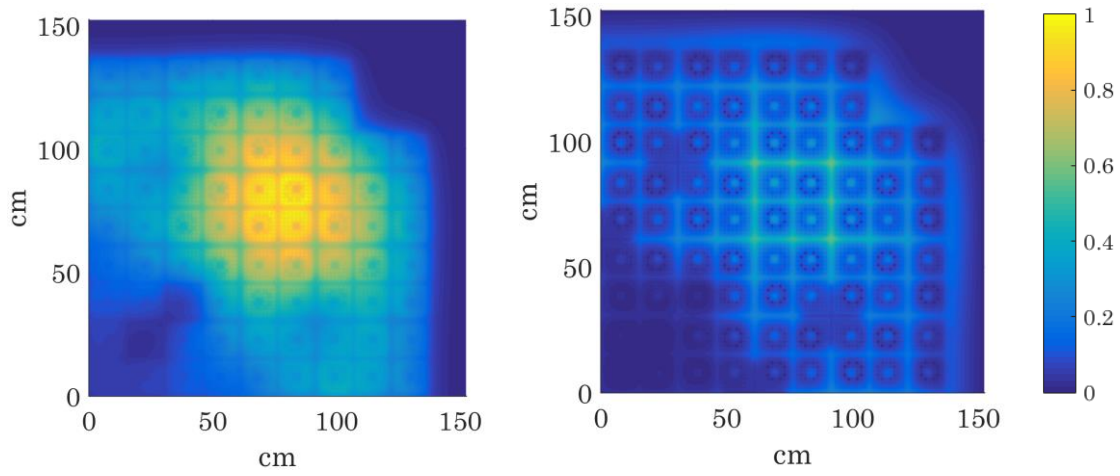

Figure 8. Reference scalar flux solution to 2-D BWR problem. Left: Fast group, right: thermal group. Flux presented in arbitrary units with maximum 1.

4.3 2-D S_N Transport Package

In order to develop proper testing grounds for the 2-D CSH and DTH methods, a neutron transport framework in which to implement the methods was first necessary. The 2-D transport tools used in order to implement the CSH and DTH methods needed to be highly adaptable and easy to access and modify in order to allow for implementation of the specifics of the CSH and

DTH methods and also in order to guarantee that any comparisons of computation time were fair. No transport packages were readily available which had the adaptability required of the CSH and DTH methods. Instead, a custom-built 2-D S_N code system was used.

This 2-D S_N code system is based up on a 'front-end' which is implemented in Python 2.7 and a 'back-end' written in compiled FORTRAN 90. The FORTRAN code has been compiled into a neutron transport Python library, which can easily be implemented into python code in order to run with any set of front-end code. This setup allows for an easy implementation of the CSH and DTH methods as all of the manipulations required by the methods can be implemented by making modifications to the Python front-end. The back-end transport routines are entirely serial in order to allow for fair comparisons of computation time, but they are otherwise robust and efficient transport solvers that can be applied to any discrete ordinates problems in two spatial dimensions.

It is the hope of the author that the highly robust code package developed in order to implement and test the CSH and DTH methods will outlive this thesis and form the backbone for future work. All of the code used in this transport package is well-documented and readily available for other researchers.

4.4 Figures of Merit

The CSH and DTH methods are evaluated for their flux accuracy, eigenvalue accuracy, and their computational efficiency. In this section, the quantities that will be used to summarize these accuracies will be described and justified. The most straightforward figure of merit will be the eigenvalue error. Eigenvalue errors will be reported as absolute differences between the calculated eigenvalue and the reference eigenvalue, in units of per cent mille (pcm) as

$$\Delta k = |k_{\text{calc}} - k_{\text{reference}}| \times 10^5 \text{ pcm}. \tag{41}$$

Three integrated quantities will be used to discuss flux errors. All three are based off of the spatially dependent relative error term, $e_g(\vec{r})$ which is defined

$$e_g(\vec{\text{r}}) = \frac{\left|\phi_g^{\text{calc}}(\vec{r}) - \phi_g^{\text{ref}}(\vec{r})\right|}{\phi_g^{\text{ref}}(\vec{r})} \times 100\%, \tag{42}$$

where $\phi_g^{\text{ref}}(\vec{r})$ is the reference scalar flux solution and $\phi_g^{calc}(\vec{r})$ is the core scalar flux solution from the CSH or DTH method. For the C5G7 benchmark, the group dependence of $e_g(\vec{r})$ is condensed into a 2-group relative error term as

$$e_c(\vec{r}) = \frac{\sum_{g\in c}\left|\phi_g^{\text{calc}}(\vec{r}) - \phi_g^{\text{ref}}(\vec{r})\right|}{\sum_{g\in c}\phi_{g,i,j}^{\text{ref}}(\vec{r})} \times 100\%, \tag{43}$$

Where c can be either the fast (f) groups or the thermal (th) groups, which are respectively groups 1-4 and groups 5-7 for the C5G7 benchmark. This group

collapse is foremost performed in order to allow for easier comparisons to previous CSH and DTH results in 1-D and also performed in order to condense the results for more concise presentation. In general, the group collapse does not largely affect the error profiles; the error profiles of the subgroups within a single collapsed group are similar in both shape and magnitude. In 1-D, a similar procedure was performed in order to collapse the error profiles from 47 groups to 2 groups without compromising the integrity of the results [1]. For the BWR benchmark case, the benchmark only specifies two groups, so the fast and thermal condensation is not necessary. The integrated flux errors, which will be used to evaluate the CSH and DTH methods, are defined as

$$\mathrm{AVG}_g = \frac{\int_{V_{\mathrm{core}}} \mathrm{d}\vec{r}\ \left|e_g(\vec{r})\right|}{\int_{V_{\mathrm{core}}} \mathrm{d}\vec{r}}, \tag{44}$$

$$\mathrm{MRE}_g = \frac{\int_{V_{\mathrm{core}}} \mathrm{d}\vec{r}\ \phi_g^{\mathrm{ref}}(\vec{r})\left|e_g(\vec{r})\right|}{\int_{V_{\mathrm{core}}} \mathrm{d}\vec{r}}, \tag{45}$$

and

$$\mathrm{MAX}_g = \max_{\vec{r}\in V_{\mathrm{core}}} \left|e_g(\vec{r})\right|. \tag{46}$$

These three metrics will be referred to as the *average error*, the *mean relative error*, and the *maximum error*, respectively. The index g refers to either the collapsed fast or thermal groups. The integrals in Equations (44), (45), and (46) are evaluated using the formula found in Equation (27). In every case, the

domains have been restricted to the core itself, neglecting the reflector regions. This is because recent work on the DTH method has shown that the spatially dependent relative flux error term $e_g(\vec{r})$ behaves poorly in the reflector regions, since if the reference solution becomes very small, the error can become arbitrarily large without necessarily being meaningful, making values of AVG and MAX useless. With the error domain restricted to core regions, AVG and MAX are useful tools for evaluating the accuracy of the method.

The CSH and DTH methods are fundamentally methods for improving computational performance of solution to the eigenvalue problem, and so the last figure of merit is the speedup factor, which is defined

$$\text{speedup} = \frac{T_{\text{ref}}}{T_{\text{calc}}}, \tag{47}$$

where T_{ref} and T_{calc} are the wall clock computational time of the reference solution and the calculation of interest. By definition, the reference solution itself will always have a speedup of exactly 1.0. Because the underlying transport code for the CSH and reference solutions are identical and written for this purpose, this metric is a fair evaluation of the calculation speedup provided by the method. Speedup is always calculated with T_{ref} and T_{calc} calculated using the same computer system, and all results from the following section are calculated on either a desktop Intel i7-3770 with a clock speed of

3.7 GHz or an Amazon EC2 system with an Intel Xeon E7-8880 v3 with a clock speed of 2.5 GHz.

CHAPTER 5. CSH RESULTS

The CSH method as defined in CHAPTER 3 has user-defined quantities for both homogeneous problem and re-homogenization stage convergence criteria, as well as user-defined quantities for homogeneous problem meshing, re-homogenization method, spatial basis function, and whether or not to apply progressively tightened convergence criteria and homogeneous mesh size. This chapter will study the effect of each of these user-defined quantities on the CSH solution to the 2-D C5G7 benchmark and the 2-D BWR benchmark. Each application of the CSH method will be evaluated for both its accuracy when compared to the reference solution and its speedup factor.

In order to more fairly compare against the reference calculation, all cases end with homogeneous problem convergence criteria equal to the criteria used to solve the reference solution, and whenever assembly fixed-source re-homogenization is used, the flux convergence criteria are chosen to be identical to the homogeneous problem flux convergence criteria that were used for that outer iteration. In previous implementations of the CSH method, spatial meshes were chosen based on 'reasonable' values, usually equal to roughly the same mesh size in mean free paths (mfp) as the largest mesh size in the heterogeneous problem. When mesh size for the homogeneous problem is chosen to be equal to the largest mesh of the heterogeneous problem in units

of mean free paths, this is considered the 'classic' meshing option. For the C5G7 benchmark, the 'classic' mesh is 0.64 mfp per homogeneous mesh. For the BWR benchmark problem, the 'classic' mesh is 0.99 mfp per homogeneous mesh.

5.1 Choice of Re-homogenization Convergence Criteria

In previous implementations of the CSH method, the re-homogenization convergence criteria as presented in Equations (15) and (16) were originally set as $\epsilon_k = 10^{-4}$ and $\epsilon_\phi = 10^{-3}$ somewhat arbitrarily, although in practice these criteria successfully led to a required number of re-homogenization iterations of about 4 or 5, and flux and eigenvalue errors tend to level out after 4 iterations, making those criteria a reasonable choice for 1-D problems. There is no reason to immediately suspect that these criteria will be reasonable in two spatial dimensions, nor if they will be reasonable for significantly more difficult problems, such as the 2-D BWR problem, which has a much higher variation of both mesh sizes and material properties than any of the 1-D benchmarks which have been tested in the past. The choice of re-homogenization iteration convergence criteria are re-investigated in this section.

As an initial investigation of the proper choice of re-homogenization convergence criteria, both benchmarks were solved with an S_2 angular order to a high number of iterations. For these two calculations, all other CSH

parameters were left as close as possible to the parameters of older 1-D implementations, meaning that homogeneous mesh sizes were chosen as the classic mesh size, and that progressive convergence criteria were used for the first four re-homogenization iterations, reducing both ε_k and ε_ϕ by a factor of 10 each re-homogenization iteration down to values of 6×10^{-6} in flux and 7×10^{-8} in eigenvalue for the fourth iteration, where they stayed for all remaining re-homogenization iterations. These calculations employed 4 core sweeps for each re-homogenization calculation, a number which was considered sufficient when this parameter was evaluated for the DTH method [3], which will also be demonstrated in Section 5.2.2. Additionally, these two calculations were performed using 1st order integral-conserving spatial basis functions both for angular flux and auxiliary source.

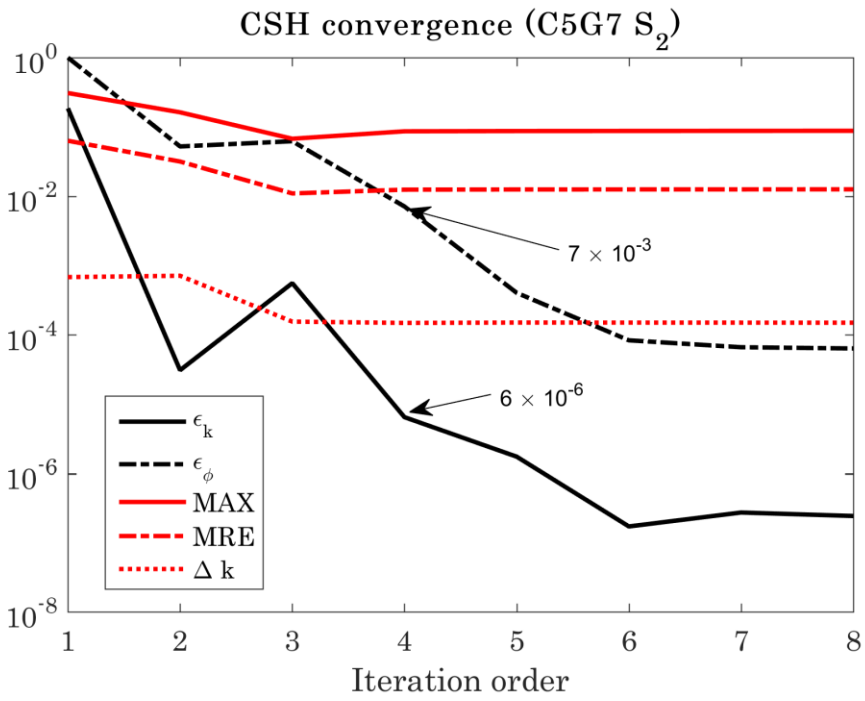

Figure 9. Re-homogenization convergence for the C5G7 core with S_2. Red lines indicate error compared to the reference at each step, and black lines indicate the current value of the within-calculation criteria.

The results of this calculation for the C5G7 core can be found in Figure 9. This plot presents the actual errors of each re-homogenization iteration as compared to the within-calculation convergence metrics. As can be seen in the graphic, the actual errors level out after four re-homogenizations, meaning that despite the fact that both ϵ_k and ϵ_ϕ continue to decline, there is no further accuracy to be gained by continuing the calculation past this point. The values of ϵ_k and ϵ_ϕ at the fourth iteration are shown to be 6×10^{-6} and 7×10^{-3} respectively, indicating that previous CSH implementations have likely been entirely controlled by the criterion for ϵ_ϕ and that the previously used value

for ϵ_ϕ of 1×10^{-3} is too tight for the 2-D implementation of the CSH method. The AVG metric is not presented in this plot, as it is overlapping with the MRE line in this calculation. In general, MAX and MRE are group-dependent values, however in these plots only the maximum group contribution is depicted. For both Figure 9 and Figure 10, these refer to the fast-group MAX and MRE metrics in all but the first several iterations.

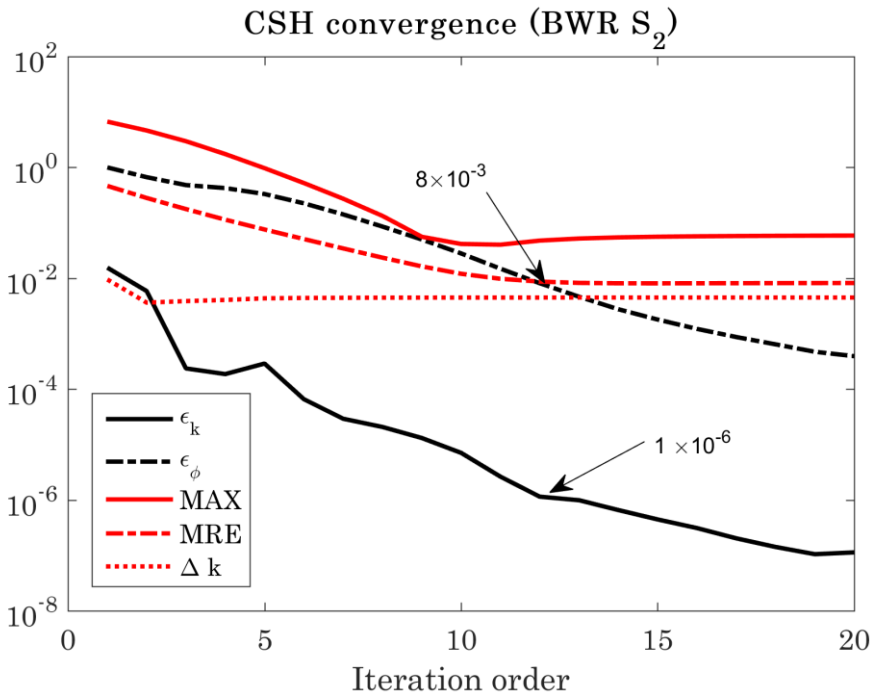

Figure 10. Re-homogenization convergence for the BWR core with S_2. Red lines indicate error compared to the reference at each step, and black lines indicate the current value of the within-calculation criteria.

When the calculation is repeated for the BWR benchmark, it is immediately clear that the previously observed effect that the CSH method generally is finished after 4-5 re-homogenizations does not hold for the BWR

benchmark. Clearly, for this core the auxiliary source is slower to converge. While the number of outer iterations required of the BWR benchmark is significantly higher, it is relieving to see that the values of ϵ_k and ϵ_ϕ are similar to the C5G7 case when the problem is converged. As indicated in Figure 10, the flux and eigenvalue errors of the CSH method visibly level off after 11 re-homogenization iterations, with values of ϵ_k and ϵ_ϕ equal to 1×10^{-6} and 8×10^{-3} respectively. Using the information of these two calculations, re-homogenization criteria of $\epsilon_\phi = 5 \times 10^{-3}$ and $\epsilon_k = 1 \times 10^{-5}$ were chosen for all subsequent calculations. In practice, every calculation was controlled by the criterion for ϵ_ϕ.

5.2 Parameter Variations

In this section, the various user-defined parameters of the CSH method are evaluated for their effect on the speed and accuracy of the CSH method. In each of the following subsections, results for the C5G7 core and the BWR core will be presented in tables where a single CSH parameter varies while the rest remain fixed. Comparison will be drawn between the behavior of the CSH method for each of the two benchmark problems, and optimal user-defined parameters will be suggested.

5.2.1 Spatial basis functions

In Section 3.3.1, three spatial basis functions for use with the CSH method were suggested. In this section, CSH results with each of the three spatial basis functions will be presented for both benchmark problems. Results for the C5G7 benchmark can be found in Table 5. These calculations were performed using four core sweeps for each re-homogenization step, with a classic (0.64 mfp per homogeneous mesh) homogeneous mesh, and progressive homogeneous convergence criteria for the first four iterations.

Table 5. Effect of spatial basis function choice on CSH solution to C5G7 benchmark. Calculation performed with four core sweeps per re-homogenization and classic homogeneous mesh.

Basis function	Iters	Speedup	Δk (pcm)	MRE (%)		AVG (%)		MAX (%)	
				F	Th.	F	Th.	F	Th.
Linear	5	2.4	1851.6	0.7	0.6	0.7	0.6	3.5	7.1
0th order	5	2.5	25.7	0.4	0.4	0.4	0.5	1.5	3.5
1st order	5	2.5	24.8	0.4	0.4	0.4	0.4	1.5	3.5

From Table 5, the poor performance of the linear B-spline basis functions in these 2-D problems with integral-conserving mesh transfers is evident. Without conserving integrals between each solution, the linear B-spline basis functions do not maintain the reaction rate of the auxiliary source, which can upset the neutron balance in the core and lead to significant errors in the eigenvalue calculation. In other observations, as expected from the brief analysis of the re-homogenization convergence behavior of the CSH method,

each calculation required five re-homogenizations to converge. For the two integral-conserving spatial basis functions, eigenvalue error is about 25 pcm, which is about a factor of ten higher than when the CSH method has been used for 1-D problems. For each of the integral-conserving basis functions, average and mean-relative flux errors are about 0.4% for both the fast and thermal energy ranges.

Table 6. Effect of spatial basis function choice on CSH solution to BWR benchmark. Calculation performed with four core sweeps per re-homogenization and classic homogeneous mesh.

Basis function	Iters	Speedup	Δk (pcm)	MRE (%)		AVG (%)		MAX (%)	
				F	Th.	F	Th.	F	Th.
Linear	15	5.3	1705.0	4.1	4.9	7.2	8.7	42.2	47.6
0th order	11	5.7	266.2	0.3	0.9	0.5	0.9	1.9	2.3
1st order	11	7.3	275.7	0.3	0.8	0.5	0.8	1.9	2.3

Table 6 contains the same results as the previous table, except applied to the BWR benchmark. The difference between the integral-conserving basis functions and the linear B-splines in accuracy is clear. In the case of the BWR benchmark, the difference in flux accuracy is stark. The linear B-spline basis functions do not result in a solution to the benchmark that can be called accurate at all, with an average flux error of over 7%. Another observation to be made from Table 6 is that the linear B-spline solutions took more re-

homogenization iterations to converge, largely due to the very high flux errors leading to trouble with the re-homogenization convergence criteria.

It should be noted that the overall eigenvalue errors for the BWR benchmark are about one order of magnitude higher than when the C5G7 benchmark is solved with CSH for nearly all cases. This is largely due to the significantly increased complexity of the 2-D BWR benchmark problem compared to the 2-D C5G7 problem, specifically the resolved cladding around each fuel pin. The extremely small mesh sizes of the cladding, on the order of 0.02 mfp are difficult or impossible to resolve in any set of spatial basis functions used for the CSH method, meaning that the effect of pin cladding is not well accounted for in the eigenvalue solution.

5.2.2 Re-homogenization process

In 1-D, the CSH method has been implemented with several different methods for re-homogenization. Initial CSH implementations [1] took the simple approach of re-using the methods employed for the single-assembly calculations but with updated boundary conditions. The solution process begins by assuming that single-assembly calculations with approximated specular reflective boundary conditions have already been performed. At each re-homogenization step, the results from the homogeneous core calculation are expanded and the calculations are repeated with updated fixed boundary

conditions and eigenvalue. Later implementations of the CSH and DTH methods in one spatial dimension gained efficiency by replacing the assembly fixed-source re-homogenization calculations with full-core transport sweeps, which have the advantage of being exceptionally quick, as long as the total number of sweeps is kept low. In the 2-D implementation of CSH, assuming an accurate spatial expansion function for angular flux, the primary function of these core sweeps is to correct for local errors introduced by the mesh transfer process, meaning that the number of core sweeps can be kept low without significantly impacting accuracy.

Table 7. Effect of different re-homogenization methods for the C5G7 core. AFS refers to assembly fixed-source re-homogenization.

# Core sweeps	Iters.	Speedup	Δk (pcm)	MRE (%)		AVG (%)		MAX (%)	
				F.	Th.	F.	Th.	F.	Th.
1	5	2.6	21.5	0.7	0.2	0.6	0.3	2.0	3.9
2	5	2.7	23.6	0.5	0.2	0.5	0.3	1.8	3.7
4	5	2.5	24.8	0.4	0.4	0.4	0.4	1.5	3.5
7	5	2.3	23.6	0.4	0.4	0.3	0.4	1.3	3.2
10	4	2.7	22.4	0.3	0.3	0.3	0.3	1.1	2.9
20	4	2.2	22.4	0.3	0.3	0.2	0.3	0.9	2.4
AFS	4	1.6	24.1	0.2	0.2	0.3	0.3	2.7	4.3

Table 7 tabulates the effect that different numbers of core sweeps per re-homogenization can have on the accuracy of the solution for the C5G7 benchmark. All of these calculations were performed using 1st order integral-conserving spatial basis functions and a 'classic' homogeneous problem mesh. For this benchmark, it can be seen that when core sweep re-homogenization is

used, the number of core sweeps per re-homogenization does not have an extremely large effect on the accuracy of the solution. When the number of core sweeps is low, one extra re-homogenization iteration is needed in order to meet the re-homogenization convergence criteria, which balances out the speedup effect, which would otherwise be expected to monotonically decrease as the number of core sweeps increases. Overall MRE and AVG flux errors as well as eigenvalue error do not change significantly with the number of core sweeps after about four sweeps per iteration, although the MAX flux error decreases as the iteration order increases. This suggests that for this benchmark, the maximum error likely occurs in a fairly isolated portion of the core, since MRE and AVG are similar.

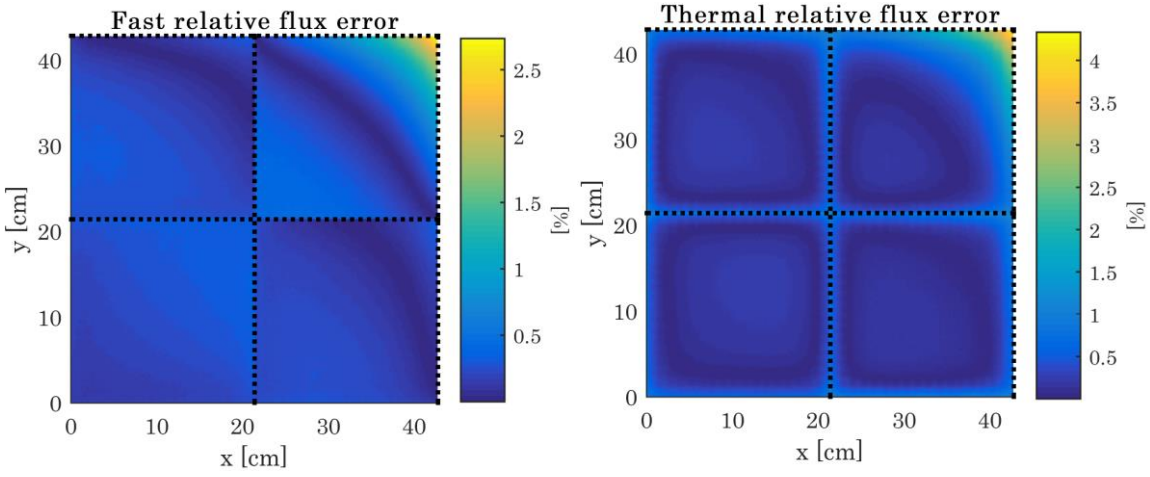

Figure 11.Within-core fast and thermal relative error profiles for the C5G7 core with AFS re-homogenization.

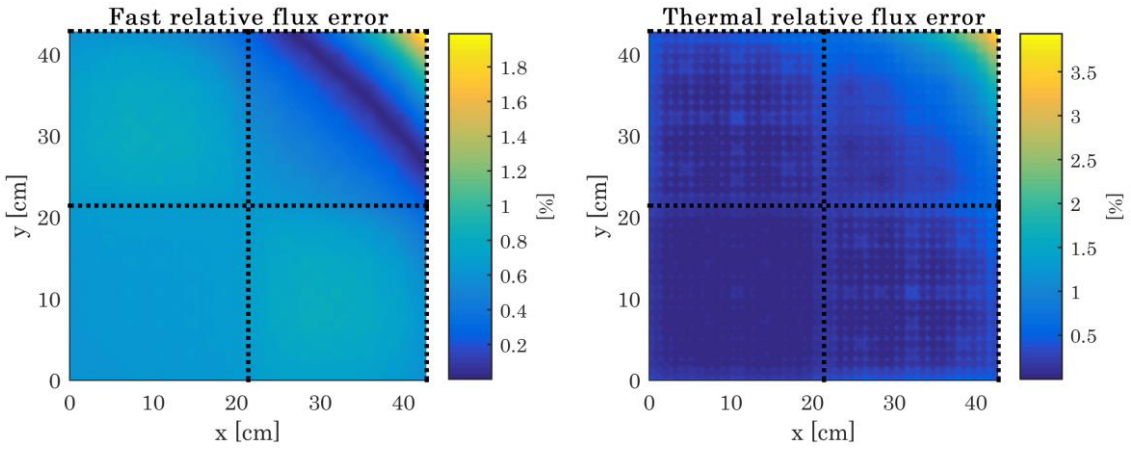

Figure 12. Within-core fast and thermal relative error profiles for the C5G7 core with one core sweep re-homogenization

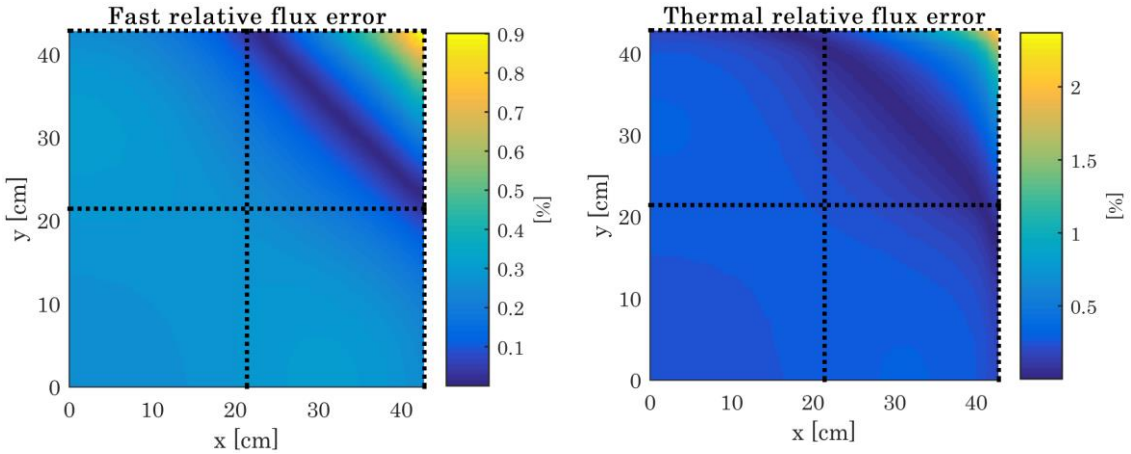

Figure 13. Within-core fast and thermal relative error profiles for the C5G7 core with 20 core sweep re-homogenization

Figures 11, 12, and 13 depict relative error profiles for the three extremes contained in Table 7. Note that because each plot has its own color scale, they should not be compared against each other in this context. Instead, these plots are provided in order to examine the behavior of the CSH method with different re-homogenization methods. In the thermal spectrum of Figure 11, error can be seen to peak near assembly boundaries, which is a behavior that was seen frequently in 1-D implementations of the CSH method. Figure 12 indicates that the use of only one core sweep for re-homogenization is insufficient. Figure 12 shows a slight checkerboard pattern within each assembly, especially visible in the thermal error profile. This indicates that the local errors have not been sufficiently reduced when only a single core sweep is employed for re-homogenization. When 20 core sweeps are used for each re-homogenization step, it becomes clear that the error profile is dominated by core-level effects, since no local effects are visible, nor is there any peaking near assembly boundaries.

Table 8. Effect of different re-homogenization methods for the BWR core. AFS refers to assembly fixed-source re-homogenization.

# Core sweeps	Iters.	Speedup	Δk (pcm)	MRE (%)		AVG (%)		MAX (%)	
				F.	Th.	F.	Th.	F.	Th.
1	11	6.5	243.6	0.4	1.3	0.5	1.3	1.9	5.0
2	11	6.4	263.0	0.4	1.1	0.5	1.1	1.9	3.6
4	11	7.3	275.7	0.3	0.8	0.5	0.8	1.9	2.3
7	11	5.2	274.8	0.3	0.5	0.4	0.5	1.9	1.6
10	11	5.0	271.1	0.2	0.4	0.4	0.4	1.9	1.6
20	11	4.4	265.8	0.2	0.3	0.3	0.3	1.7	1.5

AFS	11	4.5	271.0	0.3	0.7	0.4	0.5	1.6	2.5

These calculations were repeated for the BWR benchmark, again using a 'classic' homogeneous mesh and 1st order integral-conserving spatial basis functions, and the results are presented in Table 8. These results show a similar trend to that of the C5G7 benchmark. For the BWR benchmark, no number of core sweeps required any extra re-homogenization iterations, and the trend of decreasing speedup with increasing core sweeps is clearly demonstrated after four sweeps per iteration. For 1 and 2 core sweeps per re-homogenization, the speedup is lower; this is due to an increased number of inner iterations to solve the homogeneous problem for each re-homogenization iteration when the re-homogenization is not performed sufficiently. This indicates slower convergence of the auxiliary source term when the re-homogenization is not as accurate.

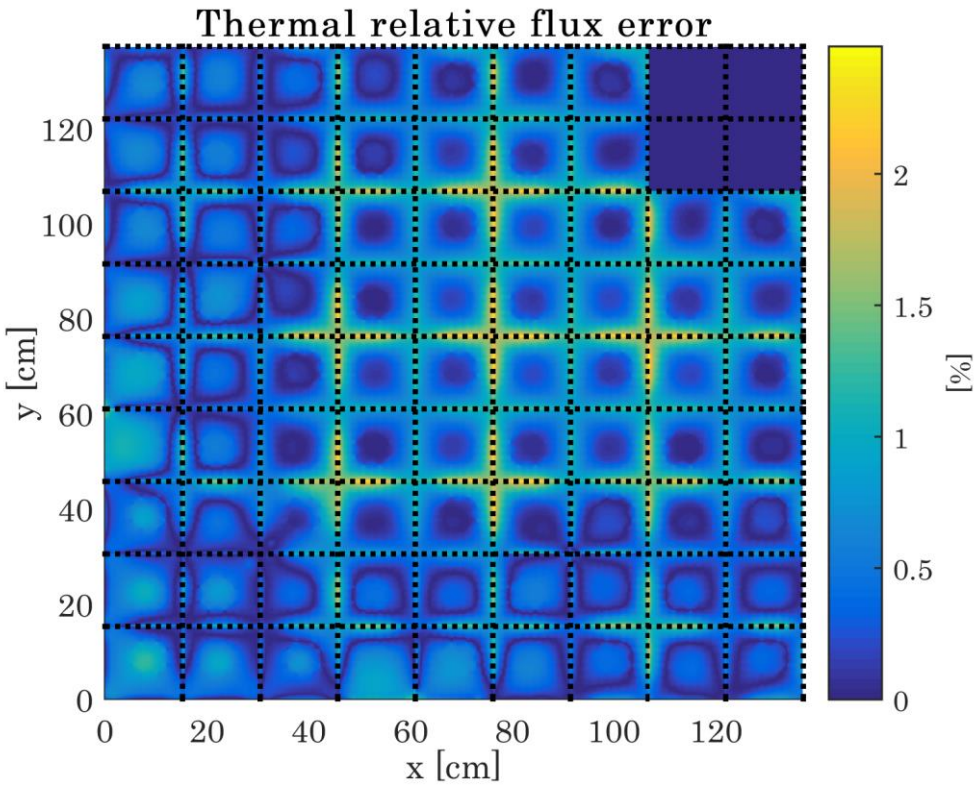

Figure 14. Within-core thermal relative thermal error profile for the BWR core with assembly fixed-source re-homogenization

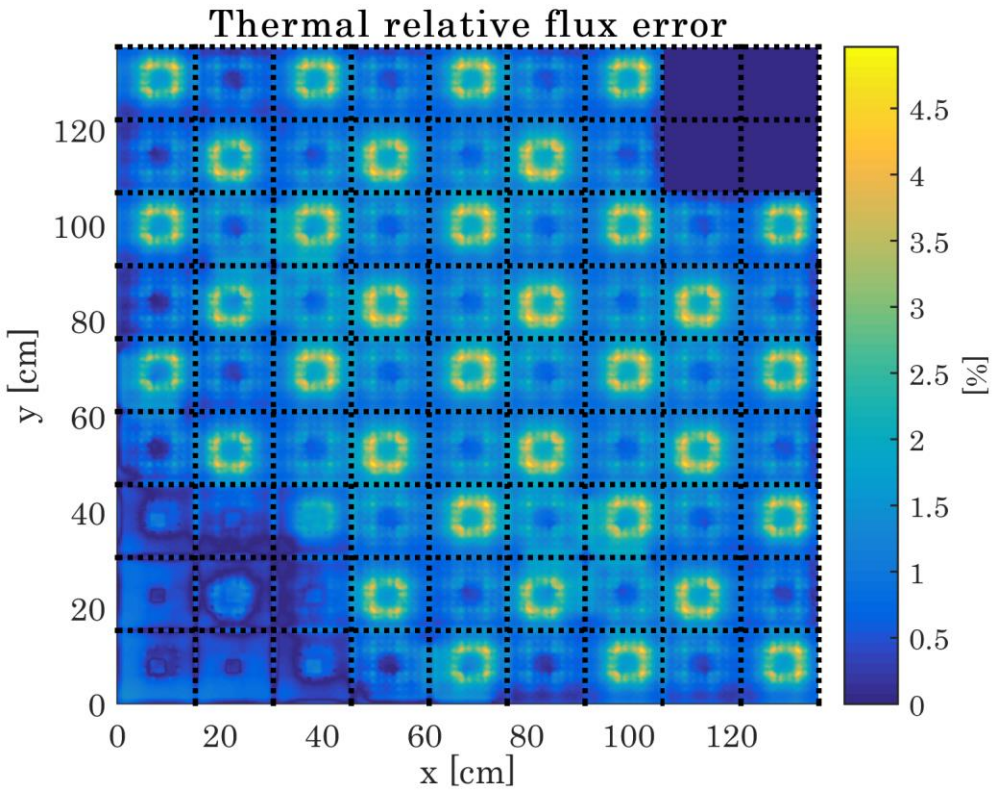

Figure 15. Within-core thermal relative thermal error profile for the BWR core with one core sweep re-homogenization

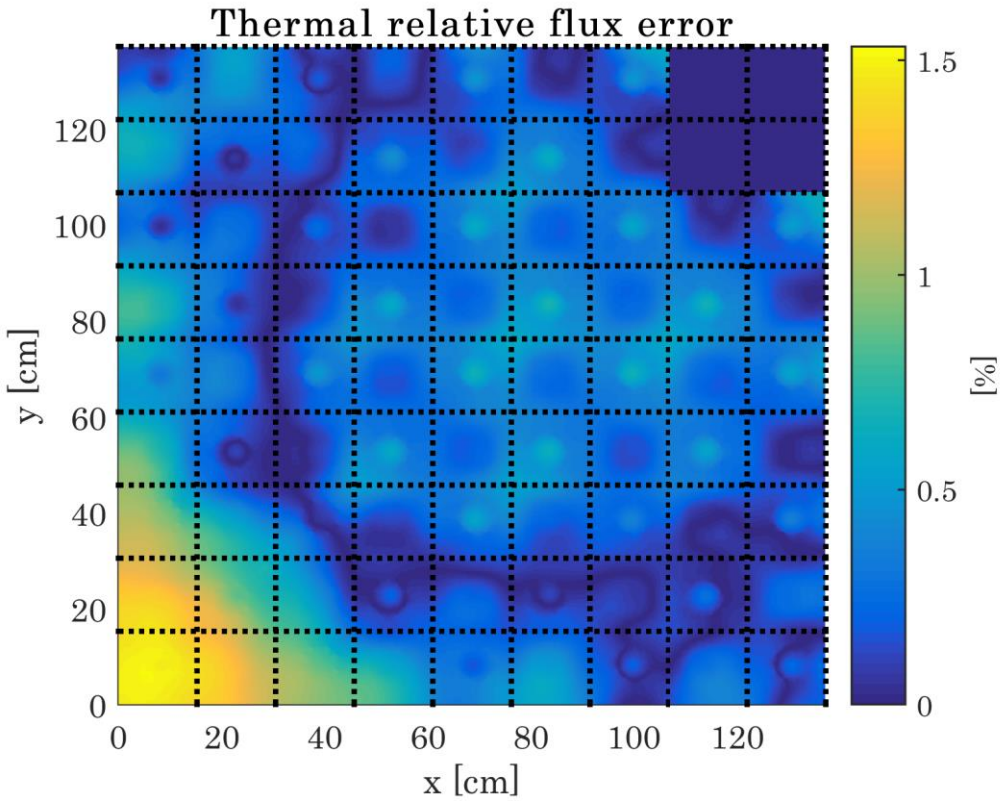

Figure 16. Within-core thermal relative thermal error profile for the BWR core with 20 core sweep re-homogenization

Figures 14, 15, and 16 depict just the relative thermal error profile for the BWR benchmark. The fast flux profiles have been omitted for succinctness. As before, each plot has its own color scale, so comparisons should not be drawn between each of the three. Like with the C5G7 case, the error profile in Figure 14 generated using assembly fixed-source re-homogenization displays the error peaking near assembly interfaces, which is common of previous CSH implementations. When only one core sweep is used per re-homogenization, the error profile is dominated by local effects, which have not been properly

reduced by an insufficient re-homogenization, and when 20 core sweeps are used per re-homogenization iteration, as seen in Figure 16, the error profile is dominated by core effects, indicating that the re-homogenization procedure is sufficiently reducing local error effects.

5.2.3 Homogeneous problem meshing

This 2-D implementation of the CSH method is the first to include homogeneous mesh size as a user-defined parameter of the method, making it prudent to investigate the effect of homogeneous mesh size on calculations.

Table 9. Effect of homogeneous problem mesh size for CSH on the C5G7 benchmark solution.

Mesh size (mfp)	Iters.	Speedup	Δk (pcm)	MRE (%)		AVG (%)		MAX (%)	
				F.	Th.	F.	Th.	F.	Th.
0.5	5	1.7	19.2	0.4	0.4	0.4	0.4	1.5	3.5
1	5	4.2	40.5	0.4	0.4	0.4	0.5	1.4	3.3
2	5	7.2	162.8	0.4	0.4	0.5	0.6	0.9	2.7
4	4	11.2	56.0	0.5	1.2	0.5	1.2	1.6	4.3
8	4	12.0	86.4	0.9	1.4	0.9	1.3	3.1	5.0
16	5	9.4	92.8	1.0	1.9	1.2	2.1	5.0	8.7
Exact	5	0.8	0.1	0.01	0.02	0.01	0.02	0.07	0.12

Table 9 summarizes the effect of homogeneous mesh size on the C5G7 benchmark solution for the CSH method. These calculations were performed using 4 core sweeps per re-homogenization and with the 1st order integral-conserving spatial basis functions. In previous tables, the C5G7 benchmark was solved using the 'classic' homogeneous mesh, which is roughly 0.6 mfp per

cell. Additionally, for comparison's sake, the calculation has also been performed using an exact heterogeneous mesh overlaid on top of the homogeneous problem. In this case, the spatial expansion functions become redundant, as the auxiliary source and angular flux can transfer between meshes without any approximation, effectively eliminating all of the error terms from Equation (40) except for e_{aux}. Recall from CHAPTER 4 and from APPENDIX B that the reference solutions are converged to roughly 0.5% maximum flux error and 0.5 pcm against the discretization error itself. As expected, when using an exact mesh overlay, the accuracy of the CSH method is effectively exact, as the calculated errors are within range of the error of the reference solution compared to discretization accuracy. However, as expected, there is no advantage in speed to calculating the solution using an exact mesh overlay.

Table 9 indicates that, to no surprise, speedup can be drastically increased by increasing the homogeneous problem mesh size, and while this results in a loss of accuracy, the loss of accuracy is not catastrophic for the C5G7 problem for fairly significant overall mesh sizes. These results indicate that previous implementations of the CSH method may have been using homogeneous mesh sizes that were significantly tighter than necessary, which may have impacted previous speedup calculations.

Table 10. Effect of homogeneous problem mesh size for CSH on the BWR benchmark solution.

Mesh size (mfp)	Iters.	Speedup	Δk (pcm)	MRE (%)		AVG (%)		MAX (%)	
				F.	Th.	F.	Th.	F.	Th.
0.5	11	1.2	322.5	0.3	0.4	0.4	0.5	2.1	2.1
1	9	6.2	355.7	0.3	0.9	0.5	0.9	2.3	3.0
2	8	19.9	336.6	0.6	1.8	0.7	1.8	2.2	5.4
4	8	33.9	731.3	1.5	5.4	1.6	5.2	7.5	17.5
8	6	58.3	649.5	3.1	9.4	3.5	8.6	13.3	28.2
16	6	61.7	1253.9	5.6	13.0	6.2	13.2	21.1	49.3
Exact	**14**	0.6	0.2	0.5	0.6	0.6	0.6	1.7	1.9

The calculations were repeated for the BWR benchmark problem, and the results are reported in Table 10. As before, these calculations were performed using 4 core sweeps per re-homogenization with 1st order integral-conserving spatial basis functions. Like for the C5G7 case, when an exact mesh overlay is used, the accuracy is effectively perfect, although speedup indicates that this is not a beneficial practice. Somewhat paradoxically, the number of re-homogenizations required appears to decrease as the homogeneous mesh size increases, however this is reasonable in the context that much of the detail of the auxiliary source term is being lost by restricting it to a very coarse mesh, making the convergence faster, since local details are washing out before every re-homogenization. Further, the results past about 8 mfp per homogeneous mesh are not well-behaved, as the accuracy of the CSH method begins to fall off catastrophically for the BWR benchmark when the mesh size is increased too high. For the BWR benchmark, both the speedup and the accuracy follow

the trends that should be expected of changing the homogeneous mesh, with accuracy falling off precipitously for very high mesh sizes.

Additionally, by applying the heterogeneous mesh overlay to the problem, the issue of the spatial basis functions being unable to properly account for the extremely small cladding meshes disappears, and the eigenvalue convergence is significantly better-behaved.

5.3 CSH with Progressively Tightened Convergence and Meshing

As discussed in the theory section, by employing a progressively tightened mesh size in addition to progressively tightened homogeneous problem convergence criteria, the CSH method can reach its optimal convergence. The optimal progression of mesh size is not a given, however, and must first be decided upon. In this section, three possible progressive mesh sizing schemes will be implemented, and the results of the CSH calculation when those progressive mesh schemes are applied will be discussed. Based on the results in Table 9 and Table 10, all three potential progressive mesh sizing schemes begin with a mesh size of four times the 'classic' mesh size. For the C5G7 benchmark, this mesh size is about 2.6 mfp per mesh, and for the BWR benchmark, this results in about 4 mfp per mesh. From there, the mesh size is decreased linearly over five re-homogenization iterations to its final value. For these calculations, building upon the parameter searches of the previous three

sections, the following user-defined parameters were employed. The spatial basis function was chosen to be 1st order integral-conserving spatial basis functions, and the re-homogenization method was chosen to be 4 core sweeps per re-homogenization step.

Table 11. Results of three separate progressive meshing schemes applied to the CSH method for the C5G7 benchmark.

Final mesh	Iters	Speedup	Δk (pcm)	MRE (%)		AVG (%)		MAX (%)	
				F	Th.	F	Th.	F	Th.
Small (het. avg.)	6	2.4	12.0	0.4	0.5	0.4	0.4	1.6	3.7
Classic (het. max)	6	2.5	24.7	0.4	0.4	0.4	0.5	1.5	3.5
Exact	**6**	**1.6**	**0.34**	**0.04**	**0.05**	**0.04**	**0.04**	**0.2**	**0.4**

In Table 11 the three progressive mesh schemes were tested. In the first, the final mesh size was chosen to be the average mesh size (in mean free paths) of the heterogeneous problem. The average mesh size of the heterogeneous C5G7 problem is 0.40 mfp per mesh, which is roughly 37% smaller than the 'classic' mesh size. This was chosen in order to attempt to achieve similar speedup factors as compared to calculations that were performed using a constant `classic' mesh for the homogeneous problem. The second method employed the `classic' mesh as the final mesh of the progressive scheme, with the idea being to achieve similar accuracy when compared to calculations performed using a constant 'classic' mesh, and the final scheme employs an exact homogeneous mesh at the last step.

The last scheme follows the same progressive meshing as the first scheme, except that after 4 re-homogenization steps, all future homogeneous meshes are instead overlaid by the heterogeneous mesh. The idea of this scheme is to cause the CSH method to act as an acceleration method rather than an approximation method, by arriving at the exact solution, but using the less accurate form of the CSH method to quickly calculate the auxiliary source on the way to that solution.

The first scheme was successful in its goal for the C5G7 benchmark. The speedup when using a smaller mesh size is very similar to calculations performed with a constant mesh size equal to the 'classic' homogeneous meshing, but the eigenvalue error is noticeably lower. However, the flux errors are not much different, when compared against a similar calculation, such as the third row of Table 7, which is identical in all but the homogeneous meshing.

The second scheme, progression to the 'classic' mesh was intended to increase speedup with no loss of accuracy over the calculations performed with classic meshing. The second scheme, while offering no drawback compared to non-progressive CSH schemes, did not succeed in its goal. The use of a progressively tightened homogeneous mesh increased the required number of re-homogenizations compared to a fixed homogeneous mesh size, eliminating any additional speedup that might have been gained over more traditional CSH calculations.

Perhaps most interesting is the third progressive scheme, where the CSH method has been employed as an acceleration method for an exact-mesh CSH final step. In this scheme, the CSH method obtained accuracy which is functionally identical to the reference solution, in that the eigenvalue error is less than 0.5 pcm, and the maximum flux error is less than 0.5%. Additionally, by using the CSH method to calculate the auxiliary source term, only two additional iterations were required after using the exact mesh overlay in order to fully converge the auxiliary source term, which effected significant speedup over the reference calculation. This speedup value is as high as any of the speedup values calculated by the most advanced implementations of the CSH method in 1-D with even higher accuracy, indicating that this is a promising application of the CSH method.

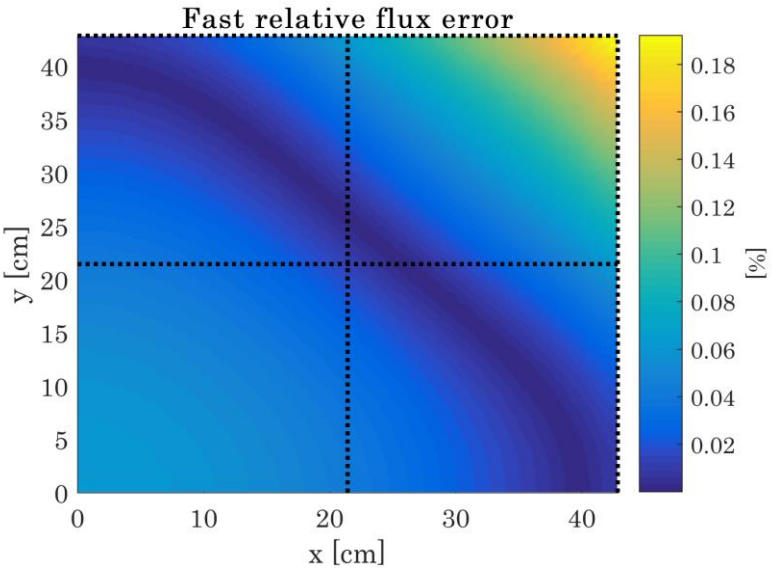

Figure 17. Fast spectrum relative flux error profile for C5G7 when solved with a progressive mesh ending with an exact heterogeneous mesh overlay.

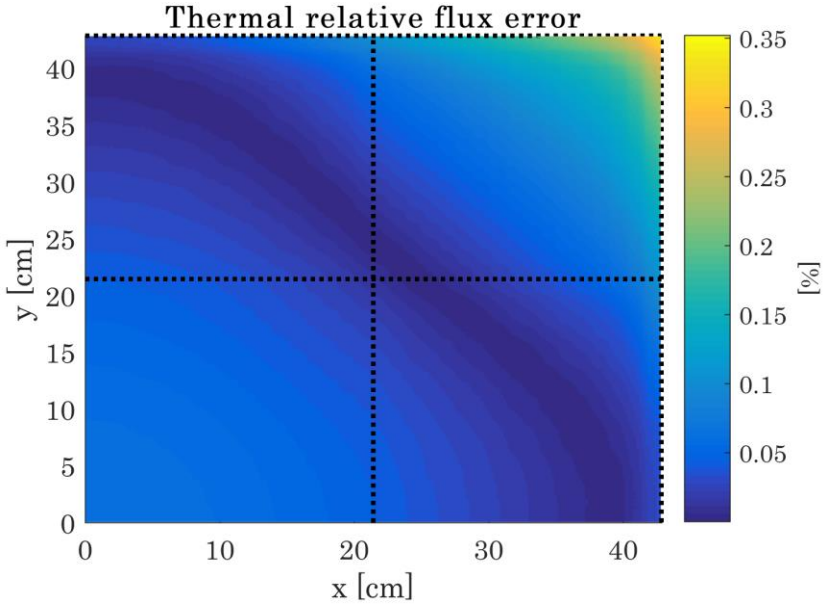

Figure 18. Thermal spectrum relative flux error profile for C5G7 when solved with a progressive mesh ending with an exact heterogeneous mesh overlay.

Figure 17 and Figure 18 depict the relative flux error profile when the CSH method is applied with a progressive mesh ending with a direct overlay of the heterogeneous mesh. In both profiles, the flux error is entirely dominated by core effects in the form of a slight diagonal tilt. Both cases have very small errors, even at their peak, with flux profiles nearly perfectly matching the reference flux solution.

Table 12. Results of three separate progressive meshing schemes applied to the CSH method for the BWR benchmark.

Final mesh	Iters	Speedup	Δk (pcm)	MRE (%)		AVG (%)		MAX (%)	
				F	Th.	F	Th.	F	Th.
Small (het. avg.)	9	2.3	304.7	0.3	0.4	0.4	0.5	2.0	2.3
Classic (het. max)	9	9.7	275.7	0.3	0.8	0.5	0.8	1.9	2.3
Exact	**11**	**1.9**	**0.2**	**0.4**	**0.4**	**0.4**	**0.5**	**1.3**	**1.4**

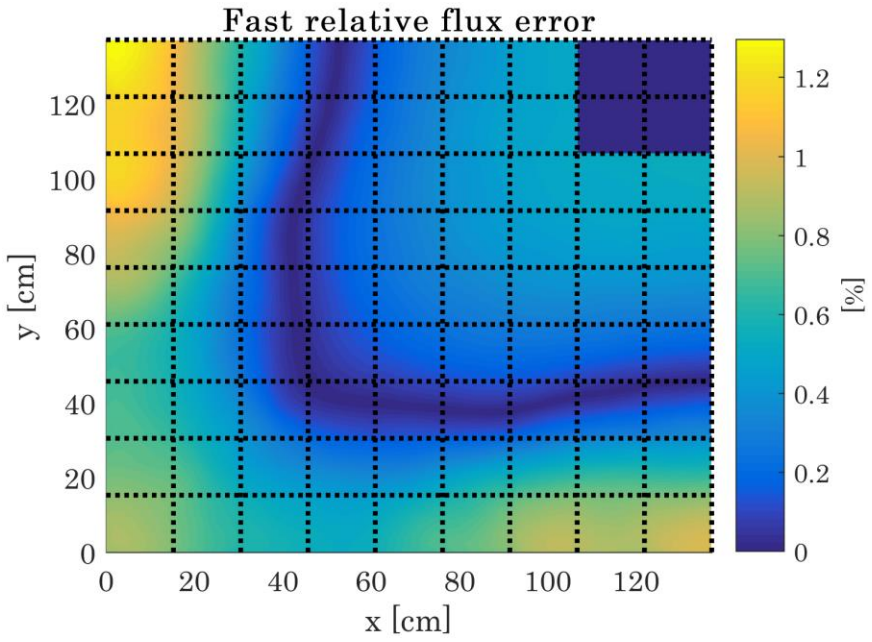

Figure 19. Fast spectrum relative flux error profile for BWR benchmark when solved with a progressive mesh ending with an exact heterogeneous mesh overlay.

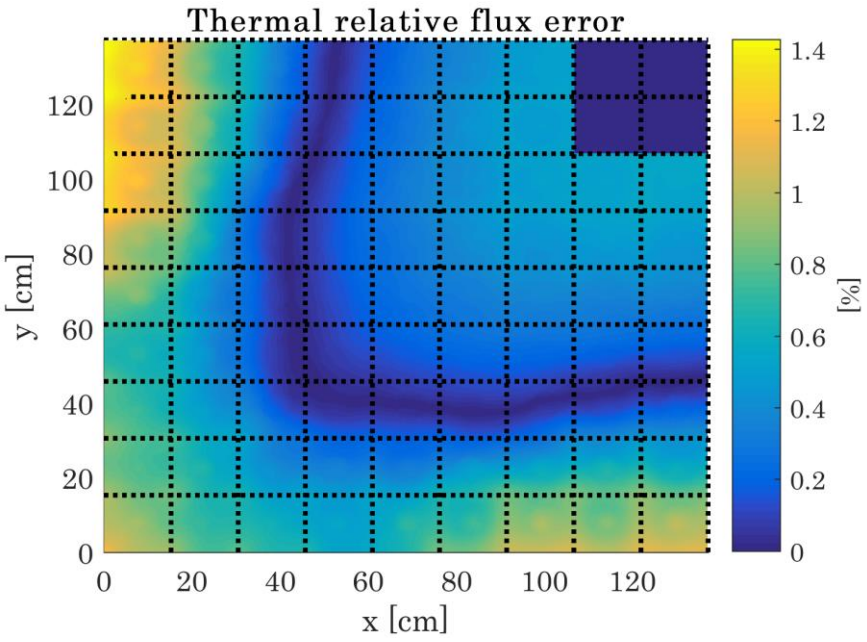

Figure 20. Thermal spectrum relative flux error profile for BWR benchmark when solved with a progressive mesh ending with an exact heterogeneous mesh overlay.

Table 12, Figure 19, and Figure 20 repeat the tests of the three progressive meshing schemes on the BWR benchmark problem. As before, these tests show that applying an exact heterogeneous mesh overlay on the homogeneous problem at the end of the progressive meshing can lead to nearly perfect solutions of the flux profile and eigenvalue. In the case of the BWR problem, the maximum flux error is not within the allowance for the reference solution, but they are of the same order, and examination of Figure 19, and Figure 20 indicates that the maximum only occurs in spatially small locations. Like the calculations that were applied to the C5G7 benchmark problem, progressing to an exact heterogeneous mesh led to speedups which are consistent with

previous 1-D implementations of the CSH method, just a little under two times. In this case, the use of a progressive homogeneous mesh for the case which ended with a 'classic' mesh did not require extra re-homogenizations, and so that calculation has about 50% more speedup than the calculation performed with a constant 'classic' meshing at each re-homogenization step.

These results indicate that progressive mesh tightening is a very effective tool with the CSH method, and that the CSH method can be reasonably tuned for either speed or accuracy by making adjustments to the mesh and progressive meshing.

CHAPTER 6. DTH RESULTS

In this chapter, results of diffusion-transport hybrid (DTH) method calculations performed on the two benchmark problems described in CHAPTER 4 will be presented. Like the CSH method, the DTH method has a number of user-defined parameters. A suite of results will be presented in this chapter in order to build a case for certain choices of user-defined parameters for the DTH method, as well as to discuss the advantages and disadvantages of various combinations of parameters.

Previous work in 1-D showed that the same 1-D convergence criteria used in CSH were appropriate for DTH [2] [3], so all DTH calculations have been performed using the same re-homogenization convergence criteria as described in CHAPTER 5. Inner calculations are performed with the same convergence criteria as used by the transport inner calculations of the CSH calculations, which is a progressive tightening of the convergence criteria over the first four re-homogenization iterations by factors of 10 until reaching the values of convergence criteria of 6×10^{-6} for flux and 7×10^{-8} for eigenvalue. Please see APPENDIX B for a further description of these values. The scalar flux convergence is said to be converged when

$$\frac{\left|\left|\phi_{i,j,g}^{\ell-1}-\phi_{i,j,g}^{\ell}\right|\right|_{\infty}}{\left|\left|\phi_{i,j,g}^{\ell}\right|\right|_{\infty}} \leq \epsilon_{\psi} = 6\times10^{-6}. \tag{48}$$

Please see APPENDIX B for a more complete review of the choice of convergence criteria for the homogeneous problem.

6.1 DTH with Progressively Tightened Convergence and Meshing

The DTH method was tested for a suite of three different progressive mesh tightening schemes, the same three schemes used with the CSH method. This section is presented before the other parameter variations in order to test what happens when all of the chosen CSH parameters are applied directly to DTH calculations with no other modifications. These calculations were performed using 4 core sweeps per re-homogenization, with 1st order integral-conserving spatial basis functions for both flux and auxiliary source term. As before, the three schemes are all progressions over the first five re-homogenizations to starting from four times the 'classic' mesh size and ending with either a `small' mesh, which has a mesh size equal to the average mesh size of the heterogeneous problem in mean free paths, the 'classic' mesh size, or an exact overlay of the heterogeneous mesh on to the homogeneous problem. All homogeneous calculations which take place after the fifth re-homogenization occur on the final mesh. This section is placed immediately

following the chapter on CSH results in order to draw contrast between the DTH and CSH methods, specifically their speedup and accuracy.

Table 13. Results of three separate progressive meshing schemes applied to the DTH method for the C5G7 benchmark.

Final mesh	Iters	Speedup	Δk (pcm)	MRE (%)		AVG (%)		MAX (%)	
				F	Th.	F	Th.	F	Th.
Small (het. avg.)	5	9.4	249.4	0.4	0.4	0.4	0.5	1.7	2.5
Classic (het. max)	5	8.7	254.5	0.4	0.4	0.5	0.5	2.2	3.0
Exact	**6**	**5.3**	**249.01**	**0.4**	**0.4**	**0.4**	**0.5**	**1.8**	**2.5**

Table 14. Results of three separate progressive meshing schemes applied to the DTH method for the BWR benchmark.

Final mesh	Iters	Speedup	Δk (pcm)	MRE (%)		AVG (%)		MAX (%)	
				F	Th.	F	Th.	F	Th.
Small (het. avg.)	9	20.9	248.6	2.7	2.9	3.4	3.8	19.1	26.3
Classic (het. max)	10	33.5	258.6	1.0	1.9	1.2	2.1	6.3	12.5
Exact	**10**	**13.1**	**254.5**	**2.5**	**2.8**	**3.1**	**3.6**	**17.6**	**24.8**

Table 13 and Table 14 contain the results of the DTH calculations performed with progressively tightened meshing schemes. It is immediately clear from these calculations that the DTH method does not benefit largely from the use of an exact mesh heterogeneous mesh overlay on to the homogeneous problem. For both benchmarks, the results when an exact overlay are no more accurate than the classic or small mesh, and require

considerably more computation time. Also apparent is that employing diffusion as the homogeneous problem solution method has resulted in a loss of eigenvalue accuracy for the C5G7 problem. In 1-D, it was shown that only the first two angular moments of the auxiliary source term are required for full accuracy of the CSH method, but the increased coupling between space and angle that 2-D adds to the calculation can cause this truncation to have a larger impact.

An interesting observation of Table 14 is that the progressive mesh tightening to a classic mesh actually has significantly more accurate results overall than for the small mesh or even the exact mesh, when applied to the 2-D BWR problem, indicating that the accuracy of the DTH method may have less to do with the overall mesh size used by the homogeneous calculation and more to do with the actual placement of mesh boundaries in each individual problem.

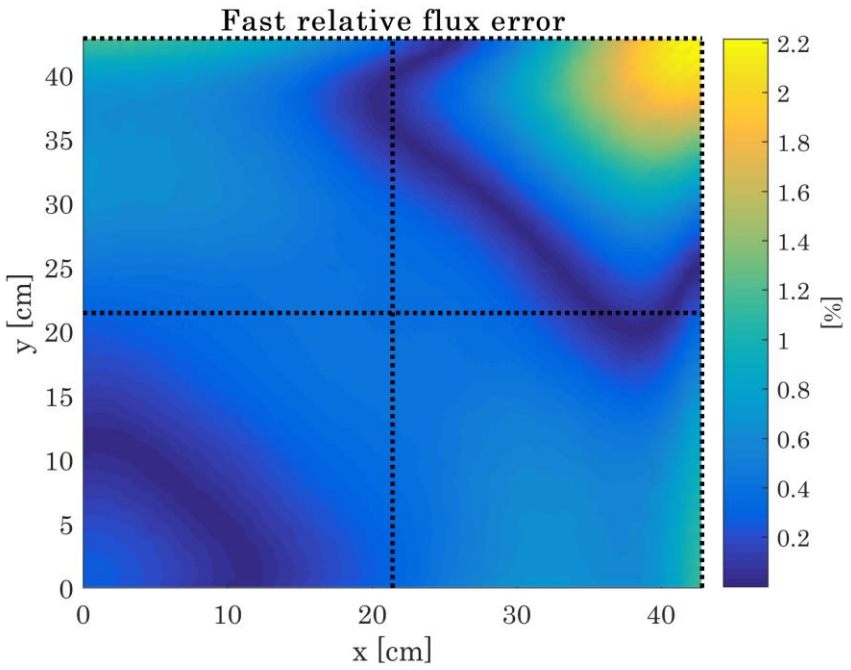

Figure 21. Fast spectrum relative error profile for the C5G7 core solved via DTH with progressive mesh to the 'classic' mesh size.

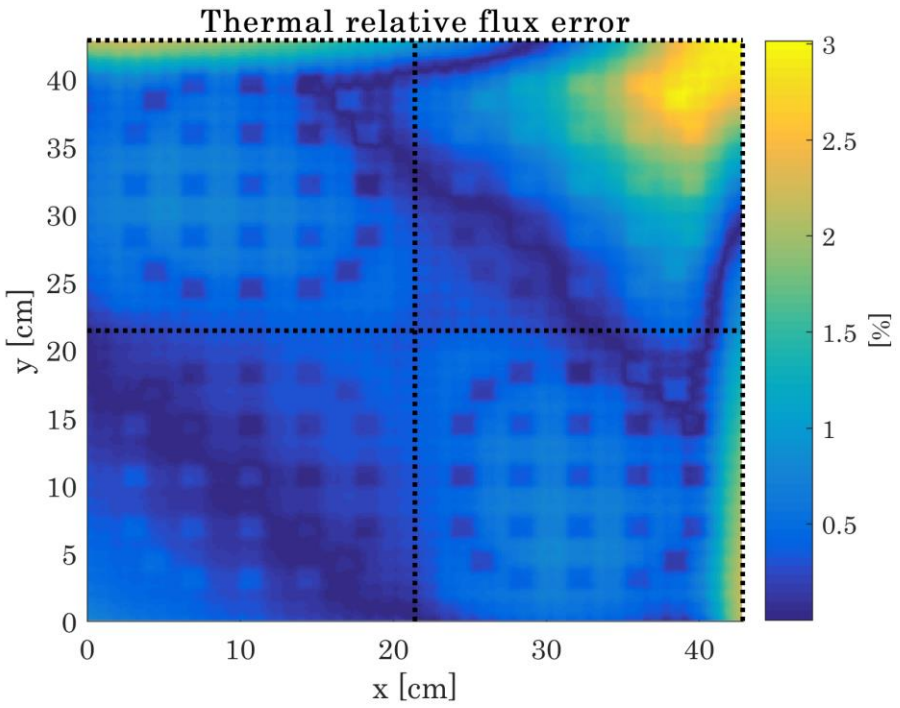

Figure 22. Thermal spectrum relative error profile for the C5G7 core solved via DTH with progressive mesh to the 'classic' mesh size.

Figures 21 and 22 contain the flux error profiles for the C5G7 core when solved via DTH with progressively tightened mesh ending with the 'classic' mesh size. The thermal flux error profile has some indication of local errors which were not properly reduced by the re-homogenization calculation, but other than that, both are dominated by core-level error effects, indicating that the calculation was successful. The error peaks are obviously fairly localized, leading to mean relative and average errors of around 0.5% in both spectra, a close match with very significant speedup.

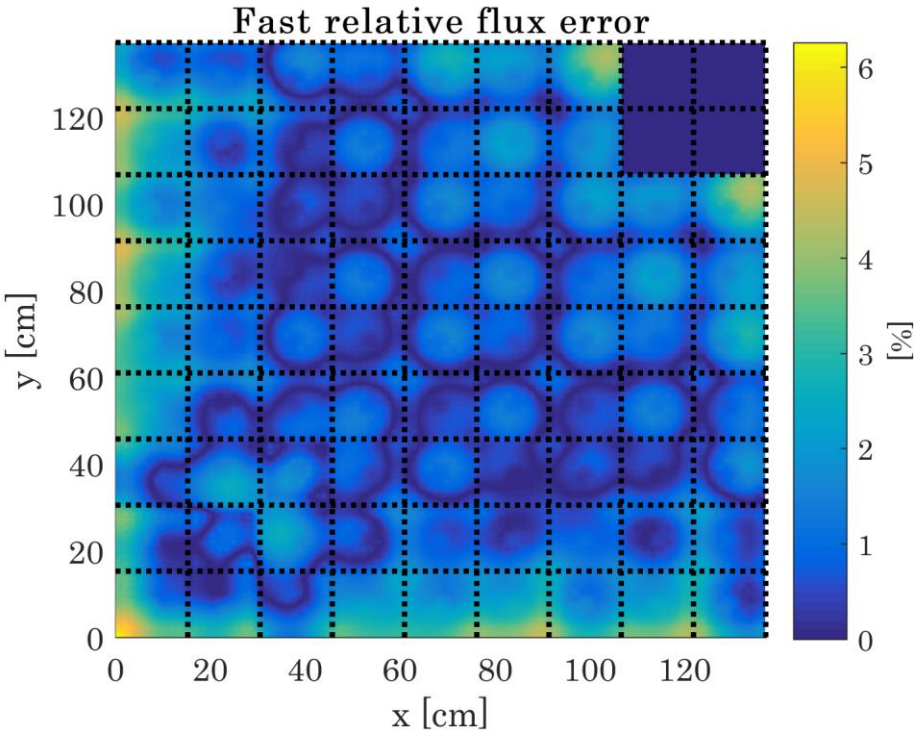

Figure 23. Fast spectrum relative error profile for the BWR core solved via DTH with progressive mesh to the 'classic' mesh size.

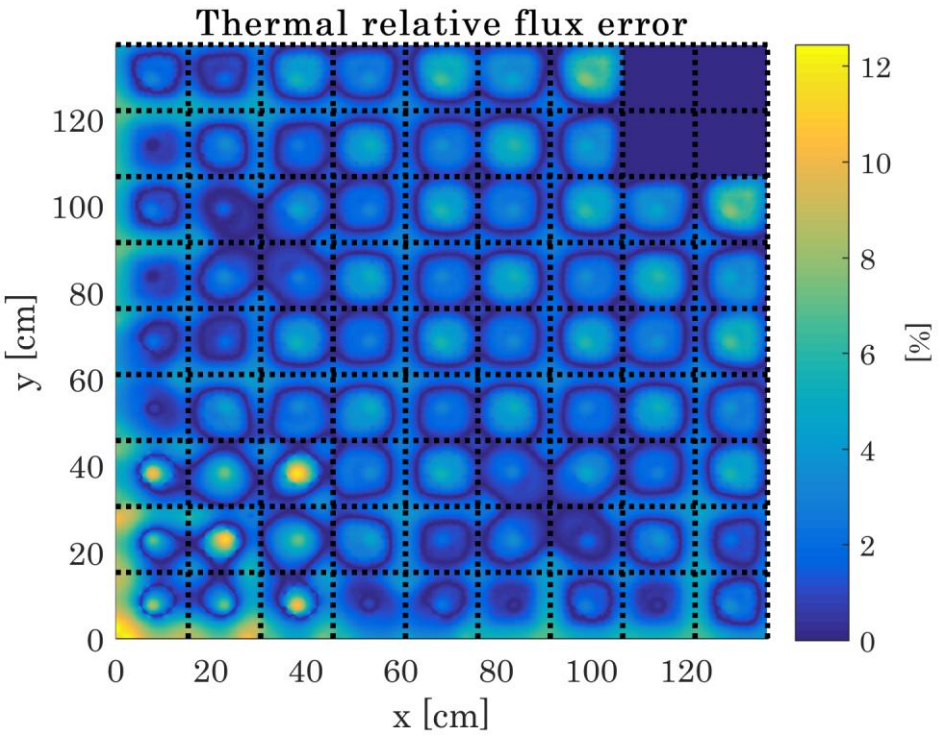

Figure 24. Fast spectrum relative error profile for the BWR core solved via DTH with progressive mesh to the 'classic' mesh size.

Figures 23 and 24 depict the relative flux error profiles for the BWR core when the DTH calculation is performed with progressively tightened meshing to the 'classic' mesh. From these plots, it is clear that the flux convergence of the DTH method for the BWR core is not very strong, and both spectra are largely influenced by local effects, indicating that higher accuracy re-homogenization may be necessary to obtain accurate DTH results.

6.2 Other Parameter Variations

In this section, as was done with the CSH method, each user-defined parameter in the DTH method will be investigated for its effect on both the

accuracy and speed of DTH calculations. Unless otherwise stated, all of these calculations have been performed using 4 core sweep re-homogenization, 1st order integral-conserving basis functions, and a 'classic' homogeneous mesh size.

6.2.1 Spatial basis functions

DTH calculations were performed using linear B-spline basis functions, as well as 0th and 1st order integral-conserving basis functions.

Table 15. The effect of different spatial basis functions on DTH calculations of the C5G7 benchmark.

Basis function	Iters	Speedup	Δk (pcm)	MRE (%)		AVG (%)		MAX (%)	
				F	Th.	F	Th.	F	Th.
Linear	8	4.8	1593.4	0.6	0.5	0.8	0.6	3.3	5.7
0th order	5	7.0	253.5	0.6	0.7	0.7	0.8	1.6	2.9
1st order	5	5.9	252.1	0.6	0.7	0.7	0.8	1.6	2.9

Table 16. The effect of different spatial basis functions on DTH calculations of the BWR benchmark.

Basis function	Iters	Speedup	Δk (pcm)	MRE (%)		AVG (%)		MAX (%)	
				F	Th.	F	Th.	F	Th.
Linear	**DID NOT CONVERGE**								
0th order	10	34.3	258.8	1.0	1.9	1.2	2.1	6.9	12.9
1st order	10	33.2	258.4	1.0	1.9	1.2	2.1	6.8	13.2

Table 15 and Table 16 summarize the effect of the different spatial basis functions on DTH method calculations. Like with the CSH method, the

integral-conserving basis functions lead to significantly better performance than the linear B-splines when mesh transfers are used in 2-D. Especially surprising is that the BWR benchmark calculation failed to converge within 20 re-homogenization iterations when linear B-splines were used for the spatial basis function, indicating complete failure of the method due to a lack of neutron balance at each re-homogenization.

A surprising result that can be found in Table 15 and Table 16 is the fact that for both benchmarks, the 0^{th} order integral-conserving basis functions resulted in higher speedup and lower error for both benchmarks. While the reason for this effect is not fully understood, it is possible that this is an effect of the fact that the 1^{st} order integral-conserving basis functions are less robust than the 0^{th} order functions. That is, if flux errors are high, or if the flux is rapidly changing, the 1^{st} order integral-conserving basis functions may result in additional error due to negative flux introduced in the mesh transfer. Any negative fluxes would be removed by the negative-flux fix-up procedure, but that process would impact the normalization, which could slow the convergence of the homogeneous problem.

6.2.2 Re-homogenization method

Results in Section 6.1 indicate that the DTH method may require more accurate re-homogenization than the CSH method. This is reasonable, since

the DTH method requires a mesh transfer of the homogeneous solution in both its angular and spatial domains, which is likely to introduce considerably more local error, and core S_N sweeps are generally Jacobi-like in their convergence of the angular and energy domains. This section contains a study and discussion on the effect of different re-homogenization methods on the DTH method's accuracy and speedup.

Table 17. Effect of different re-homogenization methods on the DTH method for the C5G7 core.

# Core sweeps	Iters.	Speedup	Δk (pcm)	MRE (%)		AVG (%)		MAX (%)	
				F.	Th.	F.	Th.	F.	Th.
1	6	7.0	217.3	0.7	1.3	0.9	1.0	2.4	3.3
2	5	6.8	240.5	0.7	0.9	0.8	0.6	2.0	2.9
4	5	5.9	252.1	0.6	0.7	0.7	0.8	1.6	2.9
7	5	4.6	253.8	0.6	0.7	0.6	0.7	1.4	2.6
10	5	3.8	253.1	0.6	0.6	0.6	0.7	0.6	0.7
20	5	2.3	254.7	0.5	0.5	0.5	0.6	1.1	1.8
AFS	**5**	1.1	259.2	0.7	1.0	0.9	0.9	2.8	4.3

Table 18. Effect of different re-homogenization methods on the DTH method for the BWR core.

# Core sweeps	Iters.	Speedup	Δk (pcm)	MRE (%)		AVG (%)		MAX (%)	
				F.	Th.	F.	Th.	F.	Th.
1	**DID NOT CONVERGE**								
2	12	29.7	225.1	0.7	2.2	0.9	2.5	7.9	14.9
4	10	33.2	258.4	1.0	1.9	1.2	2.1	6.8	13.2
7	10	21.9	294.9	1.1	1.6	1.3	1.7	4.9	9.7
10	10	21.6	319.6	1.0	1.4	1.2	1.5	4.6	7.6
20	9	16.1	364.0	1.0	1.1	1.2	1.3	4.4	5.3
AFS	**9**	11.7	391.9	0.8	0.9	1.0	1.1	6.6	8.1

As shown by Table 17 and Table 18, the increased speed of the homogeneous calculation in the DTH method compared to the CSH method implies that the overall speedup of the calculation is more strongly dominated by choice of re-homogenization method than the CSH method is. This is apparent with the precipitous change in speedup as the number of core sweeps per iteration is increased. For assembly fixed-source re-homogenization of the C5G7 core, the DTH method actually performed worse than the CSH method in this regard. This is because the assembly fixed-source calculations are performed to convergence, and the initial estimates at each re-homogenization step provided by the DTH method are weaker than those from the CSH method due to upscaling in the angle domain.

The effect of this poorer initial estimate for the re-homogenization step is especially apparent for the DTH calculation of the solution to the BWR benchmark with just one core sweep per iteration. This calculation took over twenty re-homogenization steps and did not converge. Overall, in all cases, the accuracy of the DTH method is demonstrated to be highly sensitive to the accuracy of the re-homogenization step, confirming the conclusions of [3].

6.2.3 Homogeneous problem meshing

An investigation of the effect of homogeneous problem mesh size on the DTH method calculation and results is included in this section. As with the CSH method, a wide range of mesh sizes from 1 mfp to 16 mfp have been tested.

Table 19. Effect of homogeneous problem mesh size on DTH method calculation of the C5G7 benchmark.

Mesh size (mfp)	Iters.	Speedup	Δk (pcm)	MRE (%)		AVG (%)		MAX (%)	
				F.	Th.	F.	Th.	F.	Th.
1	5	9.5	250.4	0.4	0.4	0.5	0.5	1.6	2.5
2	5	11.1	259.9	0.4	0.4	0.5	0.6	2.5	3.3
4	5	11.4	259.9	0.7	1.5	0.8	1.6	4.4	8.8
8	5	11.2	298.4	1.4	2.0	1.6	2.4	5.5	17.6
16	5	11.1	215.4	2.0	2.5	2.2	3.0	8.5	20.9

Table 20. Effect of homogeneous problem mesh size on DTH method calculation of the BWR benchmark.

Mesh size (mfp)	Iters.	Speedup	Δk (pcm)	MRE (%)		AVG (%)		MAX (%)	
				F.	Th.	F.	Th.	F.	Th.
1	10	34.5	251.8	1.0	2.0	1.2	2.1	6.7	13.0
2	10	48.4	281.8	0.6	2.1	0.7	2.2	3.3	12.4
4	9	51.4	748.3	1.0	4.2	1.0	4.5	5.4	14.4
8	8	55.3	479.9	2.8	8.2	3.0	7.7	13.3	25.7
16	6	72.5	1160.7	7.7	13.6	9.8	14.4	42.8	48.1

In Table 19 and Table 20, the calculations using an exact heterogeneous mesh overlay have not been performed, as the results of Table 13 and Table 14 indicate that the DTH method does not lead to perfect accuracy when the exact

heterogeneous mesh overlay is used. Results for the C5G7 have similar values of speedup for nearly all mesh sizes after about 2 mfp per homogeneous mesh. This indicates that the DTH calculation speedup is entirely dominated by the time it takes to perform the re-homogenization calculation (4 core sweeps in this case). As expected, the accuracy of the method declines for very large homogeneous mesh sizes, though this decline doesn't have a strong influence until the homogeneous mesh size is increased past 2 mfp. For the BWR benchmark, this leads to colossally large speedup values of nearly 50 times faster than the reference calculation.

CHAPTER 7. DISCUSSION AND CONCLUSIONS

7.1 Conclusions

The Consistent Spatial Homogenization (CSH) and Diffusion-Transport Homogenization (DTH) methods have been extended in implementation to 2-D, and have been shown to be able to successfully solve reactor eigenvalue calculations with speedup that is consistently improved over 1-D implementations of each method. In extending the implementation to 2-D, some significant improvements have been made to the method. The re-homogenization procedure has been simplified by the use of a spatial basis function expansion of the angular flux, and the spatial basis functions used in the CSH and DTH solution procedures has been updated to integral-conserving spatial basis functions. Finally, the use of a progressively tightened homogeneous problem spatial mesh has been implemented and has been shown to significantly improve the computational cost of the CSH and DTH methods without a commensurate decrease in accuracy. For the CSH method, this can even be used with an exact heterogeneous mesh overlay of the homogeneous problem to reach near perfect accuracy with modest speedup over the reference calculation, between 1.6 and 1.9, depending on the

benchmark problem. This is comparable to the speedups seen in 1-D implementations of the CSH method [1] [4].

Along the way, a robust system has been developed for 2-D transport calculations to be solved with the S_N method, implemented as an extremely portable and easy-to-use Python library, written in compiled FORTRAN 90 code. This code package is standalone, and it can be easily modified for other research projects, as well as for future extensions of the CSH and DTH methods.

7.2 Discussion

The results of the CSH and DTH parameter studies can help to determine the best set of parameters to use, depending on the needs of the user, and the success of the 2-D implementations of the CSH and DTH methods can be judged compared to the results of proof-of-concept studies in 1-D. For the CSH method, the set of parameters which yield results that can be most comfortably compared to previous implementations of the CSH method are the calculations that were performed using 4 core sweeps per re-homogenization, with a linearly progressively tightened homogeneous mesh size starting from four times the 'classic' mesh size and ending with an exact overlay of the heterogeneous mesh. When these parameters were used, the 2-D CSH method obtained 1.6 to 1.9 times speedup over the reference calculation, with less than

0.5 pcm eigenvalue and less than 1.4% maximum fine-mesh relative flux error. These errors are effectively zero, given the convergence used in the reference cases, and they are on par with both the speedup and the accuracy of 1-D CSH implementations [4]. In fact, the use of integral-conserving spatial basis functions results in a significantly closer match of reactor eigenvalue when these parameters are used, compared to 1-D implementations of the CSH method. This indicates a successful extension of the CSH, and promises excellent results for the method to be extended to further use for full-core calculations.

The results of extending the DTH method to 2-D were also successful in terms of being a clear improvement over results for the DTH method in 1-D, and DTH is promising as a method which can calculate relatively accurate fine-mesh angular flux solutions with significant speed advantages over pure heterogeneous transport calculations. By employing a progressively tightened homogeneous mesh from four times the 'classic' mesh size down to the 'classic' mesh size for DTH, calculations achieved between 8 and 33 times speedup, with less than 260 pcm eigenvalue error and between 0.5% and 2% mean relative flux error. These speedups are significantly improved over 1-D implementations of the DTH method [4], and both the eigenvalue and flux accuracy are very near to the same metrics in 1-D.

7.3 Future Work

There are several 'scales' of future work that can be undertaken for further improvements to CSH theory and implementation. In an immediate sense, there are investigations already being undertaken into alternative methods of re-homogenization that do not require full heterogeneous transport calculations. In particular, the use of High Order Diffusion [19] has shown promise as a method for re-homogenization that can still provide transport-level accuracy in calculating the auxiliary source term. This theory has been used in 1-D for a similar method undertaken on the energy domain, and the use of high order diffusion has resulted in significantly improved speedups with very little loss of accuracy. An extension of this theory to full-scale implementations of the CSH method is obvious, and it will likely have excellent results.

The results of CHAPTER 5 and CHAPTER 6 indicate that the CSH and DTH methods are viable as reactor eigenvalue solution methods. These results were performed without any sort of acceleration techniques for either the homogeneous solves or the heterogeneous core re-homogenization core sweeps. This was done in order to more fairly examine the computational benefits of the CSH and DTH methods; however, in further full-scale reactor calculations, it will be highly beneficial to implement acceleration techniques with the CSH and DTH methods for even more computational benefit.

Finally, a future investigation of the CSH and DTH theories can be done in order to focus on extending the treatment of the auxiliary source term to a full expansion in space, angle, and energy. This would involve a combination of the CSH theory with the theory of Energy Condensation [20] [21], as well as some new theory that would allow the auxiliary source term to incorporate information from a finer angular discretization as well. In such a theory, the auxiliary source term could be used to fold all of the high-order, fine-group, heterogeneous information from a given transport equation into some other low-order coarse-group homogeneous equation, which could be solved using a similar solution method to the solution method described in this thesis.

APPENDIX A. DETAILS OF 2-D BWR DISCRETIZATION GENERATION

A.1 Pin Cell Meshing Utility

The 10x10 pin cell mesh depicted in Figure 5 is a volume-conserving mesh that was generated using a custom pin cell meshing utility, which will be described in this appendix. The pin cell mesh was constrained as volume conserving in order to minimize impact on reaction rates, although restriction to an approximate mesh will always have fairly large impacts on reaction rates. Overall, the mesh was chosen in order to best approximate the physics of the reactor, without an overly large concern for maintaining the exact geometry solution.

The process by which the pin cell meshing utility works is as follows: First some user-defined number of x-direction and y-direction mesh lines are drawn that pass through points on both the fuel surface and clad outer surface. While the number of lines drawn that pass through each surface is user-defined, the points that the lines pass through are chosen to have equiangular spacing on the fuel and clad surface. At this stage, if any cell has a width that is less than 1/15 of the average cell width, then the process is repeated with a new number of mesh lines. This helps to prevent degenerate cell widths. Then, each cell of the mesh is 'colored' based on the location of its center point. If the

center point is inside the clad annulus, that cell is considered clad. If the center point is inside the fuel disc, it is considered fuel, and otherwise it is considered moderator.

The mesh defined by the previous paragraph becomes the initial estimate pin cell mesh. An error vector is calculated for the initial mesh that is the relative difference of each material's area as defined by the mesh and as defined in the exact geometry. The overall error of a mesh is taken to be the infinity norm of this error vector.

The above paragraphs describe an automatic way of generating a mesh and testing its accuracy. If the clad and fuel radii are perturbed such that $r'_{\text{clad}} = r_{\text{clad}}(1 + \delta_{\text{clad}})$ and $r'_{\text{fuel}} = r_{\text{fuel}}(1 + \delta_{\text{fuel}})$, and if these perturbed radii are used to place the mesh boundaries and 'color' cell interiors, then the process can be repeated to generate another approximate mesh of the pin cell with its own corresponding overall error value. This process can be written

$$\left|\left|e\right|\right|_{\infty} = \text{MESH}(\text{N}_1, \text{N}_2, \delta_{\text{clad}}, \delta_{\text{fuel}}, r_{\text{clad}}, r_{\text{fuel}}), \tag{49}$$

where $\left|\left|e\right|\right|_{\infty}$ is the norm of the error of a given mesh, and the function MESH performs the process described in the previous paragraphs in order to generate a pin cell mesh and its corresponding error. In Equation (49), N_1 and N_2 are the user-defined number of mesh boundaries to be placed through the fuel surface and clad surface, respectively. Only the variables δ_{clad} and δ_{fuel} are not fixed,

and so MESH is effectively a function of two variables. Once this process is defined, the nonlinear function MESH can be numerically approximated by simulating many model pin cell meshes, and it can be solved using any non-linear solution method. In generating the 10x10 BWR pin cell mesh, the solution method was chosen to be Newton's method, which converged to $||e||_\infty = 0$ within double precision. It should be noted that because MESH is nonlinear, not all choices of parameters will lead to convergent solutions with $||e||_\infty = 0$, however in practice it was found that most sets of parameters did lead to reasonable pin cell meshes.

A.2 Other Simplifications

Additional simplifications to construct the Cartesian mesh BWR model were minimal. The coolant channel within each assembly was assumed to be square and of the same size as a 2x2 grid of fuel pins, and the clad around this coolant channel was not included in any models. Additionally, the row of control rods within each cruciform control blade was assumed to be a solid blade of control material with the same diameter as the control rods themselves. This approximation increased the overall volume of control material in the core, but the overall heterogeneous problem still has an eigenvalue that is 2500 pcm higher than the exact geometry models described

in [17], suggesting that the impact of this approximation is significantly less than the impact of the Cartesian mesh discretized pin cells.

No other approximations were made for the BWR benchmark discretization.

APPENDIX B. CHOICE OF REFERENCE CONVERGENCE CRITERIA

Reference flux solutions for the C5G7 problem were calculated using convergence criteria of 6×10^{-6} in flux and 7×10^{-8} in eigenvalue, chosen so that the maximum reference flux error is 0.5% and the eigenvalue error is less than 0.5 pcm, when compared against an exact solution to the discretized equations. In this appendix, the convergence criteria will be defined, and their values will be justified.

B.1 Definition of Convergence Criteria

Within any iterative calculation, some metric must be chosen in order to measure the convergence of the calculation in order to determine when the calculation will be stopped. In general, this metric cannot be the true error of a given estimate, since generally the true solution of a calculation cannot be known while the calculation is being performed. Instead, convergence criteria usually take the form of some norm of the difference between consecutive iterations, such as

$$\left|\left|f^{n+1} - f^{n}\right|\right| \leq \epsilon. \tag{50}$$

This definition is not perfect, as it can sometimes misidentify non-convergent sequences as convergent, such as the sequence represented by the

partial sums of the harmonic series, but it is commonly used and it is suitable in practice. Ultimately, the choice of norm is arbitrary. Sometimes it is more appropriate to use a relative difference norm, such as

$$\frac{\left||f^{n+1}-f^{n}|\right|}{\left||f^{n}|\right|}\leq\epsilon. \tag{51}$$

Any norm will work, so long as it can be expected to monotonically decrease as the solution estimates converge. The convergence criteria used by the reference solutions in CHAPTER 4 are that a calculation is considered converged when following conditions hold true

$$\frac{\left|\left|\psi_{i,j,n,g}^{\ell-1}-\psi_{i,j,n,g}^{\ell}\right|\right|_{\infty}}{\left|\left|\psi_{i,j,n,g}^{\ell}\right|\right|_{\infty}}\leq\epsilon_{\psi}=6\times10^{-6}, \tag{52}$$

$$\frac{\left|k^{\ell-1}-k^{\ell}\right|}{k^{\ell}}\leq\epsilon_{k}=7\times10^{-8}. \tag{53}$$

In these convergence criteria, ℓ represents the iteration number. The use of infinity norms in Equation (52) effectively means that the flux convergence criteria are measured as maximum absolute error, normalized to the flux peak in the core.

B.2 Choice of Convergence Values

The values of ϵ_ψ and ϵ_k have meaning only insofar as they can accurately be related to the true error of the flux and eigenvalue. As such, the following calculation was performed in order to guarantee that the reference solutions were solved to within 0.5% maximum relative flux error and 0.5 pcm eigenvalue error when compared against the exact solution of the discretization. The exact solution to the discretization was approximated by running the calculation with ϵ_ψ and ϵ_k set to 1×10^{-12}.

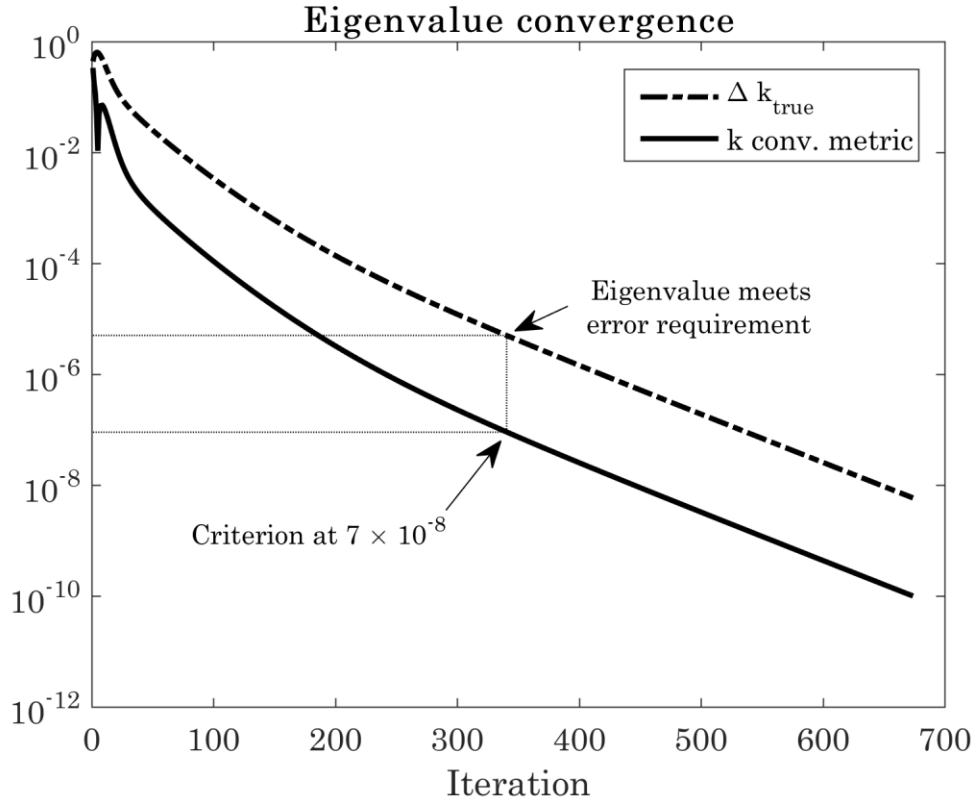

Figure 25. Result of calculation of required eigenvalue convergence criterion value

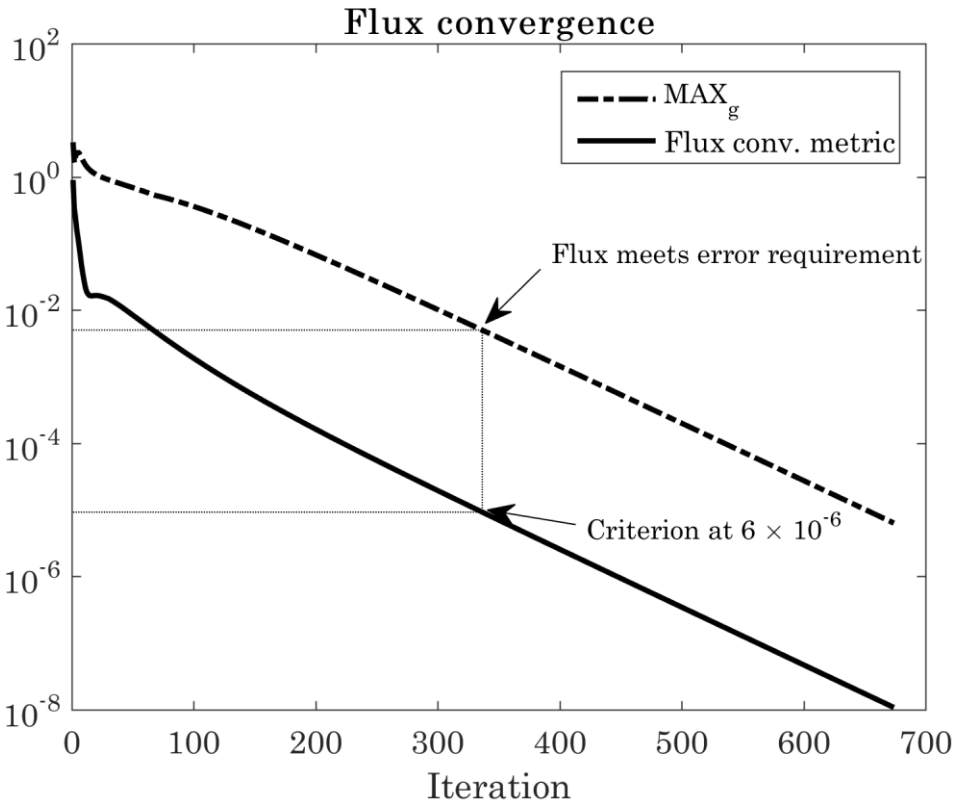

Figure 26. Result of calculation of required flux convergence criterion value

The results of these calculations are presented in Figure 25 and Figure 26. This calculation was performed on the 2-D C5G7 benchmark problem, but with S_4 angular approximation rather than S_8 due to computational limitations. There is no reason to expect dramatically different behavior for different quadrature, however. These calculations indicate that there is a well-behaved relationship between the within-calculation convergence estimates and the actual error compared against the discretization error, which can be

used to result in excellent convergence criteria. Additionally, this calculation indicates that when the chosen convergence criteria values are used, it can be expected that both the eigenvalue convergence and flux convergence will be met near the same iteration number

REFERENCES

[1] S. Yasseri and F. Rahnema, "Consistent Spatial Homogenization in Transport Theory," *Nuclear Science and Engineering,* vol. 176, pp. 292-311, 2014.

[2] G. Kooreman and F. Rahnema, "Hybrid Diffusion-Transport Spatial Homogenization Method," *Annals of Nuclear Energy,* vol. 72, pp. 95-103, 2014.

[3] G. Kooreman and F. Rahnema, "Enhanced Hybrid Diffusion-Transport Spatial Homogenization Method," *Nuclear Technology,* vol. 192, no. 3, pp. 264-277, 2015.

[4] S. Yasseri and F. Rahnema, "On the Consistent Spatial Homogenization Method in Neutron Transport Theory," *Journal of Computational and Theoretical Transport,* vol. 0, pp. 1-22, 2014.

[5] F. Rahnema and E. Nichita, "Leakage Corrected Spatial (Assembly) Homogenization Technique," *Annals of Nuclear Energy,* vol. 24, no. 6, pp. 477-488, 1997.

[6] F. Rahnema and M. S. McKinley, "High-order Cross-section Homogenization Method," *Annals of Nuclear Energy,* vol. 29, pp. 875-899, 2002.

[7] K. S. Smith, *Spatial Homogenization Methods for Light Water Reactor Analysis,* Cambridge, MA: Massachusetts Institute of Technology, 1980.

[8] D. Y. Anistratov, "Consistent Spatial Approximation of the Low-Order Quasi-Diffusion Equations on Coarse Grids," *Nuclear Science and Engineering,* vol. 149, pp. 138-161, 2005.

[9] R. Sanchez, "Assembly Homogenization Techniques for Core Calculations," *Progress in Nuclear Energy,* vol. 51, pp. 14-31, 2009.

[10] K. S. Smith, "Assembly Homogenization Techniques for Light Water Reactor Analysis," *Progress in Nuclear Energy,* vol. 17, no. 3, pp. 303-335, 1986.

[11] T. Kozlowski, Y. Xu and T. J. Downar, "Cell Homogenization Method for Pin-by-Pin Neutron Transport Calculations," *Nuclear Science and Engineering,* vol. 169, pp. 1-18, 2011.

[12] S. Yasseri and F. Rahnema, "Application of the Efficient Consistent Spatial Homogenization Method in Neutron Transport Theory to a Gas Cooled Thermal Reactor Problem," in *Physor 2014 - The Role of Reactor Physics Toward a Sustainable Future*, Kyoto, Japan, 2014.

[13] G. Kooreman, F. Rahnema and S. Yasseri, "Application of the Hybrid Diffusion-Transport Spatial Homogenization Method to a High Temperature Test Reactor Benchmark Problem," in *Physor 2014 - The Role of Reactor Physics Toward a Sustainable Future*, Kyoto, Japan, 2014.

[14] G. Kooreman and F. Rahnema, "Application of the Hybrid Diffusion-Transport Spatial Homogenization Method to a Pressurized Water Reactor Benchmark Problem," in *ANS RPSD 2014 - 18th Topical Meeting of the Radiation Protection & Shielding Division of ANS*, Knoxville, TN, 2014.

[15] E. E. Lewis and W. F. Miller, Computational Methods of Neutron Transport, Wiley-Interscience, 1993.

[16] NEA/NSC/DOC, "Benchmark on Deterministic Transport Calculations Without Spatial Homogenization, A 2-D/3-D MOX Fuel Assembly Benchmark.," *Nuclear Energy Agency, Organization for Economic Cooperation and Development,* 2003.

[17] F. Rahnema, S. Douglass and B. Forget, "2-Dimensional PWR and BWR Whole Core Benchmark Problems," in *International Conference on Mathematica, Computational Methods & Reactor Physics*, Saratoga Springs, NY, 2009.

[18] D. J. Kelley, "Depletion of a BWR Lattice Using the RACER Continuous Energy Monte Carlo Code," in *International Conference on Mathematics and Computations, Reactor Physics and Environmental Analyses*, Portland, OR, 1995.

[19] S. Yasseri and F. Rahnema, "Hybrid Subgroup Decomposition method for Solving Fine-Group Eigenvalue Transport Problems," *Annals of Nuclear Energy,* vol. 68, pp. 136-145, 2014.

[20] S. Douglass and F. Rahnema, "Subgroup Decomposition Method," *Annals of Nuclear Energy,* vol. 48, pp. 84-101, 2012.

[21] S. Yasseri and F. Rahnema, "Subgroup Decomposition Method in Diffusion Theory," *Annals of Nuclear Energy,* vol. 60, pp. 235-241, 2013.