Advanced Application of the Discrete Generalized Multigroup Method and Recondensation to Reactor Analysis

By

Matthew S. Everson

B.S. Nuclear Science and Engineering (2009)
Massachusetts Institute of Technology

M.S. Nuclear Science and Engineering (2009)
Massachusetts Institute of Technology

SUBMITTED TO THE DEPARTMENT OF NUCLEAR SCIENCE AND ENGINEERING IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY IN NUCLEAR SCIENCE AND ENGINEERING
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

FEBRUARY 2014

Signature of Author: _____

Matthew S. Everson
Department of Nuclear Science and Engineering
January 6, 2014

Certified by: _____

Benoit Forget
Associate Professor of Nuclear Science and Engineering
Thesis Supervisor

Certified by: _____

Kord Smith
KEPCO Professor of Nuclear Science and Engineering
Thesis Reader

Accepted by: _____

Mujid S. Kazimi
TEPCO Professor of Nuclear Engineering
Chair, Department Committee on Graduate Students

Advanced Application of the Discrete Generalized Multigroup Method and Recondensation to Reactor Analysis

by

Matthew S. Everson

Submitted to the Department of Nuclear Science and Engineering
on January 6, 2014 in Partial Fulfillment of the
Requirements for the Degree of Doctor of Philosophy in
Nuclear Science and Engineering

ABSTRACT

Fine-group whole-core reactor analysis remains one of the long sought goals of the reactor physics community. Such a detailed analysis is typically too computationally expensive to be realized on anything except the largest of supercomputers. Recondensation using the Discrete Generalized Multigroup (DGM) method, though, offers a relatively cheap alternative to solving the fine group transport problem. DGM, however, suffered from inconsistencies when applied to high-order spatial methods. Many different approaches were taken to rectify this problem. First, explicit spatial dependence was included in the group collapse process, thereby creating the first ever set of high-order spatial cross sections. While these cross sections were able to asymptotically improve the solution, exact consistency was not achieved. Second, the derivation of the DGM equations was instead applied to the transport equation once the spatial method had been applied, allowing for the definition of an exact corrective factor to drive recondensation to the exact fine-group solution. However, this approach requires excessive memory to be practical for realistic problems. Third, a new method called the Source Equivalence Acceleration Method (SEAM) was developed, which was able to form a coarse-group problem equivalent to the fine-group problem allowing recondensation to converge to the fine-group solution with minimal memory requirements and little additional overhead. SEAM was then implemented in OpenMOC, a 2D Method of Characteristics code developed at MIT, and its performance tested against Coarse Mesh Finite Difference (CMFD) acceleration. For extremely expensive transport calculations, SEAM was able to outperform CMFD, resulting in speed-ups of 20-45 relative to the normal power iteration calculation. Additionally, to address the growing interest in Krylov based solvers applied to reactor physics calculations, an energy-based preconditioner was developed that is inexpensive to form and can accelerate convergence.

Thesis Supervisor: Benoit Forget
Title: Associate Professor of Nuclear Science and Engineering

ACKNOWLEDGEMENTS

First, I would like express my sincere gratitude to my thesis advisor Professor Benoit Forget for all the guidance and support he provided throughout this work. I could not have accomplished what I have in this thesis had he not let me wander about in my pursuit of solutions. Thank you for having faith in me and for realizing that not all who wander are lost.

I am also very grateful for Professor Kord Smith's help in reviewing this thesis and providing insight to this work from his many years of experience in industry.

Many fellow students also helped me considerably throughout this work. To name a few, I would like to thank Lei Zhu for helping me get started with this research, Will Boyd for getting me caught up to speed with the OpenMOC code and Sam Shaner for answering all my questions about the CMFD implementation in OpenMOC. I also want to thank Nathan Gibson and Jeremy Roberts for taking the time to read through and provide helpful feedback for this thesis.

My family and close friends also provided much needed support outside of work when things got tough. To my mom and dad, thank you for supporting me and always being just a phone call away. To my undergraduate students back in J-entry, thank you for being like a second family to me while I was working on this thesis. To the rest of my friends, thank you for all your encouragement and support over these past few years. It meant so much to know that all of you were behind me 100%.

I would also like to show my sincere thanks to the most important person in my life, my wife, for putting up with all the craziness associated with life as a PhD student. Through late nights and early mornings, through the excitement of discovering something, through the sadness of something not working, you stood beside me every step of the way. You believed in me when I didn't and you helped me keep at it when I wanted to give up. I am still and always will be amazed by you.

Lastly, I would like to thank God for bringing me to where I am today. His love never failed and His blessings never ceased. His strength kept me going when everything seemed to fall apart. My completion of this thesis is a testament to His faithfulness.

1 Objectives

The end goal of this research is to make a major contribution to the field of high fidelity neutronics analysis through the production of methods and algorithms enabling fine group calculations at the whole core level. In order to achieve this goal, reliance on fine group transport sweeps to converge the fine group solution must be significantly reduced. The ability to do so in the context of high order spatial methods is also highly desirable to reduce the spatial degrees of freedom. This thesis provides the tools and methods that are key to achieving all of these objectives. However, it is important to understand both what did and did not work. Therefore, details will be provided for all the ideas and thoughts tested, which, while they may not have been breakthroughs in and of themselves, prove useful in formulating an approach that is practical, feasible and nearly universal in its application.

In Chapter 2, the primary issues surrounding whole core analysis with fine group fidelity will be discussed and some of the ways the reactor physics community has approached these issues. Chapter 3 discusses the first steps towards addressing spatial inconsistencies between the Discrete Generalized Multigroup method (DGM) and the original fine group problem and what may be the first successful attempt at producing and incorporating high order spatial cross sections in the context of multigroup collapse. Chapter 4 goes one step further and provides one of the first means of enforcing exact consistency between like spatial methods after energy condensation takes place. Chapter 5 incorporates all the aspects learned from these previous attempts and establishes the Source Equivalence Acceleration Method (SEAM). This new method defines the coarse group cross sections such that exact consistency is maintained throughout recondensation even though the coarse group problem may be using angular quadratures, orders and spatial methods different from those used in the fine group problem. Chapter 6 discusses the implementation of SEAM in OpenMOC and provides initial 2D results for the C5G7 benchmark and a 361 group problem based on the C5G7 core. The results for SEAM in both cases are compared to CMFD, the current state-of-the-art in physics-based

acceleration techniques. Chapter 7 then discusses an approach derived from DGM which allows for preconditioning of the energy problem in order to further accelerate convergence.

The ideas proposed in this thesis provide a solid foundation for further acceleration of fine group transport calculations and represent a significant step forward in realizing the goal of fine group whole core analysis.

2 Introduction: Multigroup Whole Core Analysis

Achieving high fidelity analysis of reactor neutronics is of utmost importance in analyzing the performance and safety of reactor designs. Since quantities such as power density are directly proportional to fission rate density, which in turn is dependent on the neutron distribution, accurate predictions of the distributions of neutrons in a reactor core are key to performing steady-state and transient safety analyses. These analyses are conducted by solving the neutron transport equation.

The steady-state neutron transport equation is an integral-differential equation which governs the conservation of neutrons moving through any point in space at a certain angle with a particular kinetic energy. In 3D applications, this equation has 6 dependent variables: 3 in space, 2 in angle and 1 in energy.

$$\begin{aligned}\vec{\Omega}\cdot\nabla\psi\left(\vec{r},\vec{\Omega},E\right)+\Sigma_T(\vec{r},E)\psi\left(\vec{r},\vec{\Omega},E\right)&\\ =\frac{1}{4\pi}\int_0^{\infty}\int_0^{4\pi}\Sigma_S\left(\vec{r},\overrightarrow{\Omega'}\rightarrow\vec{\Omega},E'\rightarrow E\right)\psi\left(\vec{r},\overrightarrow{\Omega'},E'\right)d\overrightarrow{\Omega'}\,dE'&\\ +\frac{\chi(\vec{r},E)}{4\pi k_{eff}}\int_0^{\infty}\nu\Sigma_f(\vec{r},E')\phi(\vec{r},E')dE'&\end{aligned} \tag{2.1}$$

Ideally, this equation would be solved exactly for a given 3D reactor model including a detailed representation of all reactor internals and their corresponding cross section data. The solution would then provide the number of neutrons at each spatial point in the core moving at a particular angle with a specific kinetic energy. In order to be accurately solved, thermal-hydraulics and fuel performance models must be coupled to the neutron transport solution in order to provide the exact neutronic conditions present in the core.

Many simplifications can be made such that this transport equation becomes solvable. If the neutron is assumed to be mono-energetic, meaning its kinetic energy remains constant, both the spatial and angular dependences can be solved. If the angular distribution is assumed to be only

linearly anisotropic, then one can derive a neutron diffusion equation which may be solved using any number of finite difference, finite element, nodal or analytic methods. In this case, the scalar (i.e. angle-integrated) fluxes are provided directly through solution of these equations. On the other hand, if one only wishes to solve this equation over a discrete set of angles, one can derive the discrete ordinates method using various treatments for the spatial dependence, step difference (SD), diamond difference (DD), method of characteristics (MOC), etc. An extra step is required in this case, since one must apply numerical integration across all angles typically using some quadrature to arrive at the scalar flux distribution.

The mono-energetic approximation, however, is extremely poor due to the nature of cross sections governing the various reactions rates. This cross section data is gathered experimentally, defined by 100,000's of data points over the energy domain and, for many of the materials used in reactors, the variation of these cross sections are extremely complex. This data may vary by many orders of magnitude, anywhere from 10^{-5} to 10^{5} barns (10^{-24} cm^2), across an energy domain that can span 12 orders of magnitude, 10^{-5} eV to 10^{7} eV. The scattering cross section is especially difficult to model accurately due to the coupled nature of scattering between energies and angles.

A stochastic approach can be taken, like in MCNP and OpenMC, in which one essentially creates a neutron and samples across the whole gambit of possible reactions according to the entire set of evaluated nuclear data. [3] Although this can provide extremely accurate analyses for reactor design, the stochastic nature of this solver leads to a standard deviation and thus an uncertainty associated with the solution which decreases only as $1/\sqrt{N}$. Here, N represents the number of neutrons interacting with a given cell within the reactor. In order to obtain an accurate representation of the fluxes across the entire core, an extremely large number of neutrons must be simulated. [3][29]

The solution of the neutron transport equation through deterministic methods is computationally cheap by comparison, since a direct solve of the transport equation directly provides the neutron fluxes across the entire reactor. However, an approximation is made which changes the energy dependence from a continuous quantity to a discrete dependence to allow solution of the neutron

spectrum in a feasible manner. This is accomplished through application of the multigroup approximation.

$$\Sigma_{T,g}(\vec{r}) = \frac{\int_{E_{g-1}}^{E_g} \Sigma_T(\vec{r},E)\phi(\vec{r},E)dE}{\int_{E_{g-1}}^{E_g} \phi(\vec{r},E)dE} \tag{2.2}$$

This approximation is applied through integration of the neutron transport equation across various bounded domains in energy, which are referred to as energy groups. Integration across each of these groups effectively provides the total reaction rates within these energy bounds. These reactions rates are then divided by the integrated fluxes to provide a set of multigroup cross sections. When applied to the continuous energy form of the neutron transport equation, this approximation provides a set of transport equations for each energy group, each of which are directly coupled to one another. These are referred to as the multigroup transport equations.

$$\begin{aligned} \vec{\Omega}\cdot\nabla\psi_g(\vec{r},\vec{\Omega}) + \Sigma_{T,g}(\vec{r})\psi_g(\vec{r},\vec{\Omega}) \\ = \frac{1}{4\pi}\sum_l (2l+1)\sum_m\sum_{g'} \Sigma_{S,l,g\leftarrow g'}(\vec{r})\,\phi_{g',l}(\vec{r}) \\ + \frac{\chi_g(\vec{r})}{4\pi k}\sum_{g'} \nu\Sigma_{f,g'}(\vec{r},E')\phi_{g',0}(\vec{r}) \end{aligned} \tag{2.3}$$

One may notice that in the production of these multigroup equations the solution to the continuous neutron transport equation is assumed to be known *a priori*. This leads us to a long standing philosophical question closely associated to whether the chicken came first or the egg. In this case, the question then becomes which comes first, the solution or the simplification? Therefore, all methods dealing with energy condensation in some way attempt to address how one obtains an approximate solution which provides an accurate solution using the simplified equations.

2.1 Cross Section Generation for Whole Core Analysis

Currently, a multi-level procedure is employed to effectively treat energy dependence in the neutron transport equation. At the first level, an extremely simplified spatial model is used to reduce the number of spatial and angular unknowns solved such that a detailed energy treatment can be undertaken. This simplification can either contain no spatial information, as in the case of an infinite homogeneous calculation, or it can represent the first level of spatial heterogeneity in core, in the case of LWRs this would be a pin cell calculation. In this calculation, the nuclear data can be represented in its original pointwise form or linear-linear interpolation of the pointwise data can also be used to solve for an extremely resolved neutron spectrum. This allows one to account for energy self-shielding effects due to the presence of resonant absorbers and, in the case of the pin cell calculation, spatial self-shielding as well. For example, in an infinite homogeneous calculation, background cross sections from either a moderator or other structural materials can be incorporated into resonance self-shielding models in order to accurately consider the effects of dilution. [15] Pin cell calculations can be used to calculate Dancoff factors in order to approximate the effects of spatial self-shielding due to the presence of fuel in an infinite repeating array. Once these important physics features are captured, energy condensation is used to produce a set of multigroup cross sections. Energy condensation allows the reaction rates from the resolved energy problem to be preserved exactly if the same spectral conditions are present in later calculations. This reduces the number of energy unknowns one needs to solve from 100,000's to a few hundred, making a more detailed spatial and angular treatment possible for larger, more complex geometries. Though theoretically few group (2-10 groups) cross sections can be produced at this point, this will lead to a very poor approximation in subsequent calculations since the spatial and angular dependence on the energy problem won't be sufficiently captured

At the second level, an assembly level calculation is conducted solving a few hundred coupled multigroup equations. This analysis incorporates the presence of different fuel enrichments, burnable poisons, control rods, and structural materials within an assembly. This step is incredibly important since any materials with significantly different properties may have a

profound effect on the reaction rates in neighboring fuel pins, due to changes in the neutron spectra. Deterministic transport methods are usually applied at this stage to adequately resolve spatial and angular dependencies not present in the ultra-fine group calculation. Once the fluxes are calculated for all the different assembly types present in the full reactor model, energy condensation is applied again, collapsing the few hundred cross sections to 2-8 group cross sections. At this point, homogenization can also be applied either across individual pin cells within the assembly or across the entire assembly to further reduce spatial degrees of freedom. While condensation and homogenization may exactly preserve reactions rates when the conditions of the assembly calculation match perfectly those in the core, preservation of net currents isn't guaranteed. This second collapse procedure makes whole core calculations feasibly by countering the increase in spatial scale and complexity through a reduction of the energy degrees of freedom.

At the third level, 2-8 multigroup equations are solved across the whole core geometry using either a low-order angular discrete ordinates approach or solving a diffusion problem. If the assumptions behind all these approximations are satisfied at the whole core level, then the solution will provide an accurate representation of the reaction rates in the core.

This approach still sidesteps the issue of needing the solution beforehand for energy condensation to preserve reaction rates exactly. If the boundary conditions of the pin cell or assembly do not match those in the full core geometry, then each application of energy condensation will introduce errors into the subsequent calculation. Unfortunately, the best assumption which can be made *a priori* is that of an infinitely repeating array of the same pin cell or assembly types. If the core is comprised entirely of the same pin cell or assembly type, then this could be a decent approximation. However, this is rarely the case. There will typically be at least 2 or 3 pin cells containing different enrichments and burnups of fuel, but others could possibly contain burnable poisons, fission chambers, water holes, etc. Assemblies of different types may be placed in close proximity to each other as well. In all of these cases, the reflective boundary conditions used to collapse cross sections can't accurately represent the interactions of ultra-fine group or fine group neutron spectra between differing adjacent materials. [1]

Therefore, additional methods have been developed to preserve some accuracy in the multi-level multigroup procedure.

2.2 *A Priori* Methods

Much of whole core analysis for the current LWR fleet is conducted through application of *a priori* methods, in which energy condensation is conducted for a single assembly using reflective boundary conditions. This is usually in conjunction with some level of spatial homogenization, either of the full assembly or the individual pin cells within the assembly. While the conditions inside the reactor may not be perfectly represented in these calculations, the approximation is typically assumed to be reasonable if there is no significant change in material properties between neighboring assemblies.

2.2.1 General Equivalence Theory

Generalized equivalence theory provided a practical and very cheap way to reduce errors accrued in the whole core calculation when differing assemblies are placed next to each other. In many of the nodal diffusion calculations conducted at the whole core level, two major constraints were typically applied: continuity of current and continuity of surface fluxes at the boundary between assemblies. However, conservation of neutrons is governed by the assembly averaged scalar fluxes and the net currents at the boundaries. Continuity of the surface scalar flux between assemblies, while providing the additional equations needed to solve the problem, doesn't really factor into neutron balance. Therefore, the surface scalar fluxes are allowed to become discontinuous between adjacent assemblies. The relaxation of this continuity condition forms the foundation of General Equivalence Theory (GET). [31]

GET's approach to defining these discontinuous surface fluxes is to compare the surface scalar flux from the heterogeneous assembly calculations to the result obtained by using the homogenized model. The ratio of the heterogeneous to the homogenized surface fluxes provides the discontinuity factor (DF). These DFs are defined using ϕ^s_{Het}, the surface scalar flux for the heterogeneous assembly calculations, and ϕ^s_{Hom}, the surface scalar flux of the assembly using the cross sections homogenized from the heterogeneous calculation.

$$f = \frac{\phi^s_{Het}}{\phi^s_{Hom}} \tag{2.4}$$

The general idea behind this approach is that the DF allows surfaces fluxes to be discontinuous in the whole core model so that the surface currents between neighboring assemblies or other homogenized regions can be more accurately captured. For example, in a diffusion calculation, two constraints are placed on the boundaries between two neighboring assemblies, continuity of the surface scalar flux and continuity of the net surface current. Since the surface scalar flux of the homogenized assembly can be much different from that of the heterogeneous assembly, the scalar flux continuity condition in the homogenized problem biases the net surface current values. When DFs are used according to Equation (2.5), the surface scalar flux continuity condition between two adjacent assemblies, *A* and *B*, is enforced using their respective DFs, f_A and f_B through multiplication with the homogenous surface fluxes, ϕ^s_A and ϕ^s_B. This forms a kind of heterogeneous surface scalar flux continuity condition for the homogenized problem, removing the bias to the net surface currents due to homogenization.

$$f_A \phi^s_A = f_B \phi^s_B \tag{2.5}$$

This produced significant improvements in assembly power distributions in whole core calculations in an extremely cheap manner. These DFs can also be produced by simultaneously homogenizing and condensing the assembly calculation to produce DFs that approximately recreate the conditions of the fine group, heterogeneous assembly calculation. Any number of pin power reconstruction methods could then be used. [31][30]

DFs have been applied to a wide variety of problems in reactor physics calculations due to their flexibility and simplicity. Their application has been extended to produce interface discontinuity factors for individual pin cells to better correct for neighboring effects at the pin cell level, as opposed to assemblies. [17] DFs have also been included in to CITATION for HTR analysis for their ability to better take into account the larger neutron currents at the edge of the reflector and in the vicinity of the control rods. [33]

2.2.2 Superhomogenization Method

A slightly different approach can be implemented through a process called Superhomogenization. For this method, homogenization and condensation occur at the pin cell level within the assembly calculation. Multiplicative factors referred to as Superhomogenization (SPH) factors, μ_k, are defined for a cell *k* by the ratio of the heterogeneous pin cell fluxes, ϕ_k^{het}, to the pin cell fluxes from the homogenous calculations, ϕ_k^{hom}.

$$\mu_k = \frac{\phi_k^{het}}{\phi_k^{hom}} \tag{2.6}$$

Since these SPH factors are determined for each of the homogenized pin cells scalar flux, and not the surface fluxes, these factors can be used to correct the homogenized cross sections for the pin cells.

$$\tilde{\Sigma}_k^{ave} = \mu_k \Sigma_k^{ave} \tag{2.7}$$

Although the heterogeneous solution is independent of the SPH corrected cross sections, the changes in the homogenized cross sections cause the homogeneous fluxes to change after each update of the SPH factors. Therefore, an iterative process must be conducted where the cross sections are updated with the SPH factors and new SPH factors are calculated using the updated cross sections. After converging, these factors should provide the necessary corrections to conserve the net current through each cell. It is important to note that there are an infinite number of SPH factors that can conserve reaction rates in this fashion, therefore an additional normalization condition is typically used to provide the closing relation. These factors have been produced at the assembly level to approximately enforce transport (S_{16}) to diffusion equivalence or transport (S_{16}) to transport (P_3) with good success. [4] [14]

In the whole core calculation, fine mesh finite-difference diffusion can be used and these SPH factors applied to their respective homogenized pin cells to enforce what should be better defined pin cell currents. While this does lead to some modest improvements in terms of assembly and pin power calculations, GET still remains the simplest and most widely used *a priori* approach.

2.3 Nonlinear Acceleration Methods

Since *a priori* methods take a one-size-fits-all approach to approximating whole core solutions in a computationally cheap manner, the spectral interactions between adjacent assemblies may not be accurately characterized, especially if simultaneous homogenization and condensation has occurred. Nonlinear acceleration methods, on the other hand, are able to capture all the spectral swapping but rely on an iterative approach in which full transport sweeps are conducted. While the full transport sweeps are expensive, the problem is solved primarily through the creation and solution of equivalent, yet cheaper, versions of the full problem. These cheaper versions can be formed using spatial homogenization, energy condensation and/or a low order angular approximation. The solutions from these equivalent problems are then used to reduce the total number of full transport calculation required to converge the fission source. This provides a relative cheap process by which the true solution of the whole core problem can be obtained. While this requires extra work relative to an assembly level calculation conducted before the fact, this provides the true solution through application of fewer full transport calculations.

2.3.1 Diffusion Synthetic Acceleration

In the early days of deterministic transport solvers, diffusion synthetic acceleration (DSA) was the primary method of use in accelerating convergence of transport problems. In DSA, an inconsistent diffusion problem is used to provide an estimate of the error between successive transport iterations (within-group or power iteration). This estimated linear error is then used to correct the transport solution to provide an acceleration of the scalar flux for the next iteration. Equation (2.8) provides an example of its use in accelerating within-group convergence. The intermediate current and fluxes are calculated using the solution from the previous transport sweep.

$$-\nabla \cdot D_g \nabla \phi_g^{l+1} + \Sigma_{R,g} \phi_g^{l+1} = Q_g^{\diamond} - \nabla \cdot J_g^{l+1/2} - \nabla \cdot D_g \nabla \phi_g^{l+1/2} \tag{2.8}$$

For Equation (2.8), D_g denotes the diffusion coefficient for coarse group g, $\Sigma_{R,g}$ is the total removal cross sections, and $Q_g^{\diamond}$ is the total source for group g neglecting within-group scattering.

The parameters $J_g^{l+1/2}$ and $\phi_g^{l+1/2}$ are the current for group g and scalar flux for group g respectively as calculated from the intermediate transport solution. The diffusion problem in Equation (2.8) is solved to provide the fully updated scalar flux, ϕ_g^{l+1}, for the next transport calculation.

This diffusion flux would accelerate the transport solution by providing a cheap means of improving the scalar flux spatial distribution. However, energy condensation and spatial homogenization were never really folded into DSA.

2.3.2 Simultaneous Homogenization and Condensation

Coarse mesh rebalance (CMR) is another method by which a simpler problem is solved but forced to be consistent with the high order angular problem. This approach, though, is typically focused on accelerating the spatial and angular complexity of high order problem, not necessarily the energy aspect. CMR not only simplifies the angular aspect of the problem, but it also incorporates spatial homogenization. A flux and volume weighted average is applied to cross sections across a collection of cells to form a larger, homogenized cell, thereby reducing the number of spatial degrees of freedom. Neutron conservation is maintained in this process by using the basic neutron balance equation and multiplying the incoming currents at the coarse mesh boundaries by a set of rebalance factors. Application to a 1D problem is highlighted in Equation (2.9). Index *i* denotes the coarse mesh, index *k* denotes a cell within coarse mesh *i* and *l* denotes the current iterate of the solution. The currents at the boundaries of cell *i*, $J_i^{+,l+1/2}$ and $J_i^{-,l+1/2}$, along with the cell averaged scalar flux, $\phi_k^{l+1/2}$, and the cell averaged source, q_k^l, are forced into exact neutron balance by solving for rebalances factors for each cell, f_i^{l+1}.

$$\left(J_i^{+,l+1/2} + J_i^{-,l+1/2} + \sum_{k\epsilon i} \Delta x_k \Sigma_k \phi_k^{l+1/2}\right) f_i^{l+1} - J_{i-1}^{+,l+1/2} f_{i-1}^{l+1} - J_{i+1}^{-,l+1/2} f_{i+1}^{l+1} = \sum_{k\epsilon i} \Delta x_k q_k^l \tag{2.9}$$

Once these rebalance factors are calculated for the previous high order iterate, the angular fluxes at the coarse mesh boundaries and the cell-averaged scalar fluxes are accelerated accordingly. This significantly reduces the total number of transport sweeps required to achieve the given tolerance. Since this approach uses a general neutron balance equation, it can be easily applied to structured and unstructured geometries.

Unfortunately, CMR is not without issues. This nonlinear approach is only conditionally stable. In many cases, when the coarse mesh is optically thick or thin, the acceleration scheme diverges. [7] However recent developments such as Coarse Mesh Angular Dependent Rebalance (CMADR) have been implemented and shown to be unconditionally stable. This approach adds additional degrees of freedom to the coarse mesh problem being solved [26][27]

Coarse mesh finite-difference (CMFD) diffusion is another widely used nonlinear acceleration scheme. In this case, the high order angular problem is represented by a low-order diffusion problem. Equivalence is enforced through the addition of current correction factors. These take on the shape of an additional set of diffusion coefficients which modify the original coefficients. These current correction factors are calculated using the net surface currents for cell i, $J_{i+1/2}^{l+1/2}$, the coarse mesh scalar fluxes, $\phi_i^{l+1/2}$ and $\phi_{i+1}^{l+1/2}$, and the normal diffusion coefficient coupling the two coarse meshes, $\widetilde{D}_{i+1/2}$. This current corrective factor between the meshes, $\widehat{D}_{i+1/2}$, produces a diffusion problem equivalent to the transport problem. Not only can Equation (2.10) take into account spatial homogenization, but it can also include energy condensation, to create a cheaper version of a fine group transport calculation.

$$\widehat{D}_{i+1/2} = -\frac{J_{i+1/2}^{l+1/2} + \widetilde{D}_{i+1/2}\left(\phi_{i+1}^{l+1/2} - \phi_i^{l+1/2}\right)}{\phi_{i+1}^{l+1/2} + \phi_i^{l+1/2}} \tag{2.10}$$

Once the current correction factors are calculated, a set of finite-difference diffusion equations are solved to provide the updated fluxes for each of the coarse meshes.

$$
\begin{aligned}
-\left(\widetilde{D}_{i-1/2}-\widehat{D}_{i-1/2}\right)\phi_{i-1}^{l+1}-\left(\widetilde{D}_{i+1/2}+\widehat{D}_{i+1/2}\right)\phi_{i+1}^{l+1} \\
+\left(\widetilde{D}_{i+1/2}-\widehat{D}_{i+1/2}+\widetilde{D}_{i-1/2}+\widehat{D}_{i-1/2}+\Delta x_i\Sigma_{T,i}\right)\phi_i^{l+1} \\
=\sum_{k\epsilon i}\Delta x_k Q_k^l
\end{aligned}
\tag{2.11}
$$

The updated coarse mesh fluxes are then used to scale the previous iterate of the fine mesh fluxes, thereby accelerating the transport problem. While this tends to accelerate transport solutions faster than CMR, it is slightly limited in application since the typical finite difference approach is more intuitive when applied to problems using structured meshes. However, recent generalizations of CMFD to unstructured meshes have been derived and used to accelerate transport calculations as well. [19] CMFD's applicability has even been extended to monte carlo codes by accelerating the convergence of the fission source distribution. [22] This acceleration of the fission source has also been included into OpenMC with the added functionality of estimating dominance ratios. [16]

CMFD also suffers from the conditional stability issues observed in CMR. In the case of CMFD, however, this usually only occurs for coarse meshes which are optically thick. The best performance in CMFD is typically observed when applied to optically thin coarse meshes. If CMFD is able to converge for an optically thick problem, little to no acceleration is observed. [21] In many applications of CMFD, a damping factor is applied to prevent such instabilities from growing and causing divergence. This damping factor is largely problem dependent but a reasonable value can be applied across a number of different problems. Partial current Coarse Mesh Finite Difference (pCMFD) is a more recent development which has also been shown to be unconditionally stable and used to accelerate MOC calculations in the NEWT transport code. This approach adds a second degree of freedom to CMFD, allowing it to create two partial current corrective factors at the boundaries of a coarse mesh instead of just recreating a single net current corrective factor. [19][18]

Generalized coarse mesh rebalance has been recently shown to bridge the gap, so to speak, between CMR and CMFD. In the derivation of GCMR, it was shown that both CMR and CMFD

are specific cases of the generalized method. This method employs a generic parameter which can be varied across a wide range of values to produce better stability than could be otherwise achieved in CMR or CMFD alone. The choice of this parameter still requires some fine tuning and the optimal choice is likely problem dependent. However, the CMFD-like method produced with GCMR tends to be very close to optimal choice for this parameter in many cases and, through this methodology, can be applied to unstructured meshes. [35]

2.3.3 Recent Developments

More recent approaches focus solely on maintaining exact consistency between a fine group and coarse group calculation. One example is a PHYSOR 2002 paper looking at enforcing consistency in the context of MOC. In order to enforce exact consistency, the outgoing coarse group angular fluxes were forced to match the collapsed fine group outgoing angular fluxes through application of discontinuity factors for each angle, segment and coarse group. This led to significant deficiencies in performance since the memory overhead became huge for larger problems. However, this method did show that it was possible to achieve full consistency, albeit at the cost of memory. [5]

A more promising approach was developed by Lulu Li at MIT. This approach formulates a low order MOC transport calculation using quadrant space-angle domains with equal weights across a pin cell homogenized problem. A low order problem equivalent to the main MOC calculation is then set up by tallying the quadrant currents at each of the pin cell surfaces. An equivalent low order source is then calculated for each of the 8 coarse mesh tracks inside each pin cell to preserve the incoming and outgoing quadrant angular fluxes. Once this lower order problem is constructed, it can then be solved and used to accelerate the high order MOC calculation. This demonstrated very good acceleration of the C5G7 benchmark, even significantly improving upon the normal CMFD approach already incorporated into OpenMOC [23]

Another approach to achieving consistency was conducted through the use of angular dependent total cross sections. These were calculated by collapsing the fine group total cross sections using the cell-averaged angular fluxes instead of the scalar fluxes. This was successfully applied using

diamond difference discrete ordinates and shown to converge to the fine group solution in that case. A reduction of the angular order by collapsing the angular problem was attempted to minimize storage costs associated with an angular dependent cross section. The result was a reduction of the high order angular quadrature to a sort of S_2 calculation using a spatially varying average cosine angle. While this approach is certainly interesting, inconsistencies lead the method to converge to the incorrect solution when the angular dependence of the coarse group total cross section was treated as an anisotropic source. [6]

2.4 Recondensation

Recondensation is slightly different from the typical nonlinear acceleration method. Rather than homogenizing and condensing a fine group problem, the focus is placed on accelerating the energy problem while staying on the same spatial discretization. Ideally, acceleration would be achieved by solving an equivalent coarse group eigenvalue problem and then using an additional method to provide a way to reconstruct the previous fine group solution using the coarse group solution. The reconstructed fine group flux would be an improved solution and subsequently used to produce new coarse group constants for another coarse group solve. If the coarse group problem and reconstruction process are equivalent to the fine group problem, then iteration of this process will lead to the fine group solution. Currently, there are two primary methods which allow for direct reconstruction of fine group fluxes: the Subgroup Decomposition Method (SGD) and the Discrete Generalized Multigroup Method (DGM).

2.4.1 Subgroup Decomposition Method (SGD)

The SGD method proposed by researchers at Georgia Tech will be covered more in-depth in Section 5.1.1.

2.4.2 Discrete Generalized Multigroup Method (DGM)

DGM is essentially a discrete form of the Generalized Energy Condensation (GEC) method, which applied a transformation through a series of continuous polynomial functions to the original fine group equations. [28] In DGM, however, a set of discrete basis functions are used to

produce a set of moment equations, referred to as DGM equations, to better represent the discrete nature of the multigroup cross sections. Once solved, these equations produce a set of flux moments which can then be used to reconstruct fine group fluxes. Sections 2.4.2.1 through 2.4.4.2, will cover the derivation, properties and usage of DGM in the context of recondensation in detail.

2.4.2.1 Derivation and Properties of DGM

Using the multigroup approximation in the group collapse process, it is assumed that the fine group spectrum within each coarse group is flat. In doing so, all information about the fine group spectrum is lost. With DGM, it is assumed that the spectrum within each coarse group is expanded using a set of orthogonal functions. While continuous functions in energy can be used to represent the fine group spectrum [28], discrete basis functions are a more natural fit for the discrete nature of the multigroup equations and constants. This discrete representation of the within-group fluxes provides the basis for the formation of the DGM moment equations. The fine group flux can be expanded in terms of these discrete basis functions within the traditional multigroup equations. After this expansion, the resulting equation is dependent on all the moments. To decouple these flux moments, the equation is multiplied by the i^{th} discrete basis function and summed. This process produces the DGM moment equations shown in Equation (2.12). [40] Although this process also works for anisotropic sources, for the purpose of this thesis, it is assumed that only isotropic scattering is taken into account.

$$\begin{aligned}\vec{\Omega}\cdot\nabla\psi_{i,g}(\vec{r},\vec{\Omega})+\Sigma_{T,0,g}(\vec{r},\vec{\Omega})\psi_{i,g}(\vec{r},\vec{\Omega})+\delta_{i,g}(\vec{r},\vec{\Omega})\psi_{0,g}(\vec{r},\vec{\Omega})\\ =\frac{1}{4\pi}\sum_{g'=1}^{G}\Sigma_{s,i,g'\to g}(\vec{r})\phi_{0,g'}(\vec{r})\\ +\frac{\chi_{i,g}(\vec{r})}{4\pi k}\sum_{g'=1}^{G}\nu\Sigma_{f,g'}(\vec{r})\phi_{0,g'}(\vec{r})\end{aligned} \tag{2.12}$$

One may notice that the multigroup constants have been redefined in Equation (2.12). This is required so that the source calculated is transformed into the same basis as the fine group flux. In this transformation, reaction rates must also be conserved. This transformation can't simply modify the fine group cross sections themselves, but must also take care to conserve the reaction rates from the fine group problem. Therefore, the fine group reactions rates are multiplied by the discrete basis functions before carrying out the typical multigroup collapse.

$$\nu\Sigma_{f,g'}(\vec{r}) = \sum_{L\epsilon g'} P_0(L,N)\phi_{g'}(\vec{r},L)\nu\Sigma_{f,g'}(\vec{r},L) \Big/ \sum_{L\epsilon g'} P_0(L,N)\phi_{g'}(\vec{r},L) \tag{2.13}$$

$$\chi_{i,g}(\vec{r}) = \sum_{K\epsilon g} P_i(K,N)\chi_g(\vec{r},K) \tag{2.14}$$

One thing to notice is that, while the definition of the fission cross section remains unchanged, the fission spectrum has been transformed using the discrete basis functions. This is because the shape of the chi spectrum is not influenced by the incoming spectrum. Only the magnitude of the outgoing fine group spectrum changes according to the total fission rate. Therefore, the summation over incoming energy groups, *L*, can be separated from the summation over the outgoing energy groups, *K*, leading to the numerator of Equation (2.13) and Equation (2.14). The coarse group fission rate on group g' is then normalized by the total flux in group *g'* to preserve the fission reaction rate. Since chi is not technically reaction rate but a spectrum, no weighting of the chi spectrum by the scalar flux is necessary in the case of energy condensation. Only the spectrum itself needs to be transformed by the discrete basis function such that the basis representing the source matches that of the flux moments.

$$\Sigma_{s,i,g'\to g}(\vec{r}) = \sum_{L\epsilon g'}\sum_{K\epsilon g} P_i(K,N)\phi_{g'}(\vec{r},L)\Sigma_s(\vec{r},L\to K) \Big/ \sum_{L\epsilon g'} \phi_{g'}(\vec{r},L) \tag{2.15}$$

The scattering cross section looks a bit different, though. Instead of the neutrons scattering from an initial fine group *L* to outgoing fine group *K*, they scatter from an initial coarse group *g'* to coarse group *g* and moment *i*. This is because the incoming energy of the scattering cross section is collapsed using the fine group flux to preserve reactions rates. Once the incoming

energy has been collapsed to their respective coarse groups, this still leaves the outgoing fine group *K*. This is weighted by the discrete basis functions before collapsing the outgoing energy. Coarse group fluxes can then be used to produce a fine group source, but one which has been transformed into the same basis as the flux moments.

$$\Sigma_{T,0,g}(\vec{r},\vec{\Omega}) = \sum_{K\epsilon g} P_0(K,N)\psi_g(\vec{r},\vec{\Omega},K)\Sigma_{T,g}(\vec{r},K) \Big/ \sum_{K\epsilon g} P_0(K,N)\psi_g(\vec{r},\vec{\Omega},K) \quad (2.16)$$

On top of the total source on the RHS of the transport equation, conservation of the streaming and removal operators must also be ensured. The streaming term itself does not explicitly depend on the energy and so the summation over the discrete basis function only produces the flux moments inside the streaming operator. However, the total removal reaction rate still needs to be conserved. The fine group total cross sections are weighted by the angular flux and collapsed to produce the cross sections. Alternatively, the scalar flux could be used to provide a coarse group total cross section, but for now the angular dependent definition will be used.

$$\Sigma_{T,i,g}(\vec{r},\vec{\Omega}) = \sum_{K\epsilon g} P_i(K,N)\psi_g(\vec{r},\vec{\Omega},K)\Sigma_{T,g}(\vec{r},K) \Big/ \sum_{K\epsilon g} P_i(K,N)\psi_g(\vec{r},\vec{\Omega},K) \quad (2.17)$$

One may notice that the total cross section was not defined using the discrete basis functions. There are two main reasons for this. First, if the moment total cross sections were defined using the flux moments as defined in Equation (2.17), the summation on the denominator could potentially become zero since only the 0^{th} order moment flux is guaranteed to be positive. This could occur if the fine group spectrum is flat within a coarse group. The first order moment total cross section would then be divided by the linear component of the spectrum, which would be zero in this case. Since computers, like human beings, tend to not like dividing by zero or numbers close to zero, this should absolutely be avoided. If left unaccounted for, this issue can lead to accrual of round off errors or divergence of the solution. Second, the use of discrete basis functions to represent the moment total reaction rates, the numerator of Equation (2.17), can result in a negative moment total cross section. This might occur if the linear component of the total reaction rates has a slope opposite that of the flux. If this is the case and a method such as

Step characteristics or MOC were being used, then the resulting solution of the first order differential transport equation produces exponential growth in the moment angular flux instead of exponential decay, possibly leading to further instability. [40] To prevent these issues the coarse group total cross sections are used for all the moment equations within a coarse group since these cross sections are always positive and finite. However, a new term must be introduced to account for the difference between using the coarse group total cross section and the moment total cross section.

$$\delta_{i,g}(\vec{r},\vec{\Omega}) = \sum_{K\epsilon g} P_i(K,N)\psi_g(\vec{r},\vec{\Omega},K)\left(\Sigma_{T,g}(\vec{r},K) - \Sigma_{T,0,g}(\vec{r},\vec{\Omega})\right) \Big/ \sum_{K\epsilon g} \psi_g(\vec{r},\vec{\Omega},K) \quad (2.18)$$

Since the coarse group total cross section is used, one can't define a distinct total cross section for each moment. A corrective term must be used that can account for this inconsistency. This is where a new term, the del term, comes into play. The del term incorporates a correction to the LHS of the DGM equations by collapsing the difference between the fine group and coarse group total cross sections. This can then be moved to the RHS of the transport equation and be treated like an angularly dependent source term.

At this point, there are still as many equations as the original fine group problem, but what has changed is how these equations are used. With DGM, the power iteration only needs to be conducted on the 0^{th} order DGM equation. This is because all of the integrated reaction rate information is contained within the 0^{th} order equation. Another nice feature is that the del term vanishes for the leading moment equation and reduces to the traditional coarse group equation. Therefore, instead of having to conduct the power iteration over all the fine group equations, this only has to be done for the coarse group equations.

$$L_{0,g}\psi_{0,g} = Q_{0,g}(k,\phi_{0,g}) \quad (2.19)$$

This by itself does not lead to an improved solution, since this still only produces a coarse group solution with the same limitations as before. However, the higher order DGM equations only depend on the coarse group scalar and angular fluxes. Therefore, once the coarse group solution has converged, these fluxes can be used to produce the sources for the higher order DGM

equations. All that's left to do is solve a set of fixed source problems to calculate the rest of the angular flux moments.

$$L_{0,g}\psi_{i,g} = Q_{i,g}\left(k, \psi_{0,g}, \phi_{0,g}\right) \tag{2.20}$$

If the moment cross sections have been properly defined for the problem, then the angular flux moments can be used to reconstruct the exact fine group solution.

$$\psi_g\left(\vec{r}, \vec{\Omega}, K\right) = \sum_{i=0}^{N-1} a_i P_i(K, N)\psi_{i,g}\left(\vec{r}, \vec{\Omega}\right) \tag{2.21}$$

2.4.3 Discrete Basis Functions

Before continuing further, it is important to discuss the choice of discrete basis function that is used to transform the original fine group equations into the DGM equations. The accuracy to which these basis functions can be constructed will in turn limit the ability of DGM to reconstruct fine group fluxes in a trustworthy manner.

2.4.3.1 Discrete Legendre Polynomials

Previous work focused primarily on using the Discrete Legendre Polynomials (DLP's) as the basis for the DGM equations. This basis is essentially the discrete analogue of the continuous Legendre Polynomials. [21] To build these DLPs, a second order recurrence relation defined by Equation (2.22) is used.

$$\begin{aligned}(m+1)(N-m)P_{m+1}(K,N) \\ = (2m+1)(N-2K)P_m(K,N) \\ - m(N+m+1)P_{m-1}(K,N)\end{aligned} \tag{2.22}$$

This relation is seeded with the initial values $P_0(K,N) = 1$ and $P_1(K,N) = (N-2K)/N$ where *N* is the total number of discrete basis functions and both *K* and *m* vary from *0* to *N-1*.

For use in DGM, normalization constants must be computed for the reconstruction of our fine group fluxes. For DLPs, these constants are defined according to Equation (2.23).

$$a_m = (2m+1)\frac{N^{(m)}}{(N+m+1)^{(m+1)}} \qquad (2.23)$$

$$N^{(m)} = N(N-1)(N-2)\cdots(N-m+1)$$

Unfortunately, forming the DLPs using this recursive relation accrues a significant amount of roundoff error for *N >50*. This is observed in Figure 1 which shows the inability of the DLPs to reconstruct a vector of random numbers using Equation (2.22). The resulting loss in orthogonality between the higher order basis functions leads to erroneous solutions when solving the DGM equations or even failure to converge.

2.4.3.2 Discrete Tchebyshev Polynomials

The use of Discrete Tchebyshev Polynomials was also proposed previously for use in DGM. These basis functions were shown to have much more favorable properties relative to the DLPs with regard to the accrual of roundoff error. This set of discrete basis functions is also defined using a recurrence relation. [42]

$$(i+1)T_{i+1}(K) = (2i+1)(2K-N+1)T_i(K) - i(N^2-i^2)T_{i-1}(K) \qquad (2.24)$$

However, the generation of discrete basis function through use of recurrence relations was dropped altogether in favor of a more direct calculation of the elements describing such a function. This was achieved through use of the Discrete Cosine Transform (DCT).

2.4.3.3 Discrete Cosine Transform

In this work, a different set of discrete basis function is used, the Discrete Cosine Transforms of Type II (DCT), for use in the DGM equations. Normally applied to problems such as JPEG file compression, the quality that makes this type of DCT suitable is that the 0^{th} order basis is a vector of 1's, which makes the leading order DGM equation still identical to the coarse group problem. [1] This type of DCT is defined by Equation (2.25).

$$P_m(K,N) = Cos\left(\frac{m\pi}{2}\left(K+\frac{1}{2}\right)\right) \tag{2.25}$$

The normalization constants for these DCTs are much more straightforward to calculate as well.

$$a_m = \begin{cases} \frac{1}{N}, & m = 0 \\ \frac{2}{N}, & m \neq 0 \end{cases}$$

Since DCTs can be built from directly evaluating cosine functions, any round off errors potentially accrued by using a recursive relation are eliminated. The only errors introduced while forming the DCTs are therefore limited to the approximations used in evaluating the cosine functions.

2.4.3.4 Best Choice of Discrete Basis

Ideally, the choice of discrete basis function should be able to expand an arbitrarily large number of fine groups within a coarse group without introducing errors into the reconstruction process. Since the orthogonality of the discrete basis functions is required to ensure accurate reconstruction of the fine group flux, then the best choice should maintain orthogonality regardless of the number of fine groups reconstructed. Therefore the ability of both DCTs and DLPs to reconstruct fine group fluxes were tested in the following manner.

1. Multiply a vector, *B*, by a matrix, *P*, whose columnspace is comprised of the basis functions
2. Multiply the resultant vector by the transpose of *P*
3. Multiply by a diagonal matrix, *a*, comprised of the normalization constants for each function
4. Subtract the final product from the original vector

$$error = B - a\left(P^T(PB)\right)$$

Here, *P*, can be any discrete basis function, but from now on the discrete basis functions will be referred by their names, *DLP* and *DCT*, which are already normalized.

If the number of basis functions, *N*, matches the length of our vector, *B*, and the discrete basis functions are exactly orthogonal to one another, the errors accrued should be on the order of machine precision. The following figure shows the results of this test using a vector of 100 random values varying from 0 to 1.

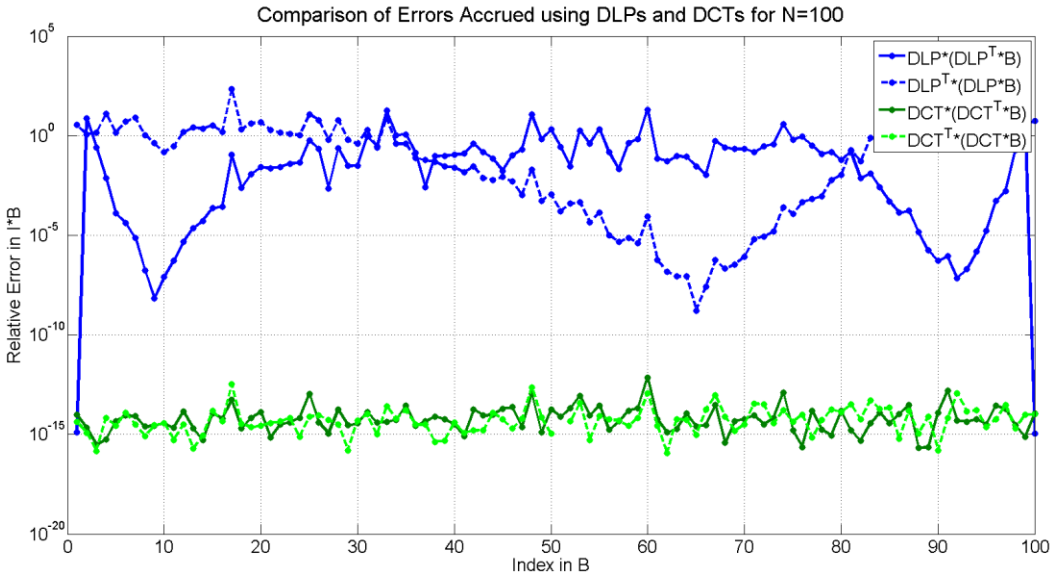

Figure 1 : Comparison of errors accrued using Discrete Legendre Polynomials and Discrete Cosine Transforms

While *DLPs* cannot reconstruct *B* to any reliable accuracy, as shown in Figure 1, *DCTs* can reconstruct all values of *B* to near machine precision. This is an unfair comparison, though, since the issue with orthogonality in our *DLPs* was already known beforehand. To be fair, the Modified Gram-Schmidt process was applied to re-orthogonalize the *DLPs* once the recursion relation has been used. While this improved the reconstruction of *B*, the combination of Modified Gram-Schmidt with *DLPs* still did not match the performance of the *DCTs*. It is important to recognize, however, that for reconstruction of a few fine groups within a coarse group, *DLPs* and *DCTs* perform similarly. Though, since the goal is to extend this application to

larger numbers of fine groups (on the order of 100s) for any given coarse group, *DCTs* is no doubt the best option.

Another very nice quality of *DCTs* is their ability to act on a vector in $O(n \log(n))$ operations using Fast Fourier Transform methods instead of the usual $O(n^2)$ number of operations. While for few fine groups this won't enhance the performance, it could prove extremely useful in preparing cross section data when using DGM to reconstruct 100's of fine groups within a single coarse group. Therefore, for the rest of our study, *DCTs* will function as our primary discrete basis function for use in DGM.

2.4.4 Recondensation with DGM

An important observation to note is that DGM by itself can't fully account for all spectral swapping at the whole core level. No matter how accurately the discrete basis functions can be evaluated, a single solve of the DGM equations still can't produce the exact fine group solution without knowing the true solution *a priori.* The root of this limitation is in the fine group spectrum used beforehand to collapse all of the moment cross sections. Even when using the moment cross sections and del terms provided by an assembly calculation, solving the DGM equations will still lead to an incorrect, albeit better, reconstructed fine group flux. Since the spectral boundary conditions for each assembly in the core aren't known a priori, the reaction rates and therefore the cross sections for all the DGM equations are not representative of the true problem. Unfortunately, the correct calculation of our cross section moments assumes *a priori* knowledge of the true fine group solution. Since nothing is gained from applying the DGM equations if the fine group eigenvalue problem has already been solved, another method is required. This led to the development of a process called recondensation.

2.4.4.1 Implementation Using DGM

To initialize the recondensation process, an initial guess is assumed for the fine group fluxes and used to calculate the cross section moments. This guess could be the result of an assembly level calculation, an infinite homogeneous calculation or simply a flat flux approximation. With these

cross sections, the power iteration is applied to the coarse group equations and the coarse group flux is used to solve the higher order DGM equations. With the new flux moments a new set of fine group fluxes is reconstructed. One additional step is added in which these reconstructed fine group fluxes are used in energy condensation to obtain a new set of moment cross sections.

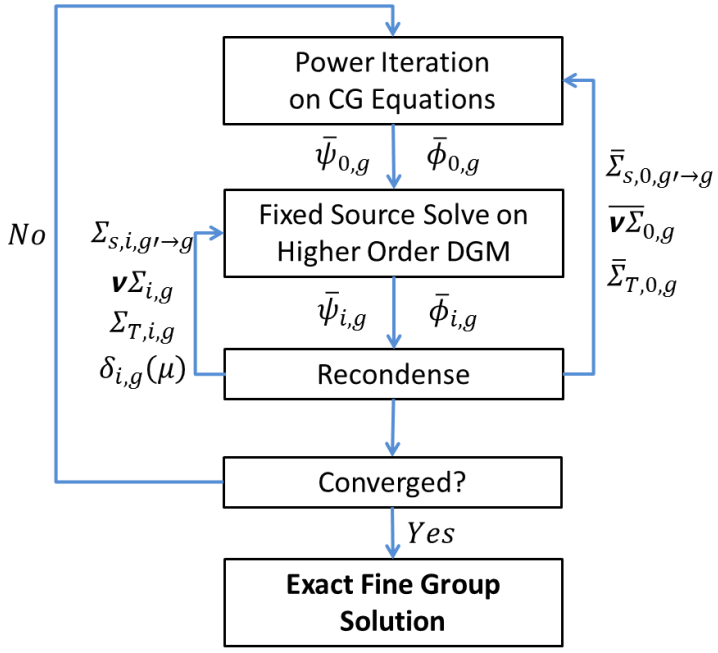

Figure 2 : Flowchart for the recondensation procedure

These updated cross sections are used in the next DGM calculation, the fine group fluxes reconstructed and the process is repeated. After a number of recondensation steps, the reconstructed fine group fluxes will converge towards the true fine group solution. This nonlinear iterative process, highlighted in Figure 2, is called recondensation. [41]

2.4.4.2 Stability and Use of Krasnoselskii Iteration

Initially, recondensation suffered from stability issues which led to divergence in many cases. This instability was first resolved by incorporating a second sweep into recondensation. One would take the fine group fluxes reconstructed form the DGM equations, construct a new fine group source with them and then carryout another transport sweep using the original fine group equations. [41]

Ideally, the recondensation process would be performed without need for any fine group sweep, which was resolved by using a Krasnoselskii iteration on the updated angular flux moments instead of employing the typical Picard Iteration. Gibson and Forget (2012) recast the recondensation process into operator notation and the Krasnoselskii iteration was applied to stabilize recondensation according to Equation (2.26).

$$\psi_{i,g}^{k+1} = (1-\lambda)\psi_{i,g}^{k} + \lambda T\psi_{i,g}^{k} \tag{2.26}$$

Here, *T* denotes the original recondensation process in operator form, *k* is the current iteration and lambda may vary from 0 to 1. The choice of lambda, of course, is problem specific, and a correct value will provide stability to the recondensation process, allowing it to converge to the true fine group solution. [12]

3 Approximate Spatial Recondensation

Much of the work conducted using DGM has been applied to low order spatial methods, such as step difference S_n. The recondensation process has already been shown to converge to the fine group solution with errors on the order of the convergence criteria set for the problem. [41] Unfortunately, when moving towards application to higher order spatial methods, DGM no longer converges to the exact fine group solution. In this section, this problem is approached through explicit inclusion of the spatial dependence from the angular and scalar fluxes.

To do this, however, requires use of high order spatial methods. Therefore, a 1D High Order Method of Characteristics (HOMOC) method is developed that achieves arbitrarily high order. This will then be used to include spatial dependences from the fluxes directly into the cross sections when group collapse is conducted. These high order spatial cross sections will then be used in a high order spatial version of the DGM equations, which should then account for much of the spatial information that is typically lost in the recondensation process when using higher order methods.

3.1 Spatial Inconsistencies in DGM

Spatial inconsistency in DGM is a direct result of the assumptions placed on the shape of the angular flux. In step difference, no shape is assumed in the calculation of the cell averaged fluxes. Therefore, the "shape" set by the streaming and removal operator is equivalent for both the fine group and DGM equations.

In the Method of Characteristics (MOC), for example, the spatial variation of the angular flux in a given cell is defined by an exponential shape. In the original fine group equations, these exponential shapes are dictated by the value of the fine group total cross section. Therefore, each of the fine group angular fluxes has a different exponential shape.

$$L_g(K)\psi_g\left(\vec{r},\vec{\Omega},K\right) = \vec{\Omega}\cdot\nabla\psi_g\left(\vec{r},\vec{\Omega},K\right) + \Sigma_{T,g}(\vec{r},K)\psi_g\left(\vec{r},\vec{\Omega},K\right) \tag{3.1}$$

The sweeping operator in the DGM equations, however, is identical across all moments since it uses the coarse group total cross section. This means that the same exponential shape is used to calculate the angular flux moments for each coarse group.

$$L_{0,g}\psi_{i,g}(\vec{r},\vec{\Omega}) = \vec{\Omega}\cdot\nabla\psi_{i,g}(\vec{r},\vec{\Omega}) + \Sigma_{T,0,g}(\vec{r},\vec{\Omega})\psi_{i,g}(\vec{r},\vec{\Omega}) + \delta_{i,g}(\vec{r},\vec{\Omega})\psi_{0,g}(\vec{r},\vec{\Omega}) \qquad (3.2)$$

While the del term that is added to try to correct for the difference in the sweeping operator, the del term is typically moved to the right hand side of the transport equation and treated as a flat source correction. This is what leads to the discrepancy between the original fine group equations and our DGM equations. The incorrect shapes assumed in DGM can't be corrected by the current definition of the del term. The lack of spatial information inside the del term introduces errors into the recondensation process, leading the method to converge to an incorrect solution.

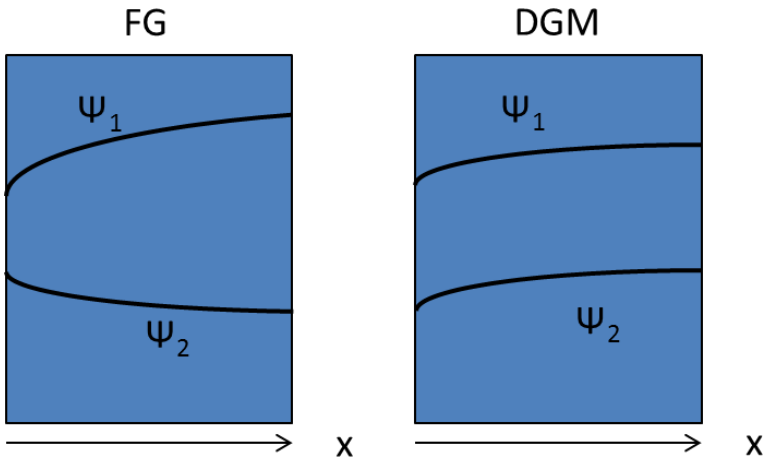

Figure 3 : Comparison of angular flux solution using the original fine group equations and the DGM equations

Figure 3 provides an example of this issue when using a characteristic type method. For the fine group case, the shapes of the angular fluxes are dependent on the total fine group cross sections. DGM, on the other hand, can only use the coarse group total cross section to define the

exponential shapes used in the calculation and therefore the same shape is applied across all moments within the same coarse group.

The consequences of this inconsistency can be seen in Figure 4. This plot provides a direct comparison of the recondensation process applied to both step difference S_n and MOC (equivalent to step characteristics in 1D) for the 47 group 1D BWR example problem defined in Appendix B. Using each of these methods, the errors in the recondensation solution relative to the fine group reference case are compared. The convergence profiles are also compared when using various convergence criteria for the fine group problem and the recondensation problem. This provides a good picture of whether or not DGM is consistent with the current spatial method being applied. If the spatial method being used in recondensation is consistent, then tightening the convergence criteria will result in a smaller difference between the reference case and the converged solution. If the method is inconsistent, then the deviation of the converged solution will reach an asymptotic value even though the convergence criteria are tightened further.

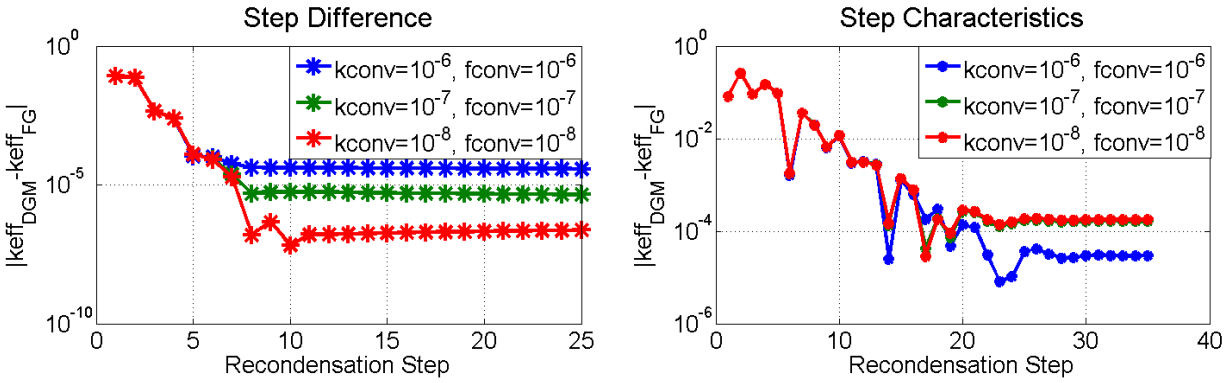

Figure 4 : Recondensation with varying convergence criteria using step difference and step characteristics

For the step difference method, as the convergence criterion for k_{eff} and fluxes is decreased, the difference between the fine group and DGM solution decreases. This shows that DGM can be applied to step difference S_n in a fully consistent manner. As predicted before for step characteristics, the errors in the recondensation solution approach an asymptotic value even when the convergence criteria are tightened.

One other interesting behavior to notice is that when the criteria are set to 10^{-6} in the case of step characteristics, this produced an unanticipated error cancellation in the fluxes and resulted in a falsely accurate eigenvalue. Therefore, it should be noted that the convergence criteria also plays a very important role in recondensation.

While implementation of step difference is straightforward, this spatial method is only 1^{st} order accurate and requires a much finer spatial discretization to achieve a high fidelity solution relative to higher order spatial methods. Since it would be widely beneficial if DGM could be made consistent with any spatial method applied in the whole core analysis, a modified definition of the del term is proposed which will incorporate local spatial dependence with the goal of correcting this inconsistency.

Typically spatial dependence is referred to in a global sense where a set of cross section or material properties are described as homogenous within a given cell but can between cells. In this case, though, it is suggested that there also exists a local spatial dependence within each cell that is typically ignored in the recondensation process. Therefore, the goal is to define a del term that includes this local spatial dependence and recaptures the spatial information typically lost. Introducing a high order spatial definition of the del term, though, requires use of a high order spatial method. In this work, the 1D High-Order MOC (HOMOC) method previously developed will be used for the analysis. [10]

3.2 A 1D High Order MOC Method for Testing

In a typical MOC calculation, the exponential shape of the angular flux is spatially averaged prior to being tallied to the cell's scalar flux. This simplifies the problem considerably since the source in each spatial cell becomes flat as well. However, in order to apply DGM in a consistent manner using 1D MOC, a way must be found to include the spatial dependence lost when the angular flux is spatially averaged. This requires a high order spatial method that provides the information needed to define a correct del term. In this section, the derivation and the convergence properties of the 1D HOMOC method will be discussed. 1D HOMOC will then be used in building a new definition of the del term that can include this missing spatial information.

3.2.1 Derivation of 1D HOMOC

To derive a high order MOC method, the spatial dependence within the 1D transport equation is cast into a dimensionless form. The change of variable is conducted along the length of a single track l_k. To do this, the variable s spanning $[0,l_k]$ is changed to x spanning [-1, 1].

$$x = \frac{2}{l_k}s - 1;\ \frac{dx}{ds} = \frac{2}{l_k} \tag{3.3}$$

This change of variable leads to the following form of the 1D transport equation.

$$\frac{d\psi(x,\mu)}{dx} + \frac{\Sigma l_k}{2\mu}\psi(x,\mu) = \frac{l_k}{2\mu}Q(x) \tag{3.4}$$

This 1D High Order MOC method assumes first that the source term on the right hand side of the 1D transport equation can be expanded with a series of Legendre polynomials up to order N.

$$Q(x) = \sum_{n=0}^{N} Q_n P_n(x) \tag{3.5}$$

In this case, the exact form of the solution is known.

$$\psi(x,\mu) = \sum_{n=0}^{N} a_n(\mu)P_n(x) + be^{-cx}\ where\ c = \frac{\Sigma l_k}{2\mu} \tag{3.6}$$

The exact form in Equation (3.6) is then substituted back into Equation (3.4), multiplied by Legendre polynomial of order l and integrated along the interval of orthogonality.

$$\int_{-1}^{1} P_l(x)\left[\frac{d}{dx}\left(\sum_{n=0}^{N} a_n(\mu)P_n(x) + be^{-cx}\right) + c\left(\sum_{n=0}^{N} a_n(\mu)P_n(x) + be^{-cx}\right) = \frac{l_k}{2\mu}\sum_{n=0}^{N} Q_n P_n(x)\right]dx \tag{3.7}$$

Upon taking the derivatives with respect to *x*, Equation (3.8) is reached. Note that the exponential term has been canceled out since this solves the homogeneous form of the 1D transport equation.

$$\int_{-1}^{1} P_l(x)\left[\left(\sum_{n=0}^{N} a_n(\mu)\frac{dP_n(x)}{dx}\right) + c\left(\sum_{n=0}^{N} a_n(\mu)P_n(x)\right) = \frac{l_k}{2\mu}\sum_{n=0}^{N} Q_n P_n(x)\right]dx \tag{3.8}$$

In order to remove the dependence on the derivative, an identity relation is used, Equation (3.9) , to represent the derivative of a Legendre polynomial of order *n.*

$$\frac{dP_n(x)}{dx} = \sum_{m=0}^{flr\left(\frac{n-1}{2}\right)} (2(n-2m)-1)P_{n-2m-1}(x) \tag{3.9}$$

This representation of the derivative is substituted back into the Equation (3.8).

$$\int_{-1}^{1} P_l(x)\left[\left(\sum_{n=0}^{N} a_n(\mu) \sum_{m=0}^{flr\left(\frac{n-1}{2}\right)} (2(n-2m-1)-1)P_{n-2m-1}(x)\right) + c\left(\sum_{n=0}^{N} a_n(\mu)P_n(x)\right) = \frac{l_k}{2\mu}\sum_{n=0}^{N} Q_n P_n(x)\right]dx \tag{3.10}$$

Now that all terms are expressed in a series of Legendre polynomials, orthogonality properties can be used to simplify Equation (3.10).

$$\sum_{n=0}^{N} a_n(\mu) \sum_{m=0}^{flr\left(\frac{n-1}{2}\right)} (2(n-2m-1)+1)\frac{2}{2l+1}\delta_{l,n-2m-1} + c\frac{2}{2l+1}a_l(\mu) = \frac{l_k}{2\mu}\frac{2}{2l+1}Q_l \tag{3.11}$$

For the left most term to be non-zero, *n-2m-1* must be equal to *l*, therefore all possible permutations of *n* and *m* that satisfy this equation must be found. For any given *n*, *m* must be equal to *(n-l-1)/2*, therefore, all possible *n* satisfying this condition from *m = 0* to the maximum possible value of *flr(N-l-1/2)* are found. The floor function (*flr*) function ignores all decimal information and only uses the integer digits of the number being acted upon. This gives us the final form of the moment equation of order *l*.

$$(2l+1) \sum_{m=0}^{flr\left(\frac{N-l-1}{2}\right)} a_{l+2m+1}(\mu) + ca_l(\mu) = \frac{l_k}{2\mu}Q_l \tag{3.12}$$

In Equation (3.12), solving for *a* is equivalent to solving a linear system of equations.

$$\bar{\bar{M}}\vec{a} = \vec{Q} \tag{3.13}$$

Once the particular solution is found, the homogeneous solution can be calculated using the incoming angular flux condition.

$$\psi_{in} = \sum_{n=0}^{N}(-1)^n a_n + be^c \tag{3.14}$$

The matrix notation can be substituted into this boundary condition to solve for *b* in terms of *M*, *Q* and the incoming angular flux, ψ_{in}. These equations provide the following relation.

$$b = \left(\psi_{in} - \overrightarrow{Opp} \cdot \bar{\bar{M}}^{-1}\vec{Q}\right)e^{-c}$$
$$\overrightarrow{Opp} = [1,\ -1,1,\cdots,(-1)^n] \tag{3.15}$$

Then the shape of the angular flux within the cell is defined by Equation (3.16).

$$\psi(x) = \psi_{in}e^{-c(x+1)} + \left(\overrightarrow{P(x)} - e^{-c(x+1)}\overrightarrow{Opp}\right) \cdot \bar{\bar{M}}^{-1}\vec{Q} \tag{3.16}$$

Evaluation of the outgoing angular flux leads to the following set of equations.

$$\psi_{out} = A\psi_{in} + \vec{B} \cdot \vec{Q}$$
$$A = e^{-2c} \tag{3.17}$$
$$\vec{B} = (\vec{1} - \overrightarrow{Opp}\, e^{-2c})\bar{\bar{M}}^{-1}$$

In this formulation, *A* represents the typical streaming term governing the contribution of the incoming flux to the outgoing flux, but one may notice that *B* has changed from a scalar value (used in ordinary MOC calculations) into a vector that determines the correct contributions of the source's spatial moments to the outgoing flux.

The angular flux spatial moments are then calculated by multiplying Equation (3.16) by the n^{th} Legendre polynomial and integrating over the cell. These moments are then found by using the following equations.

$$\vec{\psi} = \vec{C}\,\psi_{in} + \bar{\bar{D}} \cdot \vec{Q}$$
$$\vec{C} = \vec{e}\ e^{-c} \tag{3.18}$$
$$\bar{\bar{D}} = (\bar{\bar{I}} - e^{-c}\left(\vec{e} \otimes \overrightarrow{Opp}\right)) \cdot \bar{\bar{M}}^{-1}$$

Here, $\vec{e}$ is defined according to the Equation (3.19).

$$\vec{e} = [\int_{-1}^{1} P_0(x)\, e^{-cx}dx,\ \int_{-1}^{1} P_1(x)\, e^{-cx}dx, \cdots \int_{-1}^{1} P_n(x)\, e^{-cx}dx] \tag{3.19}$$

The previously scalar quantities *C* and *D* have become a vector and matrix respectively. In this new method, the vector $\vec{C}$ projects the shape of the incoming angular flux's exponential shape into Legendre polynomials and adds them to the angular flux moment. The matrix $\overline{\overline{D}}$ maps the spatial moments of the source directly to the spatial moments of the angular flux.

All of these equations, though, hinge on the correct evaluation of the exponential moments in Equation (3.19). Exact integration to find these Legendre moments of the exponential term is not necessary and in some cases can be detrimental to the calculation. This is observed for optically thin cells. In Table 1, the results are shown from analytic computation of these moments for a cell where $c = 0.05$ using both double and quad precision arithmetic.

Table 1 Comparison of Legendre moments for exponential using double and quad precision arithmetic.

	Analytic Method (double)	Analytic Method (quad)
P=0	1.000416718753101	1.000416718753100306060923165559 9
P=1	-5.0012501116115(-2)*	-5.001250111612310001289874986308 6(-2)
P=2	8.33482161510801 0(-4)	8.33482153191529014740841491081 54(-4)
P=3	-8.33438207337166 9(-6)	-8.33449080650463304594092911670 34(-6)
P=4	1.818989403545856(-7)	5.95305739182782641474428970721 33(-8)

* (-2) notation used to represent multiplication with 10^{-2}

At first, there doesn't seem to be any problem just using the analytic form derived by conducting the appropriate integration. Both the 0^{th} and 1^{st} order Legendre polynomial coefficients at double precision match the quad precision result down to 13 digits. When the Legendre order is increased further, though, the double precision result deviates substantially from the quad precision result, until no digits in the double precision calculation actually match.

The source of this error turns out to be inherent to the analytic form. The main issue is that for these optically thin cells, the integrands evaluated at x=-1 and x=1 grow by two orders of magnitude with each increase in the Legendre order while the magnitude of the difference decreases.

Table 2 Legendre moment integrands for the exponential function using double precision arithmetic

	Integrand at x=1	Integrand at x=-1
P=0	**-9.512294245007139**	**-10.512710963760240**
P=1	-5.992**745374354497**(2)	-5.992**245249343337**(2)
P=2	-5.997501**521477001**(4)	-5.997501**604825217**(4)
P=3	-8.3979004373**21818**(6)	-8.3979004373**13481**(6)
P=4	-1.511730033745314(8)	-1.511730033745314(8)

Since double precision codes can only handle precision up to roughly 16 digits, there reaches a point where the difference between the integrands falls beyond this range and produces results which are incorrect. This behavior can be seen in Table 2 above.

An alternative approach using Taylor series expansion can be applied instead to produce accurate results while still using double precision arithmetic. While the Taylor series expansion for the exponential function is very well known, the Legendre polynomial expansion is not known and must be determined for this method. To do so, a linear system of equations is formed knowing that the Taylor series expansion and the Legendre polynomial expansion must produce the same function for the infinite series of each.

$$e^{-cx} = \sum_{n=0}^{\infty} e_n(c) P_n(x) = \sum_{n=0}^{\infty} \frac{(-c)^n}{n!} x^n \tag{3.20}$$

Solving an infinite linear system of equations for the exact coefficients is obviously infeasible, therefore both the Taylor and Legendre expansion are truncated to their first 30 terms. Using this truncation, one can then solve the system of equations to determine approximate Legendre coefficients for the exponential.

Table 3 Comparison of the Taylor Series Method with the Analytic Method using quad precision arithmetic

	Taylor Expansion Method (double)	Analytic Method (>quad)
P=0	1.000416718753100	1.000416718753100306060923165559**9**
P=1	-5.0012501116123(-2)	-5.00125011161231000128987498630**8**(-2)
P=2	8.334821531915293(-4)	8.33482153191529014740841491081**5**4(-4)
P=3	-8.33449080650463**5**(-6)	-8.33449080650463304594092911670**3**(-6)
P=4	5.953057391827828(-8)	5.95305739182782641474428970721**3**3(-8)

This truncation was tested on the same optically thin cell using $c = 0.05$ and the previously specified truncation. The Taylor expansion coefficients provide the right hand side of the linear system of equations, the form of the Legendre polynomials provide the matrix on the left hand side and the Legendre coefficients are the vector of unknowns. Table 3 shows that solving for the Legendre coefficients in this manner provides accurate results in the first 5 moment values to within double precision for this thin cell problem. It is very interesting that this provides extra stability in the thin cell limit relative to the analytic form, but further research is beyond the scope of this thesis.

3.2.2 Convergence Results Using 1D HOMOC

A very simple 1D LWR fuel pin cell using 4 group cross sections was used to test each of these high order methods. For the coarsest mesh, the fuel was split into 2 equally sized cells and the coolant was modeled as one cell on each side of the fuel. The convergence criteria for both k_{eff} and the flux updates were set to 10^{-9}. Details of the geometry and group constants used can be found in **Appendix A**.

For this case, the errors in k_{eff} relative to the fully converged solution were observed to match well with the calculated l^2 and l^∞ errors of the scalar fluxes, therefore the h-p convergence results for k_{eff} were used. The reference value for the fully converged k_{eff} was found using a standard MOC solver. Beginning with the 4 cell model, each cell was split into 2 equally sized cells after satisfying the convergence criteria. These calculations were repeated until $4x2^7$ cells were used to discretize the geometry. With the data for each successive refinement, a polynomial fit was carried out on the relationship between k_{eff} and the cell size. The resulting polynomial was then

extrapolated to define a fully converged k_{eff} at the limit where the cell size approaches zero. This fitted value was then used to evaluate the convergence rates of the 1D HOMOC method.

The eigenvalue problem was solved for various cell sizes and using spatial moments beginning from the normal calculation with a 0th order spatial representation and ending at 6th order. Figure 5 shows the convergence of k_{eff} for each of these calculations.

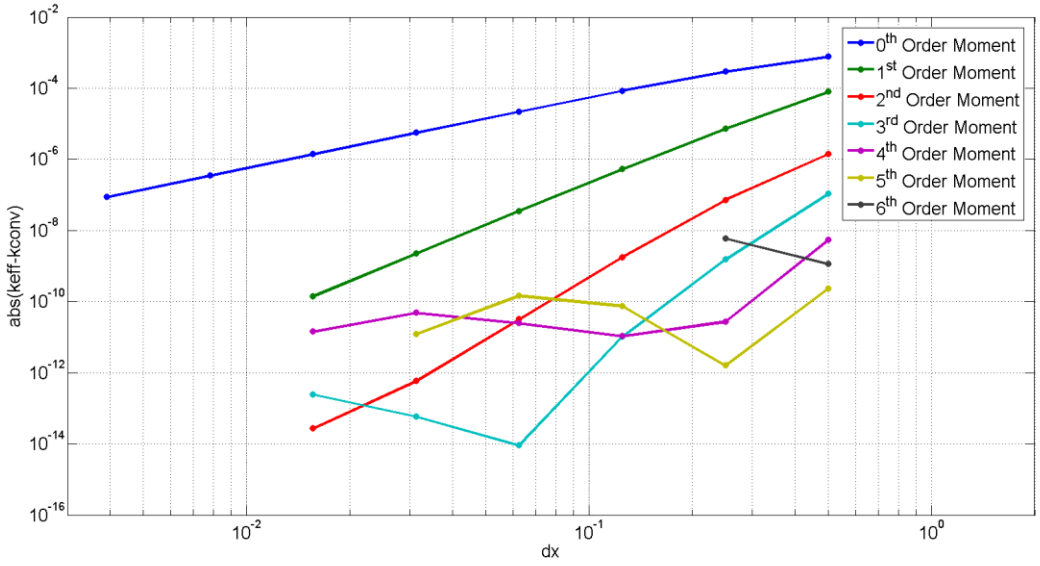

Figure 5. Convergence properties of the M-LC method

In addition, the log-log slope for the data was taken for every order up to the 3rd order, only stopping the slope calculation when the differences in the eigenvalue drop down to the order of machine precision. Greater than the 3rd order method, convergence behavior worsens until they ultimately fail to converge to any value of k_{eff} when using a 6th order spatial representation of the flux.

Table 4 Minimum, maximum and average convergence rates (log-log slopes) for M-LC method

Order of Method	Minimum Slope	Maximum Slope	Average Slope
0	1.3926	1.9997	1.8733
1	3.4631	3.9941	3.8197
2	4.3030	5.7975	5.2977
3	6.1159	10.1406	7.8363

The order of the method for this particular test case seems to go approximately as 2(*P*+1), as is observed in Table 4. When values of k_{eff} are near machine precision in accuracy, the method stagnates and no longer exhibits h-p convergence. Actual divergence is observed for the 6th order calculation, though this is likely due to both the accumulation of round-off errors in the high order equations and amplification of these errors when the moments are solved. This amplification comes directly from the system of equations being solved. When the moment equations are solved at high order and/or for optically thin cells, the system of equations becomes ill-conditioned. The explicit form of the moment equations for an 7th order calculation is provided in Equation (3.21) as an example.

$$\begin{bmatrix} \frac{\Sigma l_k}{2\mu} & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & \frac{\Sigma l_k}{2\mu} & 3 & 0 & 3 & 0 & 3 \\ \vdots & 0 & \frac{\Sigma l_k}{2\mu} & 5 & 0 & 5 & 0 \\ \vdots & \vdots & 0 & \frac{\Sigma l_k}{2\mu} & 7 & 0 & 7 \\ \vdots & \vdots & \vdots & 0 & \frac{\Sigma l_k}{2\mu} & 9 & 0 \\ \vdots & \vdots & \vdots & \vdots & 0 & \frac{\Sigma l_k}{2\mu} & 11 \\ 0 & 0 & 0 & 0 & 0 & 0 & \frac{\Sigma l_k}{2\mu} \end{bmatrix} \begin{bmatrix} a_0 \\ a_1 \\ a_2 \\ a_3 \\ a_4 \\ a_5 \\ a_6 \\ a_7 \end{bmatrix} = \frac{l_k}{2\mu} \begin{bmatrix} Q_0 \\ Q_1 \\ Q_2 \\ Q_3 \\ Q_4 \\ Q_5 \\ Q_6 \\ Q_7 \end{bmatrix} \tag{3.21}$$

When the cross section and/or track length decrease, the entire diagonal of the matrix becomes very small, significantly increasing the condition number of the system. Also, the addition of an

extra order to the system of equations increases the condition number of the system, even if the cell sizes and cross sections remain the same. Therefore, when applying the 1D HOMOC method to spatial recondensation, the maximum order method was set to *P=3*.

With this high order method explained, let us return our focus to correcting the del term definition to provide spatial consistency between DGM and the original fine group problem

3.3 Treatment of the Spatially Dependent Del Term

To incorporate local spatial variation, a del term must be defined that incorporates the spatial dependence of the angular flux within each cell, which can now be done using 1D HOMOC. This spatial dependence has been explicitly included into the del term's definition in Equation (3.22). Note that we have replaced $\vec{r}$ with *x* to differentiate global (element to element) from local (within element) spatial dependence. *x* is defined as spanning the length of a track segment from 0 to *l*.

$$\delta_{i,g}(x,\mu) = \sum_{K\epsilon g} P_i(K,N)\psi_g(x,\mu,K)\left[\Sigma_{T,g}(K) - \Sigma_{T,0,g}(\mu)\right] \Big/ \sum_{K\epsilon g} P_0(K,N)\psi_g(x,\mu,K) \qquad (3.22)$$

This new del term takes the spatially dependent angular flux and uses this to weight the difference between the fine group and coarse group total cross sections. This produces a spatially varying del term which can be used as a high order spatial source term on the RHS of the transport equation.

$$\Sigma_{T,0,g}(\mu) = \frac{1}{l}\int \Sigma_{T,0g}(x,\mu)dx = \frac{1}{l}\int \sum_{K\epsilon g} \psi_g(x,\mu,K)\Sigma_{T,g}(K) \Big/ \sum_{K\epsilon g} \psi_g(x,\mu,K)\, dx \qquad (3.23)$$

The coarse group total cross section is also defined as a spatial average of the spatially dependent cross section. This still maintains consistency in our 0th order coarse group equation since this definition causes the del term to vanish upon numerical integration in the *i*=0 case.

An unfortunate side effect of this del term and coarse group total cross section definition is that they now contain rational functions with respect to the position in a given cell. The nature of such a rational function prevents usage of Legendre moment expansion since this will not guarantee the necessary quantities will be preserved. The following thought experiments provide an understanding of why this is.

3.3.1 Moment Expansion

Assume that a del term is defined by a linear function divided by another linear function, each of which are defined using the exact angular flux. Then subsequent multiplication of this del term by the exact flux, produces $ax+b$. Multiplying this by the normalized 0^{th} order Legendre polynomial and integrating produces b, the quantity we want to be conserved.

$$\int_{-1}^{1} a_0 P_0(x)\delta(x)\psi(x)dx = \int_{-1}^{1} a_0 \left(\frac{ax+b}{cx+d}\right)(cx+d)dx = b \tag{3.24}$$

Unfortunately, if the del term in Equation (3.24) were expanded into its 0^{th} and 1^{st} continuous Legendre polynomial moments beforehand, then multiplication by $cx+d$ and integration, would not result in the exact value b. The full spatial dependence of the del term, since it is a rational function, can't be fully captured using a finite number of Legendre moment expansions.

$$\int_{-1}^{1} a_0 P_0(x)(\delta_1 x+\delta_0)(cx+d)dx \neq b \tag{3.25}$$

This means that neutron conservation would not be satisfied with this del term and would still cause the recondensation process to converge to an errant solution.

3.3.2 Numerical Quadrature

Instead of moment expansion, numerical integration is used to incorporate the spatial del term into our high order method. This requires the evaluation of all the components (the del term, the angular flux and the Legendre polynomials) at specific spatial quadrature points in the cell. These evaluations are multiplied by their corresponding quadrature weights and then summed to

conduct the numerical integration. For this work, the Gauss-Legendre quadrature is used for its ability to use *N+1* point evaluations to exactly integrate polynomials up to order *2N+1*. Therefore, as long as enough point evaluations are conducted, the Gauss-Legendre quadrature is able to exactly integrate the following test case.

$$\sum_{gp} w_{gp} a_0 P_0(x_{gp}) \delta(x_{gp}) \psi(x_{gp}) =$$

$$\sum_{gp} w_{gp} a_0 \left(\frac{a x_{gp} + b}{c x_{gp} + d}\right)(c x_{gp} + d) = \tag{3.26}$$

$$\sum_{gp} w_{gp} a_0 (a x_{gp} + b) = b$$

Evaluation at these quadrature points, as shown in Equation (3.26), allows the denominator of our del term to cancel exactly with the angular flux, leading to integration of the numerator. This is exactly what needed to happen to ensure neutron conservation. The only remaining issue is to ensure that enough points are chosen to carry out this integration exactly. Since integration of a 1st order polynomial is required in this case, 2 quadrature points is enough for exact integration. But for a given *P* order problem, the maximum order polynomial needed to be integrated is of order *2P*, a result of multiplying the P^{th} order Legendre polynomial with the P^{th} order numerator of the del term. Therefore, using *P+1* quadrature points should always be enough to evaluate our integrals exactly, at least in 1D.

3.4 1D Benchmark Results Using the Spatial Del Term

Now this new del term definition must be included into the 1D HOMOC formulation. To do this the contribution of the del term is treated as a high order spatial source. This is done using numerical integration with the Gauss-Legendre quadrature. First the total coarse group cross section is calculated for each cell and then the del term is evaluated at the spatial quadrature points. Then numerical integration is conducted after multiplication with the new iterate of the angular flux. This produces an additional source on the RHS of our transport equation which acts as a corrective term for the DGM equations.

$$Q_{del,l,i,g}(\mu) = \int P_l(x) Q_{del,i,g}(x,\mu) dx \approx \sum_{gp} w_{gp} P_l(x_{gp}) \delta_{i,g}(x_{gp},\mu) \psi_{0,g}(x_{gp},\mu) \tag{3.27}$$

This del source must be defined for each spatial moment *l*, DGM moment *i* and coarse group *g*. Once this source is calculated it can be included to the total source according to Equation (3.28).

$$Q_{l,i,g} = \frac{1}{2}\sum_{g'} \Sigma_{s,i,g'\to g}\phi_{l,0,g'} + \frac{\chi_{i,g}}{2k}\sum_{g'} \nu\Sigma_{f,0,g'}\phi_{l,0,g'} - \frac{(2l+1)}{2}\sum_{gp} w_{gp} P_l(x_{gp}) \delta_{i,g}(x_{gp},\mu) \psi_{0,g}(x_{gp},\mu) \tag{3.28}$$

The ability of this method to correct for the spatial inconsistencies inherent to DGM was tested on a simple BWR assembly described in Appendix B. To determine the effectiveness of the new del term, the errors in k_{eff} are compared for each like order solution. This means the converged 1st order fine group solution is compared to the converged 1st order DGM solution, the 2nd order fine group solution to the 2nd order DGM solution, and so on.

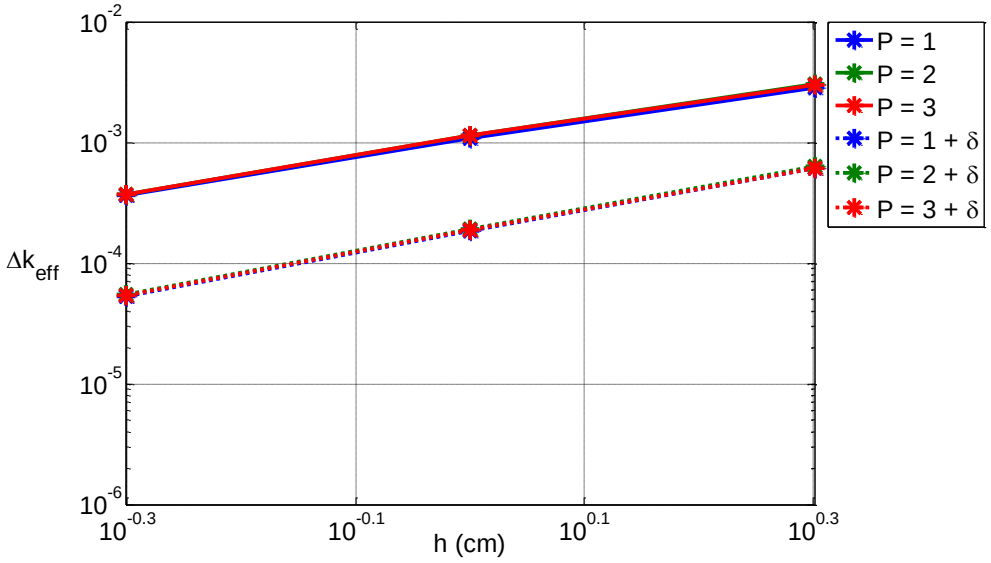

Figure 6 : Spatial convergence of errors in k_{eff} with and without the high spatial order del term

Figure 6 shows that the del term can improve k_{eff} by a significant digit, although this improvement is likely problem dependent. Also, if these corrections were fully consistent, different orders of spatial convergence would be observed, but this is not the case. This means that another source of spatial inconsistency exists that has not yet been taken to account. To determine where these errors are originating, it is necessary to take another look at the fine group equations.

3.5 Spatially Dependent Fission and Scattering Cross Sections

For the original fine group equations, the geometry is discretized such that the total, scattering and fission cross sections are spatially flat within each cell. After group collapse of the fine group cross sections, if the fine group fluxes are assumed to have some spatial variation within each cell, then the coarse group cross sections will become spatially dependent as well. This spatial dependence is usually removed by collapsing the cross sections using the spatially averaged fine group fluxes instead. This approximation results in propagation of additional errors in the coarse group calculation which can't be corrected by the spatial del term alone. Therefore it was concluded that the lack of consistency observed when only using the del term could be addressed if the local spatial dependence of the scattering and fission cross sections are also taken into account. To modify the definitions of the fission and scattering cross section, this spatial dependence of the scalar flux is included into energy condensation.

$$\nu\Sigma_{f,0,g}(x) = \sum_{K\epsilon g} P_0(K,N)\phi_g(x,K)\nu\Sigma_{f,g}(K) \Big/ \sum_{K\epsilon g} P_0(K,N)\phi_g(x,K)$$
$$\Sigma_{s,i,g'\to g}(x) = \sum_{L\epsilon g'}\sum_{K\epsilon g} P_i(K,N)\phi_{g'}(x,L)\Sigma_s(L\to K) \Big/ \sum_{L\epsilon g'} P_0(L,M)\phi_{g'}(x,L) \quad (3.29)$$

To provide clarity for these new cross section definitions, *L* is an index pointing to a particular fine group within the incoming coarse group *g'* and *K* is an index pointing to a particular fine group within the outgoing coarse group g.

This leads to the same issues that arose for the del term since these cross sections are now defined by a rational function. For reasons previously discussed in Sections 3.3.1 and 3.3.2,

numerical integration with the Gauss-Legendre quadrature is used to form the necessary spatial moments for the sources. These spatially dependent cross sections are stored at the gauss points within each cell for numerical integration with the next iterate of the scalar flux. As before, the fission and scattering sources are assumed isotropic.

$$Q_{l,f,i,g} = \int P_l(x) Q_{f,0,g}(x) dx = \frac{\chi_{i,g}}{2k} \sum_{g'} \sum_{gp} w_{gp} P_l(x_{gp}) \nu\Sigma_{f,0,g'}(x_{gp}) \phi_{0,g'}(x_{gp})$$
$$Q_{l,s,i,g} = \int P_l(x) Q_{s,i,g}(x) dx = \frac{1}{2} \sum_{g'} \sum_{gp} w_{gp} P_l(x_{gp}) \Sigma_{s,i,g'\to g}(x_{gp}) \phi_{0,g'}(x_{gp}) \quad (3.30)$$

These new definitions provide additional information to our transport equation regarding the spatial variation of the fission and scattering sources within each cell. At this point, all the spatial information that is lost in the typical multigroup collapse has been placed back into these cross section definitions.

3.6 1D Benchmark Results Using Spatial Cross Sections

These new definitions of the fission and scattering sources were incorporated into the 1D HOMOC code and tested on the 47 group 1D BWR assembly problem defined in Appendix B

. The exact fine group solution of each given spatial order calculation is used to collapse the fine group cross sections to provide the spatially dependent coarse group values for each coarse group calculation of matching spatial order (a 1st order fine group calculation provides the values for the 1st order coarse group calculation for example). The eigenvalue solve on the coarse group problem is applied using these cross sections. This may seem like a waste, since in a realistic calculation one would not begin with any knowledge of the fine group solution, but it serves as an important metric for the purposes of identifying consistency. Full consistency dictates that if the exact fine group solution is used to condense these new cross sections, then the coarse group eigenvalue problem should converge to the exact same eigenvalue. If this is satisfied, then recondensation will also be able to converge to the true fine group solution.

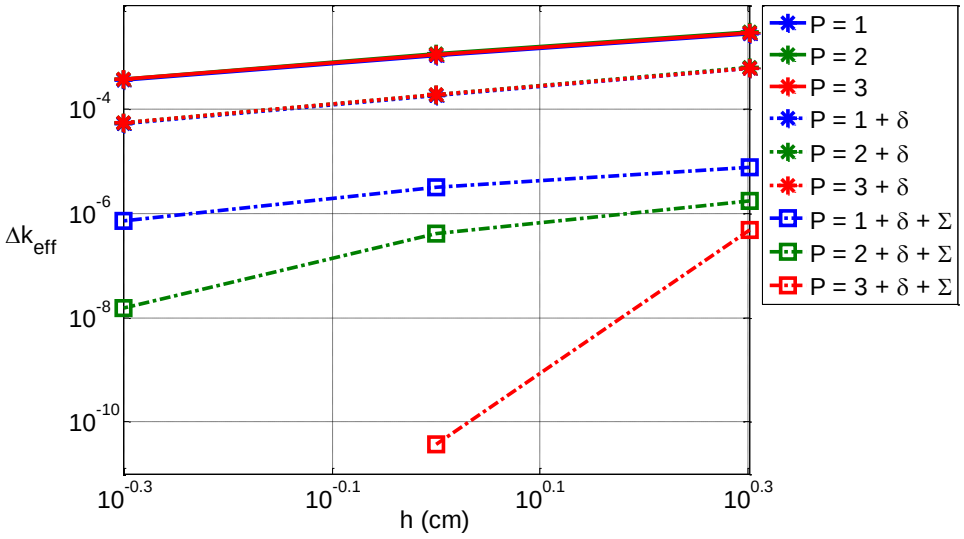

Figure 7 : Spatial convergence of errors in k_{eff} (y-axis) relative to the inverse mesh spacing (x-axis) with and without using high spatial order del term and cross sections

The result is an asymptotically consistent set of coarse group equations that can correct for the spatial inconsistencies between the fine group equations and the leading order DGM equation when using a high order spatial method. It is important to note, however, that exact consistency is not attained with this approach. The errors accrued when solving the DGM equations instead of the original fine group equations are due to the difference in the streaming operators, which are defined using exponentials in MOC. These differences between exponential terms produce an error defined by an infinite sum, which can only be partially accounted for using the N^{th} spatial order corrections applied through the cross sections and del terms. Since these high order calculations do not achieve exact consistency, it is less clear from these examples that the recondensation procedure accompanying the DGM equations will converge closer to the fine group solution. These new definitions for the group constants will now be incorporated into the recondensation process.

Each spatial order recondensation calculation is conducted starting with 0^{th} order and ending with 3^{rd} HOMOC using successively tighter convergence criteria for k_{eff} and the fluxes. Every calculation uses the same discretization of the core and uses 100 iterations for the recondensation process. While this may seem excessive, it is a direct result of the poor choice of lambda for the

Krasnoselskii Iteration and coarse group structure, the optimal choice of which is beyond the scope of this section. Absolute errors are calculated relative to the true fine group solution for the same spatial order calculation using the same convergence criteria (e.g. 1st order spatial DGM solution compared to 1st order fine group solution). Any persisting inconsistencies in the recondensation procedure will cause the solution to converge to an incorrect solution even if the convergence criteria are tightened. Thus, if increasing consistency is achieved with this high order spatial recondensation process, consecutively lowering the convergence criteria should lead to decreases in errors relative to the other low order problems. This behavior is exactly what is observed in Figure 8.

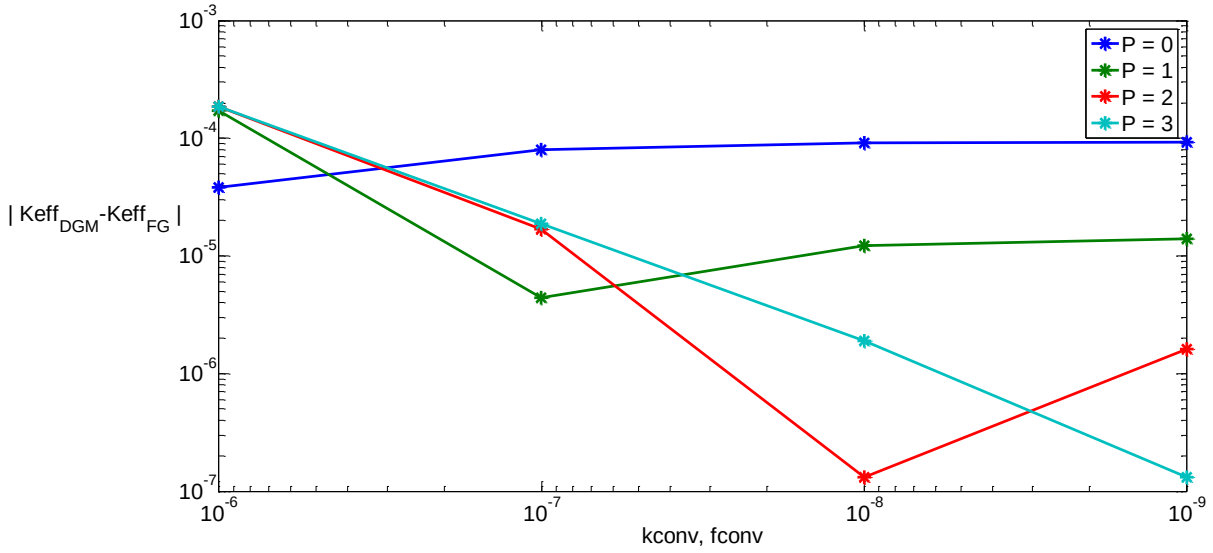

Figure 8 : Comparison of errors in k_{eff} (y-axis) for different convergence criteria (x-axis) using various spatial order transport calculations

Figure 8 shows that when moving from 0th order to 3rd order spatial recondensation, the converged value of k_{eff} improves by approximately 1 significant digit for each increase in spatial order. There is some curious behavior observed, however, since there is no monotonic trend in the error as the convergence criterion is tightened until we reach the 3rd order calculation. Error cancellation is the most likely culprit, since k_{eff} is a global quantity that depends on a sum over

all scalar fluxes. Therefore, to show that our methodology is truly working, it is necessary to show the root mean square (RMS) errors for the reconstructed fine group scalar fluxes instead.

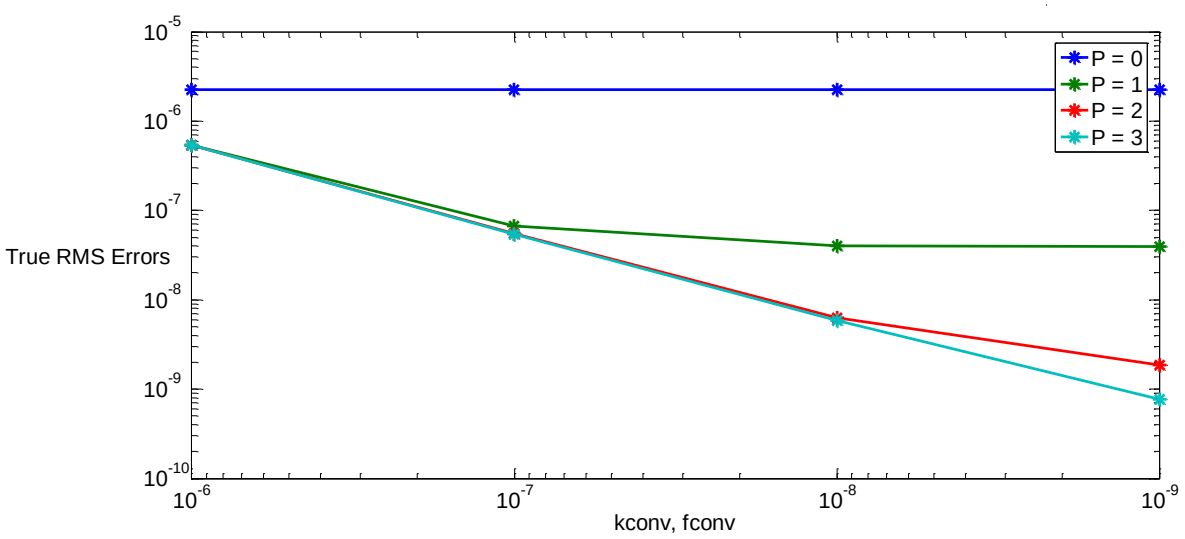

Figure 9 : Comparison of group 1 flux RMS errors (y-axis) for different convergence criteria (x-axis) using various spatial order transport calculations

For the group 1 reconstructed flux, Figure 9 shows that the converged errors behave similar to what was observed in the absolute errors in k_{eff}. Absolute RMS errors in the flux converged to $2.2\text{x}10^{-6}$ for the 0^{th} order calculation, $4.0\text{x}10^{-8}$ for the 1^{st} order, $1.6\text{x}10^{-9}$ for the 2^{nd} order and $3.6\text{x}10^{-10}$ for the 3^{rd} order. This same behavior was found in every reconstructed fine group except for group 19 fluxes.

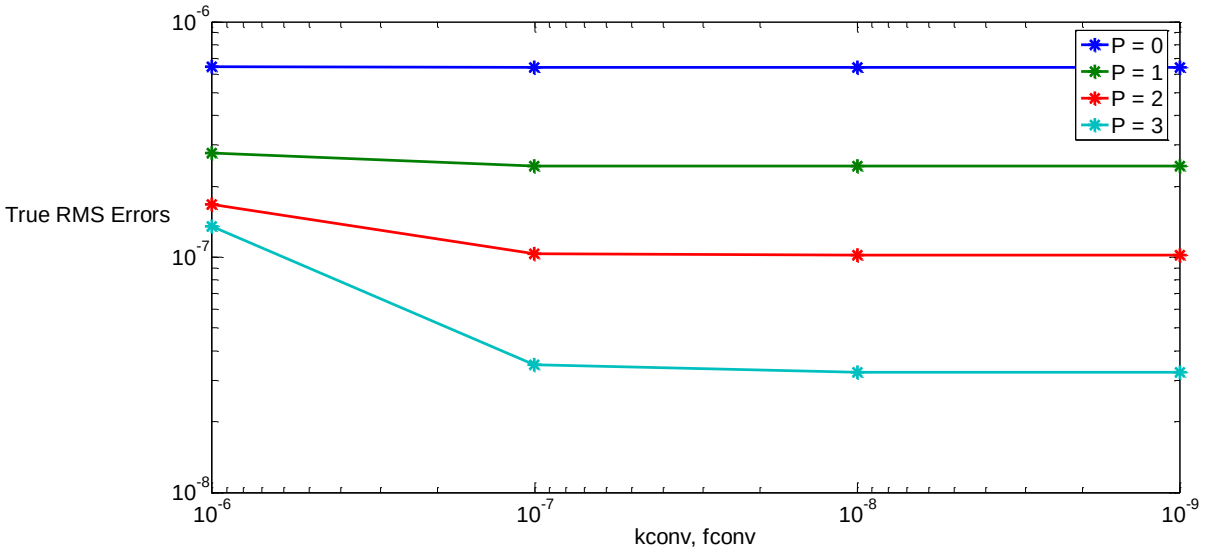

Figure 10 : Comparison of group 19 flux RMS errors (y-axis) using different convergence criteria (x-axis) for various spatial order transport calculations

Figure 10 shows that the RMS errors in group 19 flux decrease much slower than the other groups. This is likely due to the fine group total cross section in group 19 being significantly higher than the total cross section for the coarse group it lies within. Therefore, the del term requires a higher order spatial calculation in order to capture the necessary spatial information for a better reconstructed flux. Perhaps the 0^{th} order calculation is already a good enough approximation for this particular fine group, since moving to a high order spatial method in a high absorbing material can lead to numerical difficulties in some cases. It is apparent, though, from Figure 8 that this has little effect on the accuracy of the global solution.

Therefore, for this simple 1D case, it has been demonstrated that using high spatial order material properties in DGM can recapture spatial information lost in the typical group collapse procedure.

3.7 Improving the 0^{th} Order Coarse Group Solution

While these new cross section definitions have driven the recondensation solution closer to the fine group solution of an equivalent order, coupling a 0^{th} order coarse group to the 0^{th} order fine

group problem in a consistent manner still needs to be addressed. While a lot of work has been put into showing consistency between the true fine group solution and the DGM solution in order to validate the methods described here, it should be noted that the eventual goal is to improve the coarse group solution with as little extra work as possible. Therefore, a new approach must be established to couple the high order spatial DGM equations to the 0^{th} order coarse group equation to drive the recondensation process towards the 0^{th} order fine group solution.

3.7.1 Methods for Using High Order Spatial DGM for Corrections

Unfortunately, conducting a purely 0^{th} order recondensation step, by definition, results in a del term that is always 0. Spatial information from the cross sections must be passed into the 0^{th} order equations even though only spatially averaged fluxes are produced by the eigenvalue solve of the coarse group problem. Therefore, a post-processing step is added after the power iteration to grab spatial information lost when the angular flux is spatially averaged. This can be done using some of the new terms derived in 1D HOMOC. More specifically, the full *C* vector and the first column of *D* are used to obtain higher order spatial moment contributions from ψ_{in} and the flat source respectively. Once the coarse group spatial moments are calculated, the high spatial order DGM equations can use them to produce spatially dependent cross sections which are then averaged and passed back into the 0^{th} order coarse group equations.

To keep the coarse group problem consistently 0^{th} order, the fission and scattering source are limited to their 0^{th} order definitions while loosening the restriction placed on the del source, allowing use of high order del sources. But these high order sources' contributions to the 0^{th} order angular fluxes are still needed. Therefore, the only parts of matrix *D* used in subsequent calculations are the first column, which will provide the spatial information necessary for the high order spatial DGM equations, and the first row, which incorporates the spatially dependent del term into our 0^{th} order calculation.

$$\overline{\overline{D_{LO}}}\vec{Q} = \begin{pmatrix} D_{00} & D_{01} & D_{02} & D_{03} \\ D_{10} & 0 & 0 & 0 \\ D_{20} & 0 & 0 & 0 \\ D_{30} & 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} Q_0 \\ -Q_{del,1} \\ -Q_{del,2} \\ -Q_{del,3} \end{pmatrix}$$

This mixed low order, high order recondensation process can then be depicted according to Figure 11.

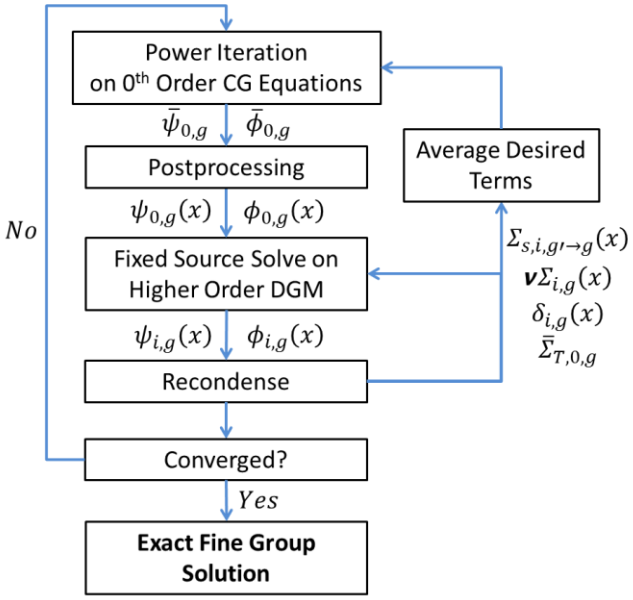

Figure 11 : Flowchart for including high order spatial information into the 0th order recondensation procedure

The result is a decoupling of the high order DGM equations from the 0th order coarse group equations, allowing for a flexible treatment of the material properties. Any number of combinations of averaged and spatially dependent cross sections or del terms can be used in the coarse group equations. The following combinations were tested to identify the level of spatial dependence necessary to improve upon the original 0th order recondensation solution. It should be further noted that including more spatial information will lead to increases in the memory requirements of the method.

1. Only use spatial dependence in the del term.
2. Use the spatial dependence of angular flux and del term to evaluate the del source.
3. All spatial dependences are considered, including fission, scattering and scalar fluxes.
4. All spatial dependences are considered and the full matrix *D* is used in the postprocessing step that takes place in each iteration within the coarse group eigenvalue solve.

3.7.2 Results for 1D Benchmark

Each of these steps successively adds more spatial information from the high order DGM solution into our 0^{th} order coarse group equations. The fourth test case is used to validate the use of the low order *D* matrix to maintain consistency with the 0^{th} order problem.

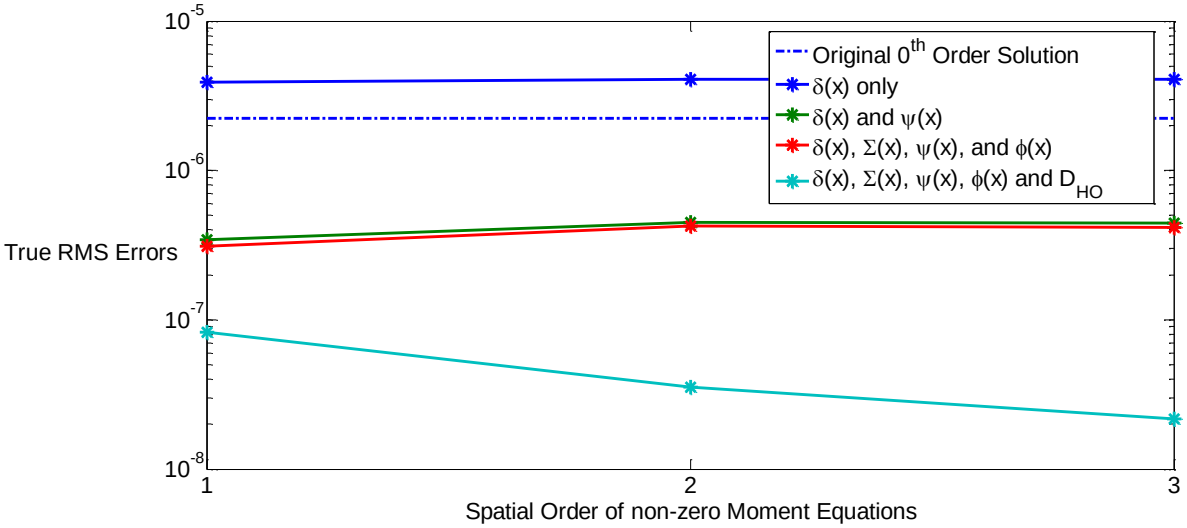

Figure 12 : Comparison of group 1 RMS errors (y-axis) for different spatial order transport calculations (x-axis) using various spatial approximations to represent the cross sections

In Figure 12, it is found that the best and most consistent reduction in group 1 errors is achieved when we include all the spatial dependences of our cross sections and when using the full definition of the *D* matrix from our 1D HOMOC. This makes sense since local errors in the the fast fluxes will have a greater impact across the rest of the core due to the longer mean free paths of fast neutrons. Also, the mean free path is larger than the average cell size, allowing better resolution of spatially dependent behavior. Therefore, including all spatial information into the recondensation procedure will produce better agreement in fluxes across the core and lead to much lower core averaged RMS errors.

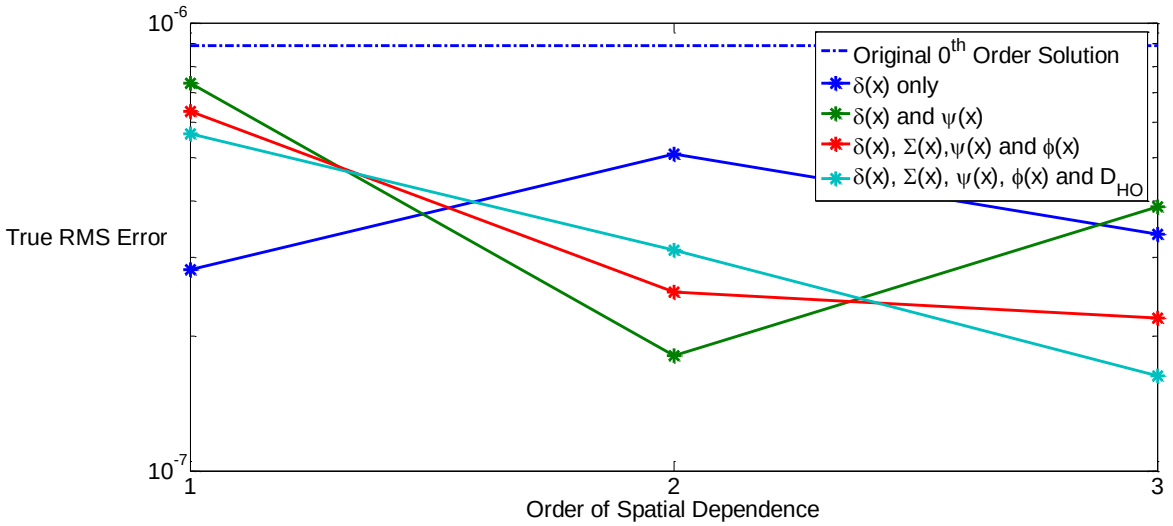

Figure 13 : Comparison of group 47 RMS errors (y-axis) for different spatial order transport calculations (x-axis) using various spatial approximations to represent the cross sections

For group 47 fluxes, however, the behavior is quite different. For all 4 of the tested combinations the original solution was improved. This is likely because the local errors do not propagate very far from their originating cells and the mean free path of the group 47 flux is smaller than the size of the current mesh. All things considered, it is apparent that the case taking into consideration the most spatial information is the only consistent procedure.

Therefore, this work has shown that high order spatial DGM equations can be used to provide corrections to our 0^{th} order coarse group equation and drive the coarse group solution closer to the true 0^{th} order fine group solution. Unfortunately, doing so in a consistent manner requires use of all spatially dependent material parameters and fluxes, making our coarse group equations no longer 0^{th} order in the conventional sense. The high order spatial equations are required for the "0^{th}" order coarse group solve in order to incorporate the spatial information of the del term. Since the cost is similar to that of improving consistency for the higher order spatial methods, this spatial recondensation does not appear well suited for low order spatial method.

3.8 Memory Considerations for Spatial Recondensation

While these methods have demonstrated the ability of high order spatial recondensation to correct for the spatial inconsistencies in DGM, they are still inadequate. Increasing the spatial order of the angular flux expansion requires increased memory storage. When these methods are applied 2D and 3D geometries, the number of gauss points will roughly increase as the order squared and cubed respectively. This will put an enormous burden on the memory of a given system. However, moving towards whole core analysis, this procedure leads to major memory issues that need to be addressed.

Since computational power is generally considered to be cheaper than storage capacity, storage is typically replaced by on the fly calculations to reduce these memory requirements. Therefore, a redefinition of the material properties for DGM is conducted to replace storage with on the fly calculations. This is done by joining together two previously separate steps of the recondensation process. The first step is the explicit reconstruction of the fine group fluxes after solving the DGM equations and the second is the generation of the new material properties using the reconstructed fine group fluxes. Since the reconstruction process itself does not change, this calculation can be explicitly placed within the generation of the new material properties. The del term can be redefined according to Equation (3.31).

$$\begin{bmatrix} \delta_{0,g}(x,\mu) \\ \delta_{1,g}(x,\mu) \\ \vdots \\ \delta_{N-2,g}(x,\mu) \\ \delta_{N-1,g}(x,\mu) \end{bmatrix} = \frac{1}{\psi_{0,g}(x,\mu)} \left(\overline{\overline{\Sigma_{T,g}}} - \overline{\overline{\mathbb{I}}} \Sigma_{T,0,g}(\mu) \right) \begin{bmatrix} \psi_{0,g}(x,\mu) \\ \psi_{1,g}(x,\mu) \\ \vdots \\ \psi_{N-2,g}(x,\mu) \\ \psi_{N-1,g}(x,\mu) \end{bmatrix} \tag{3.31}$$

Here, $\overline{\overline{\Sigma_{T,g}}}$ is defined as the total cross section mass matrix. It is formed using the fine group total cross sections for a given material and the discrete basis functions used in DGM. In this summation, a_j is the normalization constant for the j^{th} discrete basis function. This mass matrix can be thought of as mapping the contribution of the j^{th} flux moment to the i^{th} cross section moment.

$$\overline{\overline{\Sigma_{T,g}}} = a_j \sum_{K \in g} P_i(K,N)\Sigma_{T,g}(K)P_j(K,N)$$

This redefinition of the del term requires only the total cross section mass matrix and the coarse group total cross sections be stored to calculate the del terms on the fly. The mass matrix only needs to be stored for each coarse group and material type; resulting in negligible storage requirements (for the whole core problem the maximum storage would be approximately 3 MB). The major storage requirement therefore comes from the coarse group total cross section (10 groups assumed), which is stored for each cell, angle and coarse group. This shifts the storage requirement on del to the angular fluxes.

The fission and scattering cross sections can be redefined in a similar fashion to minimize storage requirements. For these material properties, only their respective mass matrices and vectors need to be stored for each material type, requiring storage of roughly a few MB for the same example. For these definitions, only the scalar fluxes need to be stored, leading to significantly less memory consumption relative to the del term.

The spatial del term is still the most memory intensive property in the entire calculation. For this term, it is necessary to store the angular flux values for each moment within each coarse group (the total is the same as the total number of fine groups), for each angle, for each cell and each spatial moment. To put this in perspective, imagine a 2D whole core geometry which spatial recondensation is applied to. This core is discretized into a million cells, each of which contain 3 spatial moments for a 1st order spatial problem. For the transport calculation 32 total angles are swept for each of the 200 fine groups (a total of 200 DGM moments). The required storage for all the del terms (assuming double precision) for such a problem would be roughly 150 GB. For the fine group problem in this case, the highest memory requirement is the scalar flux, which only needs to be stored for each cell and fine group, leading to 1.5 GB of storage in the case of considering only isotropic scattering. This would certainly increase if anisotropic scattering were taken into account. Increasing the spatial order of calculation also causes the number of

spatial moments to increase quadratically in 2D as opposed to linearly in 1D, further increasing storage requirements.

In addition, the angular dependence may also be represented in terms of continuous basis functions, which would further reduce memory requirements but could lead to an increase in the errors introduced into the recondensation process.

3.9 Summary

To our knowledge, this is one of the first successful attempts at placing local spatial dependence in the the multigroup collapse procedure. Using numerical integration with the Gauss-Legendre quadrature, high order spatial cross sections can be coupled to a high order spatial method. At this point, only 1D testing has been conducted with this approach. However, the numerical integration techniques used in our 1D example are already widely used in solving PDEs using finite element methods. Therefore, this work is able to provide much insight into conducting the multigroup collapse procedure in a consistent manner when working with a finite element representation of a 2D core geometry.

For our purposes, these spatially dependent cross sections have been successfully used in DGM equations to drive the converged flux moments closer to the fine group solution. However, the issue of spatial consistency has only been partially addressed because the spatial recondensation approach is only asymptotically consistent. Higher and higher spatial order representations of the angular fluxes can be used, but full consistency between DGM and the fine group problem can't be feasibly achieved with this approach.

There is also the very real problem of storing the angular fluxes. In order for this work to be competitive, memory costs ideally would be lower than that of the fine group problem. At this point, the memory requirements are at least two orders of magnitude greater in term of storage requirements when only considering isotropic scattering in the fine group problem. On top of that, there is currently no consistent equivalent of 1D HOMOC in the arena of 2D solvers. The best choice moving forward with this approach would be implementation into a finite element based neutron transport code. However, the goal of this thesis was to make DGM and

recondensation applicable to a wide range of methods and problem. In order to avoid being limited to finite element methods, it was necessary to find a new way to resolve the inconsistencies in DGM without resorting to the use of high order spatial cross sections.

4 Exact Spatial Recondensation

In all previous work, the DGM equations were derived using the integro-differential form of the neutron transport equation. Typically this is the best place to start because this is our primary source for ensuring conservation of neutrons. However, in this particular case, the assumption that no particular spatial method has been applied limits the current DGM formulation. The result is a del term that neglects the assumptions of the spatial method applied after formulating the DGM equations. While the original del term becomes generalized and can technically be applied to any spatial method, full consistency is only observed for 1st order methods assuming a flat angular flux. Therefore, instead of forming a generalized del term, we decided to form a del term specifically for MOC by conducting the spatial integration on our transport equation first and then derive a set of DGM equations for this method.

4.1 Derivation of Exact Recondensation for 1D MOC

Instead of starting with the fine group equations in their integro-differential form, we assume that MOC has been applied. This provides two equations that govern neutron conservation along a single 1D track superimposed on a given problem.

$$\begin{aligned} \psi_{out,g}(K) &= A_g(K)\psi_{in,g}(K) + B_g(K)Q_g(K) \\ \overline{\psi_g}(K) &= \frac{\mu}{l}\left(B_g(K)\psi_{in,g}(K) + C_g(K)Q_g(K)\right) \end{aligned} \tag{4.1}$$

In the first part of Equation (4.1), the incoming angular flux and the flat source are multiplied by respective coefficients, *A* and *B* to determine the outgoing angular flux. The second equation uses coefficients, *B* and *C*, along with the incoming angular flux and the cell averaged source to determine the cell averaged angular flux for fine group *K* within coarse group *g*. [15] The notation denoting angle and segment has been dropped to maintain clarity.

If we want to apply DGM consistently, we need to form a del term that can reproduce the exact outgoing and average fine group angular fluxes while instead using the the coarse group total

cross section to define the *A*, *B*, and *C* coefficients. While the incoming-outgoing angular flux relation is directly related to the cell averaged angular flux relation, both are shown to highlight the need for two different correction terms to be introduced since difference streaming coefficients are used between the two equations in Equation (4.2).

$$\begin{aligned} \psi_{out,g}(K) &= A_g \psi_{in,g}(K) + B_g Q_g(K) + \Delta_g^{io}(K) \\ \overline{\psi_g}(K) &= \frac{\mu}{l}\left(B_g \psi_{in,g}(K) + C_g Q_g(K)\right) + \Delta_g^{avg}(K) \end{aligned} \tag{4.2}$$

These correction factors, $\Delta_g^{io}(K)$ and $\Delta_g^{avg}(K)$, will be called del sources, since they act like source corrections to the MOC equations. These del sources can be defined by setting Equations (4.2) and (4.3) equal to each other and then subtracting the necessary quantities. The del sources are now defined according to Equation (4.3).

$$\begin{aligned} \Delta_g^{io}(K) &= \delta A_g(K) \psi_{in,g}(K) + \delta B_g(K) Q_g(K) \\ \Delta_g^{avg}(K) &= \frac{\mu}{l}\left(\delta B_g(K) \psi_{in,g}(K) + \delta C_g(K) Q_g(K)\right) \end{aligned} \tag{4.3}$$

Now, instead of using the differences in the total cross sections to define the del term, the del term is now calculated using the difference in the fine group streaming term and the coarse group streaming term.

$$\begin{aligned} \delta A_g(K) &= A_g(K) - A_g \\ \delta B_g(K) &= B_g(K) - B_g \\ \delta C_g(K) &= C_g(K) - C_g \end{aligned} \tag{4.4}$$

Now a new set of moment equations are obtained by multiplying the modified fine group equations by the discrete basis functions and then sum over all fine groups *K* within coarse group *g*. This leads us to the following equations defining the outgoing and cell averaged angular flux moments.

$$\psi_{out,i,g} = A_g \psi_{in,i,g} + B_g Q_{i,g} + \Delta_{i,g}^{io}$$
$$\overline{\psi_{i,g}} = \frac{\mu}{l}\left(B_g \psi_{in,i,g} + C_g Q_{i,g}\right) + \Delta_{i,g}^{avg} \tag{4.5}$$

These sets of equations are unsuspectingly difficult to actually solve. This is due to the definition of the corrective terms on the right hand side. The terms are defined by the action of a set of coefficient matrices on all the incoming angular flux moments for that particular coarse group.

$$\Delta_{i,g}^{io} = \overline{\overline{\delta A_{g,\imath J}}}\,\overrightarrow{\psi_{\imath n,J,g}} + \overline{\overline{\delta B_{g,\imath J}}}\,\overrightarrow{Q_{J,g}}$$
$$\Delta_{i,g}^{avg} = \frac{\mu}{l}\left(\overline{\overline{\delta B_{g,\imath J}}}\,\overrightarrow{\psi_{\imath n,J,g}} + \overline{\overline{\delta C_{g,\imath J}}}\,\overrightarrow{Q_{J,g}}\right) \tag{4.6}$$

The coefficient matrices are produced through application of the discrete basis function on the differences between the fine group and coarse group streaming coefficients.

$$\overline{\overline{\delta A_{g,\imath J}}} = a_j \sum_{K \in g} P_i(K,N)\delta A_g(K) P_j(K,N)$$
$$\overline{\overline{\delta B_{g,\imath J}}} = a_j \sum_{K \in g} P_i(K,N)\delta B_g(K) P_j(K,N) \tag{4.7}$$
$$\overline{\overline{\delta C_{g,\imath J}}} = a_j \sum_{K \in g} P_i(K,N)\delta C_g(K) P_j(K,N)$$

These equations are almost in the same form as the original DGM equations now. The cell averaged source moment is still defined by the scattering cross section and chi moments calculated during the recondensation process according to the work conducted by Zhu and Forget. [41] The only difference in these equations is that the del source is now defined in terms of the incoming angular flux moments and cell averaged source moments. This is still not quite a DGM method, since these sources still contain couplings between all equations *i* and *j* within

the same coarse group. Therefore, we divide the del source by the cell averaged angular flux and create a new del term that is fully consistent with the fine group MOC equations.

$$\begin{aligned} \psi_{out,i,g} &= A_g \psi_{in,i,g} + B_g Q_{i,g} + \delta_{\mathrm{i,g}}^{MOC,io} \overline{\psi_{0,g}} \\ \overline{\psi_{\imath,g}} &= \frac{\mu}{l_k}\left(B_g \psi_{in,i,g} + C_g Q_{i,g}\right) + \delta_{\mathrm{i,g}}^{MOC,avg} \overline{\psi_{0,g}} \end{aligned} \tag{4.8}$$

Our new set of DGM equations defined in Equation (4.8) is very similar to the original formulation, except that in the case where *i*=0, the new del term does not vanish as it did in the original DGM equations. This allows the del term to correct for the spatial errors without the need for additional high order spatial equations. Although we have a new term added to the *i*=0 equation, the sources for all the other higher order DGM equations still only depend on the coarse group solution. Therefore, this new formulation still allows us to conduct the power iteration on the leading order equation, *i*=0, and then do a single sweep on the rest of the DGM equations for fine group flux reconstruction. Once fine group reconstruction takes place, we can then recondense our cross section moments as before and use Equation (4.9) to update the del terms.

$$\begin{aligned} \delta_{\mathrm{i,g}}^{MOC,io} &= \frac{\overline{\overline{\delta A_{g,\imath J}}}\,\overrightarrow{\overline{\psi_{\imath n,J,g}}} + \overline{\overline{\delta B_{g,\imath J}}}\,\overrightarrow{\overline{Q_{J,g}}}}{\overline{\psi_{0,g}}} \\ \delta_{\mathrm{i,g}}^{MOC,avg} &= \frac{\frac{\mu}{l_k}\left(\overline{\overline{\delta B_{g,\imath J}}}\,\overrightarrow{\overline{\psi_{\imath n,J,g}}} + \overline{\overline{\delta C_{g,\imath J}}}\,\overrightarrow{\overline{Q_{J,g}}}\right)}{\overline{\psi_{0,g}}} \end{aligned} \tag{4.9}$$

We have now produced a new set of DGM equations built specifically for 1D MOC that enables the recondensation procedure to converge to the exact fine group solution.

4.2 1D Benchmark Results for Exact Recondensation

This new method is tested by using the true fine group solution from the 1D BWR problem to provide all the cross section moments and del terms for the DGM equations. The leading order DGM equation is used in the power iteration until the same convergence criteria used in the fine group calculation are satisfied. The higher order DGM equations are then solved and the fine

group fluxes are reconstructed. This process was conducted for consecutively finer meshes. If these equations are exactly consistent, then the errors remaining should be on the order of the convergence criteria and should not depend on the mesh size for the problem. If an inconsistency remains, then the errors should start at some level above the convergence criteria and drop as the spatial problem is refined.

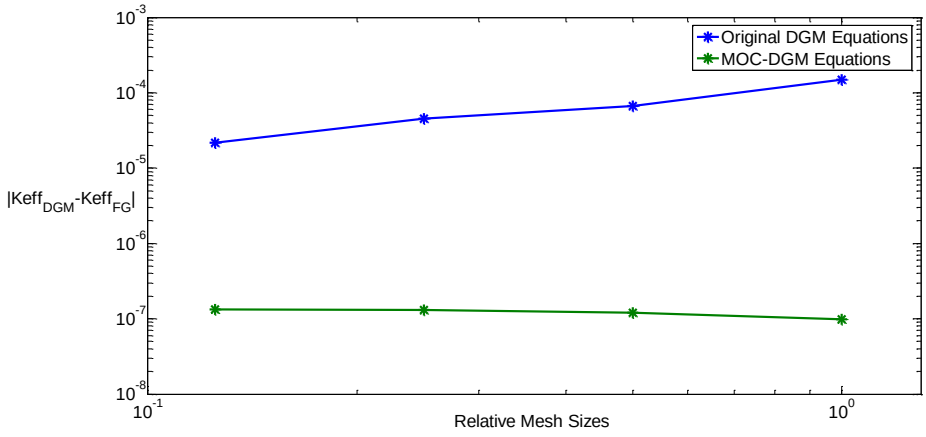

Figure 14 : Comparison of errors in k_{eff} using the original del term definition and exact definition for varying mesh size

From Figure 14, we see that while the original DGM solution slowly converges to the fine group solutions as the mesh size decreases, the MOC-DGM solution converges to the near exact solution. We also see very little variation in the magnitude of the persisting errors as the mesh size decreases, further showing that the MOC-DGM equations allow the recondensation process to converge (within the convergence criteria, of course) to the exact fine group solution without need for any additional high order equations.

4.3 Convergence Properties of Exact DGM

Up to this point, no effort has been made to optimize the convergence properties of the recondensation process. For all previous results, convergence of the recondensation process took roughly 100 iterations or more and required roughly the same number of transport sweeps as the full fine group problem. The lambda value for the Krasnoselskii iteration was set to a very low

value, no effort was made to optimize the choice of coarse group structure and the convergence criterion used on the coarse group eigenvalue problem was extremely tight regardless of how inaccurate the initial guess was.

The coarse group structure was modified according to the algorithm proposed by Gibson and Forget [11]. They suggest choosing the coarse group structure such that the ratio of the maximum and minimum fine group total cross sections within that coarse group remains below 2. This selection process is conducted one group at a time for all materials present in the example core. Each material will likely have a different choice for the first coarse group, so the smallest one is chosen to ensure this ratio is maintained across all materials in the core. This process begins again until the whole fine group domain has been distributed between the coarse groups. For the 1D BWR geometry, a 6 coarse group model was able to roughly satisfy this constraint using 4, 4, 7, 27, 3 and 2 fine groups for each of the respective coarse groups. This new coarse group structure allows us to use a λ of 0.9 for the Krasnoselskii iterations.

4.3.1 Choosing the Eigenvalue Convergence Criteria

In addition to choosing a different coarse group structure, we can also optimize how we solve the coarse group problem within recondensation. Most of the work using DGM in recondensation supposes that fixing the convergence criteria for the coarse group to that used for the fine group calculation is the best option. Unfortunately, if a poor initial guess is used to initialize recondensation and a tight convergence criterion is used, this will cause an unnecessary amount of work to be placed on converging the coarse group problem when the cross sections themselves may be way off. Ideally, an algorithm needs to be put in place that can anticipate the convergence criteria needed to improve the coarse group solution enough so that source moments for our DGM solve are improved.

Therefore, the recondensation procedure was modified such that the convergence criteria placed on the coarse group eigenvalue problem could change as the convergence of the reconstructed fine group fluxes improved. In this work, the convergence criteria were set to the maximum change in the reconstructed fine group fluxes between recondensation steps and then divided by 10. The idea behind setting this value is that one would be setting their coarse group calculation

to converge the fluxes just enough so that the DGM solve will reconstruct a better fine group flux next time. For these calculations, the convergence criteria set for the outer iteration (the recondensation iteration) is 10^{-8}.

$$\varepsilon_{k_{eff}}, \varepsilon_{\phi} = \frac{\left\| \phi_g^r(K) - \phi_g^{r-1}(K) \right\|_{\infty}}{10} \tag{4.10}$$

This approach was taken to affect the recondensation process in two major ways. First, it reduces the computational burden placed on each eigenvalue solve by limiting the number of transport sweeps used to converge the coarse group fluxes. This is achieved by assuming that the maximum difference in the fluxes between the previous two recondensation steps provides a decent estimate of the error remaining in the recondensation process. Since there is not enough information to inform this method at the beginning, the criteria are set to 10^{-3} for the first two recondensation steps. Second, dividing the maximum difference between recondensation steps ensures that the improvements between recondensation steps are not held back because of inaccurate eigenvalue solutions.

This algorithm was tested using the 1D BWR example with a flat initial. Since this initial guess is very poor, the value of λ for the first recondensation step was set to 0.3 to ensure stability and then increased to 1.0 for all other steps. A plot of the number of power iterations applied for each recondensation step provides a good indication of whether the desired goal has been achieved.

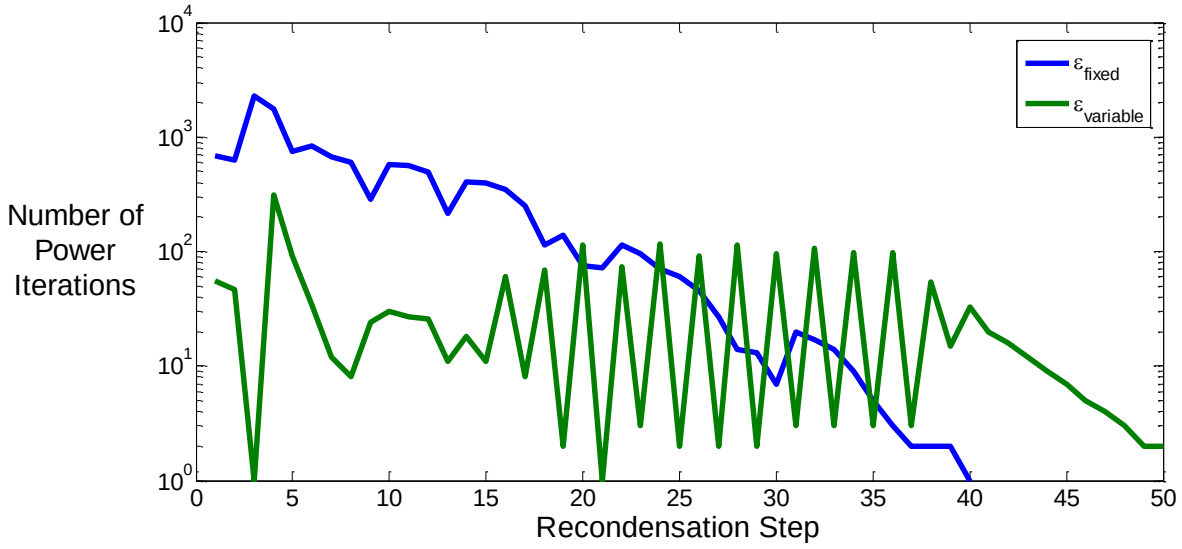

Figure 15 The number of power iterations used at each recondensation step for the fixed and variable convergence criteria case

According to Figure 15, this method does provide a significant reduction in the total number of power iterations applied to the coarse group solve and a more even distribution of the power iterations throughout recondensation. However, these improvements do come at the cost of an increase in the number of recondensation steps needed to converge the outer iteration. It is important to remember that the true measure of improvement is in the total number of transport sweeps used through this process. For the fixed convergence criteria case, the total number of transport sweeps required converge to the fine group solution is 77,744. This is almost twice as much as the original fine group problem which required roughly 44,000 transport sweeps. This is the predominant reason why recondensation has not been considered as a stand-alone method for accelerating convergence to the fine group solution. However, when the variable convergence criteria are used, the recondensation process sheds a significant fraction of the transport sweeps, reducing to 14,063.

While this algorithm does work, it is apparent that it is not ideal. This method suffers from an instability that occurs when a small change in the reconstructed fine group flux occurs early on in the recondensation process. Due to the non-monotonic convergence of the fluxes, the change between subsequent steps can be much smaller than the actual error at the current iterate. Any

small changes between steps will cause the convergence criteria for the following eigenvalue calculation to be set unnecessarily tight. For example, let's look at the case when a small change results in a difference of 10^{-6} in the fine group fluxes even though the recondensation process is still 10^{-3} away from the true solution. Using Equation (4.10) the convergence criteria for the next eigenvalue calculation will be set to 10^{-7}. This significantly increases the number of power iterations used in the following coarse group eigenvalue solve. Such a large emphasis on the eigenvalue problem will likely lead to big changes in the reconstructed fluxes for the next step and produces a loose convergence criterion for the next eigenvalue solve. This loose criterion results in very few power iterations in the next eigenvalue solve, leading again to a small change in the reconstructed flux. Eventually, though, this oscillation is damped because the recondensation process is still driving the fluxes toward the fine group solution.

To dampen the oscillations occurring using Equation (4.10), another scheme is introduced that uses a weighted average of the previous two recondensation steps.

$$\varepsilon_{k_{eff}}^{k+1}, \varepsilon_{\phi}^{k+1} = \frac{\omega \left\| \phi_g^k(K) - \phi_g^{k-1}(K) \right\|_\infty + (1-\omega) \left\| \phi_g^{k-1}(K) - \phi_g^{k-2}(K) \right\|_\infty}{\beta} \quad (4.11)$$

It can be seen that Equation (4.10) is recovered when we set β=10 and ω=1. To evaluate the impact this has on the eigenvalues solves, we compared the total number of power iterations at each recondensation step for various values of ω.

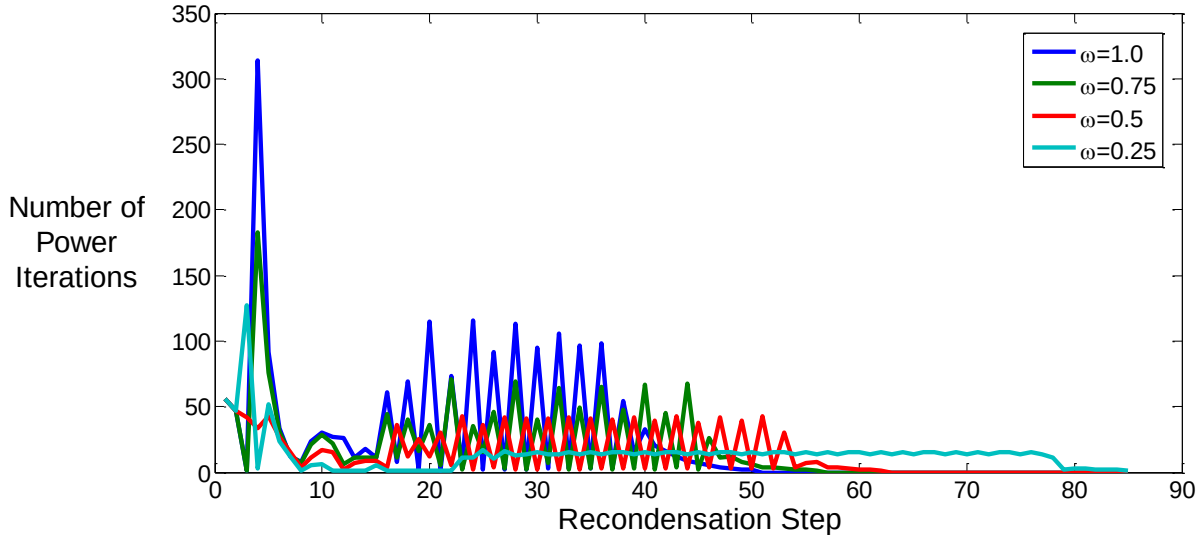

Figure 16 The number of power iterations used at each recondensation step for various values of omega

In Figure 16 we see that using Equation (4.11) can better distribute the number of power iterations used over the course of the recondensation process. However, we must keep in mind that placing too little weight on the power iterations at each recondensation step will shift the focus to our DGM solves and significantly increase the number of transport sweeps. Therefore, we must compare the total number of transport sweeps required to converge the outer iteration in order to establish the true performance of this method. We also compare these results when using different factors of β in Equation (4.11).

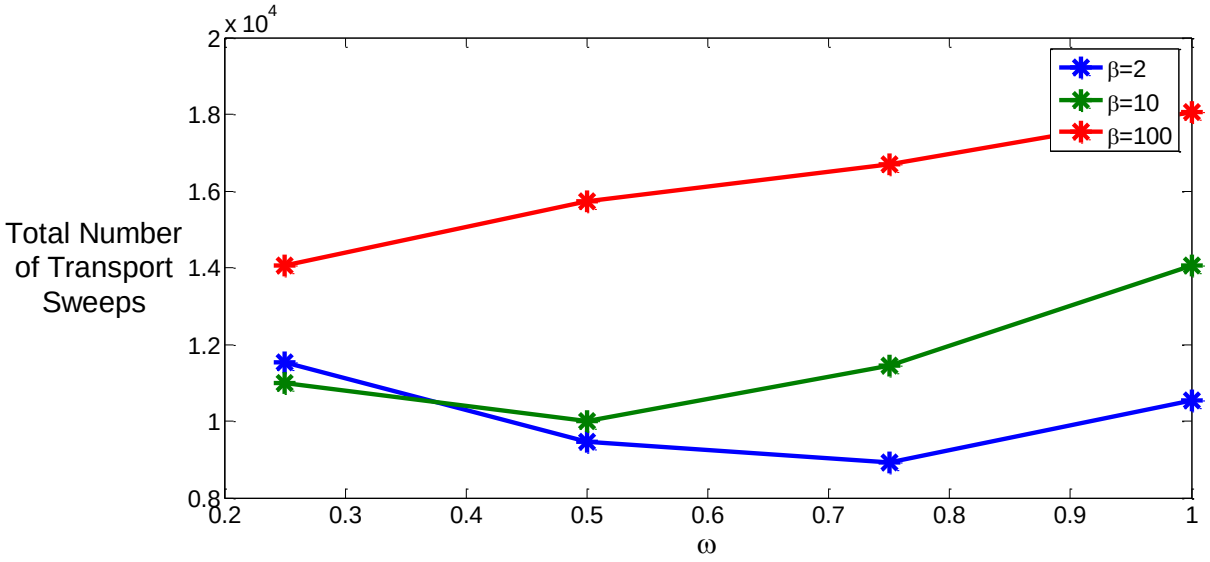

Figure 17 Comparison of total number of transport sweeps for various parameters used in Equation (4.11)

As expected, when we tighten the convergence using a high value for β, here we use 100, the total number of transport sweeps substantially increases. This makes sense because the convergence criteria placed on the eigenvalue solve gets set very tightly and therefore too much effort is placed on converging the eigenvalue at each recondensation step.

When we reduce β to 10, the total number of transport sweeps decrease. This is because we are now converging the eigenvalue to a value that is much closer to the error remaining between the current iterate and the fine group solution. The number of sweeps further decreases as ω is reduced until reaching a minimum at 0.5. Further decreasing ω sets the convergence criteria to the maximum difference two recondensation steps before, placing too lax of a constraint on the eigenvalue for the current coarse group calculation. This requires more recondensation steps and therefore more DGM solves.

Another decrease is observed when we set β to a very low value, in this case 2. The total number of sweeps drops beneath 10,000 and reaches a minimum when ω is 0.75. This is curious because one would expect that the transport sweeps would increase for such a low β value because more emphasis is being placed on solving the DGM equations than the coarse group eigenvalue

problem. Therefore, the distribution of power iterations is analyzed in Figure 18 when setting β=2 to better understand this behavior.

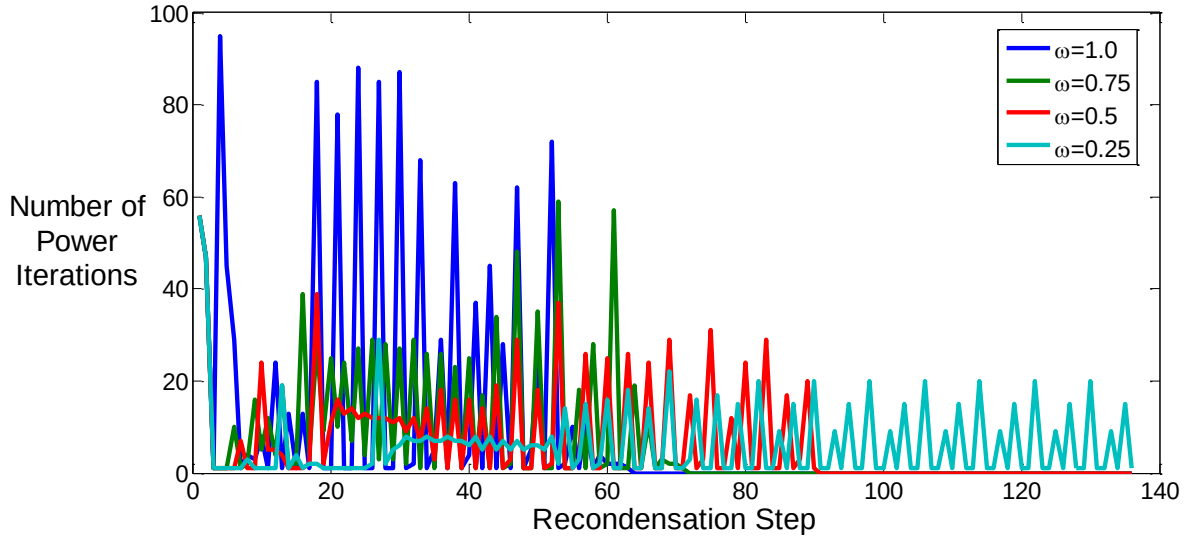

Figure 18 Comparison of the number of power iterations per recondensation step for various values of ω for β equal to 2

One of the consequences of setting the convergence criteria too low is an increase in the instability observed before Equation (4.11) was applied. While changing the value of ω does seem to reduce the oscillations, it is not nearly as effective as it was in the β=10 case. These oscillations also have an effect on the convergence in k_{eff} between recondensation steps. When comparing the β=2 and ω=0.75 case to the β=10 and ω=0.25 case in Figure 19, we see that while the β=2 case does converge faster overall, the convergence behavior is much more oscillatory throughout the recondensation process. It is not readily apparent that this behavior is worrying, though. Recondensation is still being driven to the fine group solution regardless of how the power iterations are distributed. While divergence of the solution due to this behavior has not yet been observed, further work should be conducted to determine if these oscillations can indeed lead to divergence in the outer iteration. Optimal choice of these factors will likely be problem dependent base since the appropriate choice of omega will depend on how quickly the recondensation process converges and beta will depend on how expensive the higher order DGM sweeps are relative to the coarse group sweeps.

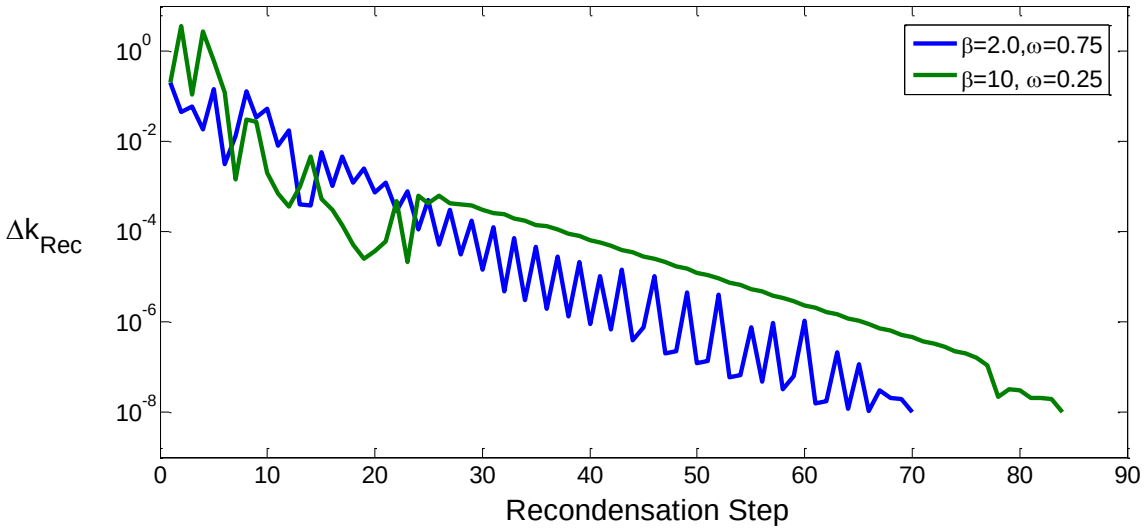

Figure 19 Difference of eigenvalues between consecutive recondensation steps

It is also important to understand how the initial guess affects recondensation's convergence behavior as well. For all of these cases so far we have assumed that the flux is flat both in space and in energy. This provides an extremely poor initial guess to start the recondensation process. For a realistic problem, though, we will likely have cross sections and fluxes from a previous assembly level calculation to work with, so for the next set of tests we will be using these fluxes as an initial guess. There is no guarantee that we have all the incoming angular fluxes to all the cell faces from the assembly calculation to calculate our del terms, so the cell-averaged angular flux from an adjacent cell is used instead as the incoming angular flux. While this is an approximation that is not consistent with step characteristics, it should capture much of the spatial effects from the assembly calculation needed to make the initial guess better.

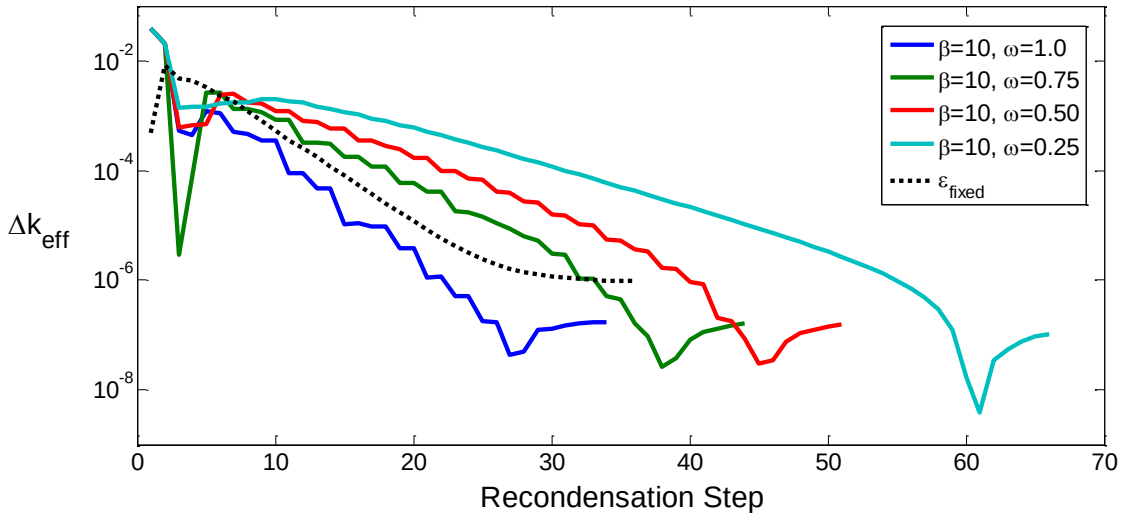

Figure 20 Comparing convergence profiles when initial criteria set to 10^{-3}

In Figure 20, we see that choosing loose convergence criteria for the beginning of recondensation introduces a significant amount of oscillation in the convergence especially when $\omega = 0.5$ or greater. Interestingly, the overall convergence profile shifts closer to the fixed eigenvalue case and when $\omega = 1.0$ the method converges faster than the fixed case. One should remember, though, that since we are starting from better approximation we should also be using tighter convergence criteria from the beginning. Any poorly converged solutions at the beginning of recondensation may be causing error cancellations and lead to false acceleration. To test this hypothesis, we ran these cases again but this time setting the initial convergence criteria to 10^{-4}.

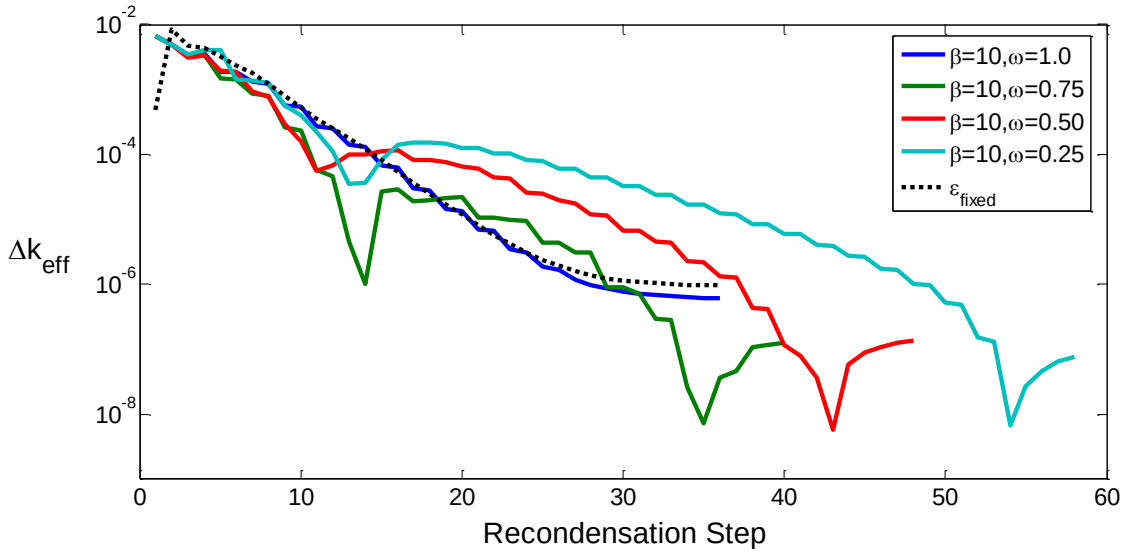

Figure 21 Comparing convergence profiles when initial criteria set to 10^{-4}

In Figure 21, we observe that using tighter convergence criteria reduces the number of recondensation steps it takes to converge. In fact, we see that when ω=1.0 the convergence profile almost exactly matches the fixed convergence criteria case. When comparing both Figure 20 and Figure 21, it becomes apparent that the non-monotonic convergence is really an oscillation about the convergence profile for the fixed criteria case. This oscillation appears to be dampened when the initial convergence criteria is tightened. If this is the case, then further tightening the criteria should result in better agreement with the fixed criteria case.

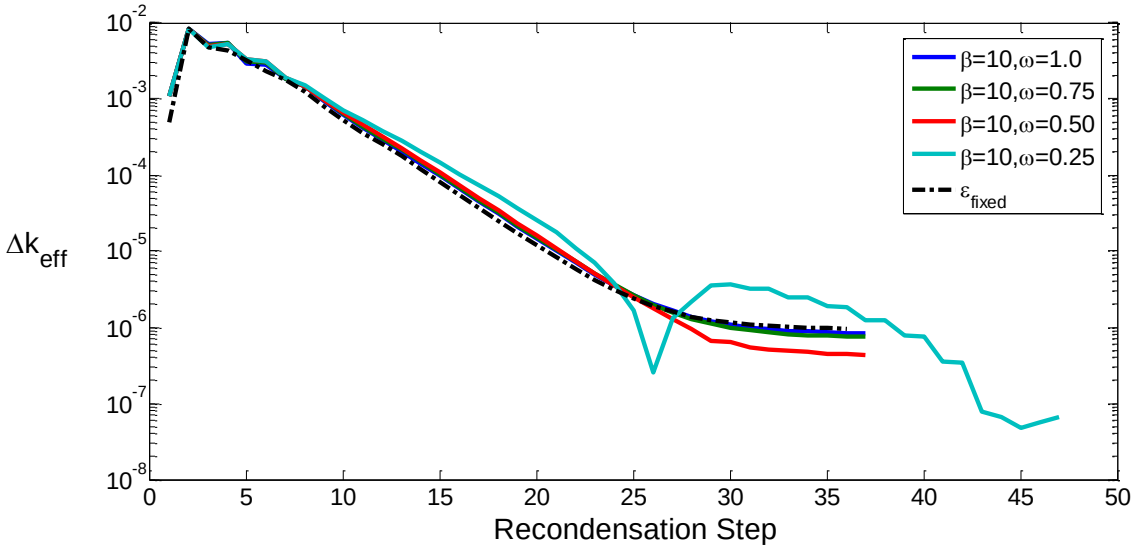

Figure 22 Comparing convergence profiles when initial criteria set to 10^{-5}

In Figure 22 we see that the convergence profiles match much closer to the fixed criteria case. As in the previous cases, much better agreement with the fixed criteria case is observed when the value of ω increases. We have now shown that errors from using loose convergence criteria in the first couple eigenvalue solves can propagate through the recondensation process and lead to poor convergence properties in recondensation. Therefore, it is important to keep in mind how well the first couple coarse group eigenvalue problems are converged.

While it is certainly important to make sure we are optimally converging in the outer iteration of recondensation, we must also look at the total number of transport sweeps recondensation applied to converge the solution. For all the cases shown in Figure 20, Figure 21 and Figure 22, we plot the dependence of the total number of transport sweeps on ω and the initial convergence criteria.

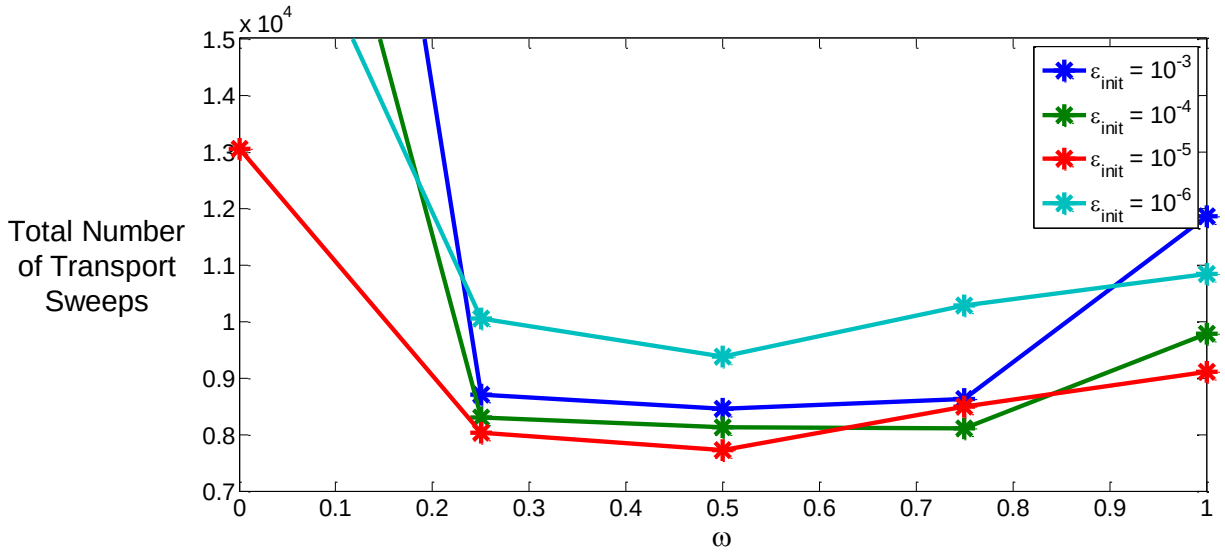

Figure 23 The effect that omega and the initial convergence criteria have on the total number of transport sweeps used in recondensation

We see that the minimum in all cases is still achieved between $\omega = 0.5$ and 0.75, although it seems that placing more emphasis on the initial convergence criteria actually reduces the total number. This makes sense in a way because when the initial criteria produced a poorly converged coarse group solution, this brought more errors on top of those the multigroup approximation introduced and resulted in the oscillatory behavior. So, in the case where the initial criteria were set to 10^{-3}, convergence required more outer iterations and therefore required more DGM solves. When the initial criteria were tightened, the errors leading into recondensation were much smaller. This reduced the oscillations observed previously and brought the convergence profile much closer to the case when the criteria were fixed at 10^{-8}. The effect was a reduction in the number of DGM solves. While the Krasnoselskii iteration is being used, the fixed source sweep across all high order DGM equations requires roughly a fine group number of sweeps over the problem, therefore a decrease in the number of recondensation steps leads to a significant reduction in the total number of transport sweeps. When compared relative to the fixed tolerance case, on the order of 77,000, the total number of transport sweeps has been reduced by over a factor of 7 in almost all cases.

While this method for choosing the convergence criteria at each recondensation step is shown to work well, it is by no means the only possible approach. Theoretically, one could use any L-norm to set the convergence criteria or apply the L_∞ norm to the cross section values at different iterates instead. There are many possibilities that could lead to further improvement, but this work shows that the L_∞ norm defined in Equation (4.11) provides a good estimate in a quick and cheap manner.

4.3.2 Defining a Natural Del Term

One of the interesting qualities of the exact del term is how flexible it is in application. This is best seen when we look back at how it was first derived. We started with the original fine group equations with our spatial method already incorporated, but then stated that we wanted to use the coarse group streaming operator instead of those used for the fine group problem. We then had to calculate the difference between the two sets of operators and then transform these differences using the discrete basis functions. This is what led to the formation of the coefficient matrices used to calculate the del term. From the definitions in Equation (4.7), the coefficient matrix can be expanded according to Equation (4.12) in terms of the fine group quantities and coarse group quantities.

$$\overline{\overline{\delta A_{g,\iota J}}} = a_j \sum_{K \in g} P_i(K,N)\left(A_g(K) - A_g\right)P_j(K,N) \tag{4.12}$$

After applying the summations two matrices are produced. The first is dense and corresponds to the transformed set of fine group streaming operators. The other matrix is diagonal and contains the values of coarse group streaming operator. This same expansion can be applied to the coefficient matrices for *B* and *C*, as well.

$$\overline{\overline{\delta A_{g,\iota J}}} = \overline{\overline{A_{g,\iota J}}} - diag\left(A_g\right) \tag{4.13}$$

The purpose of these coefficient matrices is to map the contributions of all the flux moments within a coarse group to the overall correction needed for that particular moment, angle and spatial cell to drive the recondensation process to the fine group solution. Since it is extremely unlikely that the diagonal of the fine group streaming matrix is the same as that for the coarse

group, the coefficient matrix ends up containing non-zero values along its diagonal. This means that at the current recondensation step, every del term contains information about the incoming angular flux and source for the same moment from the previous step. For example, if we look at the incoming-outgoing del term for DGM moment *i* and coarse group *g*, we see that the diagonal terms of the coefficient matrix provide information about the i^{th} moment of the incoming angular flux from the previous step. Take note that *r* denotes the recondensation step, not the iteration within an eigenvalue calculation.

$$\Delta_{i,g}^{io,r+1} = \overline{\overline{\delta A_{g,\iota J}}}\,\overrightarrow{\psi_{in,J,g}^{r}} + \overline{\overline{\delta B_{g,\iota J}}}\,\overrightarrow{Q_{J,g}^{r}} \tag{4.14}$$

At first, this observation does not seem to be of any significance, but if we insert this del term with its functional dependence into the incoming-outgoing equation, we observe something very interesting. For our purposes, we assume that the coarse group values have been sufficiently converged in the eigenvalue calculation going into the DGM solve.

$$\psi_{out,i,g}^{r} = A_g\psi_{in,i,g}^{r} + B_g Q_{i,g}^{r} + \delta_{i,g}^{io,r}\left(Q_{j\neq i,g}^{r-1}, \psi_{in,j\neq i,g}^{r-1}, Q_{i,g}^{r-1}, \psi_{in,i,g}^{r-1}\right)\bar{\psi}_{0,g}^{r} \tag{4.15}$$

The first thing we notice in Equation (4.15) is that the del term for moment *i* contains information about the sources and incoming angular fluxes for all moments not equal to *i*. These dependences are very important because, although these values are taken from the previous recondensation step, they provide the necessary spectral information to improve the solution. Their incorporation into the del term is also what decouples all the flux moments from each other so that the moments can be solved independent of one another.

The second observation is that the current outer iteration on the outgoing angular flux also depends on the incoming angular flux and source for moment *i* from this iteration and from the previous iteration. This dependence is very odd because, in a typical transport sweep, the outgoing angular flux only depends on the current iterate of the incoming angular flux and source. Therefore, since we're now using a combination of the current and previous iteration, the convergence of the solution very well could be lagging behind what it could achieve otherwise. To test this hypothesis, a slight modification was made to the exact del term to

remove the dependence on the previous iterate but still provide the necessary spectral information from all the other moments.

The primary source of this problem is that the coefficient matrices have nonzero elements along their diagonals. This is because we assumed in our previous definition of the del term that we wanted to use the coarse group streaming operator for each of the moment equations. While transforming the equations in this fashion is much more intuitive and straightforward, there is no reason we can't use some other value entirely. By definition of the exact del term, we can use practically any value in place of the coarse group streaming operator and still maintain consistency with our fine group problem throughout the recondensation process.

To remove the redundancy in our del term definition, we apply the discrete basis functions to our fine group equations without first assuming that we want to use the coarse group streaming operators. This leads to the following set of equations for the outgoing angular flux and cell-averaged angular flux moments. We expand the matrix-vector form of the equation into its summation notation for the sake of clarity.

$$\begin{aligned} \psi_{out,i,g} &= \sum_{j\in g} A_{g,ij}\psi_{in,j,g} + \sum_{j\in g} B_{g,ij}Q_{j,g} \\ \overline{\psi_{\iota,g}} &= \frac{\mu}{l}\left(\sum_{j\in g} B_{g,ij}\psi_{in,j,g} + \sum_{j\in g} C_{g,ij}Q_{j,g}\right) \end{aligned} \tag{4.16}$$

To remove all the contributions of the i^{th} flux moment from the del term definition for that moment, the diagonal terms from the *A* and *B* matrices must be removed.

$$\begin{aligned} \psi_{out,i,g} &= A_{g,ii}\psi_{in,i,g} + B_{g,ii}Q_{i,g} + \sum_{j\neq i} A_{g,ij}\psi_{in,j,g} + \sum_{j\neq i} B_{g,ij}Q_{j,g} \\ \overline{\psi_{\iota,g}} &= \frac{\mu}{l}\left(B_{g,ii}\psi_{in,i,g} + C_{g,ii}Q_{i,g} + \sum_{j\neq i} B_{g,ij}\psi_{in,j,g} + \sum_{j\neq i} C_{g,ij}Q_{j,g}\right) \end{aligned} \tag{4.17}$$

By taking the diagonal terms out of the summation, we have defined a new streaming term for every flux moment equation whereas before we had only defined one term for all these moment

equations. While these streaming terms produce a set of equations that are DGM-like, they are markedly different because no group collapse action on the fine group total cross sections is required. The new streaming terms for use in our DGM equations are calculated using only the fine group streaming terms and the discrete basis functions without any weighting by angular or scalar fluxes. Since the spatial method has already been applied, conservation is no longer placed on a coarse group removal reaction rate *per se* but rather it is placed on the streaming and removal that occurs within a given cell.

The primary benefit of this new approach is that these streaming terms do not need to be changed after each recondensation step. Since this method uses the fine group streaming terms without any weighting by flux moments, no update of the coarse group total cross section is required. This "natural" del term can then be defined as follows.

$$\delta_{\mathrm{i,g}}^{io} = \frac{\sum_{j\neq i} A_{g,ij}\psi_{in,j,g} + \sum_{j\neq i} B_{g,ij}Q_{j,g}}{\overline{\psi_{0,g}}}$$
$$\delta_{\mathrm{i,g}}^{avg} = \frac{\frac{\mu}{l_k}\left(\sum_{j\neq i} B_{g,ij}\psi_{in,j,g} + \sum_{j\neq i} C_{g,ij}Q_{j,g}\right)}{\overline{\psi_{0,g}}} \quad (4.18)$$

Now, the del term for coarse group *g* and moment *i* only takes in information from all other moments not equal to *i* within the same coarse group. We first test this new definition of the del term using the fixed convergence criteria case and compare the results to the previous exact del term.

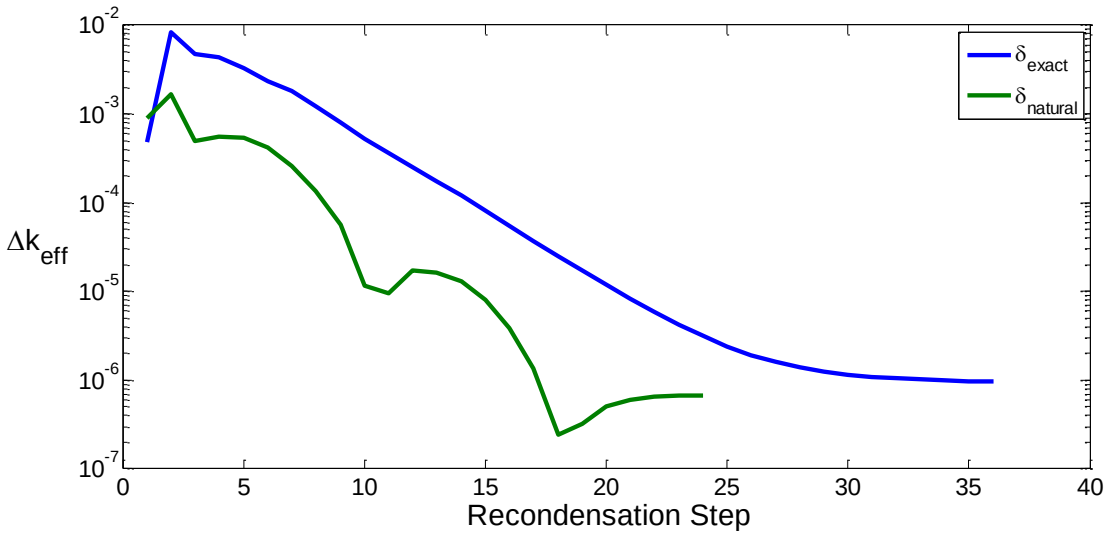

Figure 24 Comparing the convergence profiles of the natural del term and the previous del term

The resulting difference between the two definitions is very interesting. Using the natural del term instead of the previous exact definition reduces the number of recondensation steps by about a third. This definition does, however, introduce some oscillation to the convergence profile that was not there before. The reason for this non-monotonic convergence will be explained later in Section 5.1.4. While this oscillation is a drawback of this approach, this is outweighed by the fact that this natural del term dropped the total number of transport sweeps from 30,564 to 20,460. Since the reduction in total sweeps is directly proportional to that of the recondensation process, this tells us is that the speed up in convergence is driven solely by a decrease in the number of fixed source DGM sweeps and not in the number of coarse group sweeps.

Another noteworthy observation is that this does not appear to be an acceleration *per se*, since the average slope of the convergence profile for the natural del term roughly matches that of the previous case. Instead, what is occurring is a faster improvement in the reconstructed fine group solution at the beginning of the recondensation process. However, it is important to ensure that the initial decrease in Δk is not due to some error cancellation. Therefore, we also looked at the convergence profile for the core RMS errors for the 4 fine group fluxes in the first coarse group.

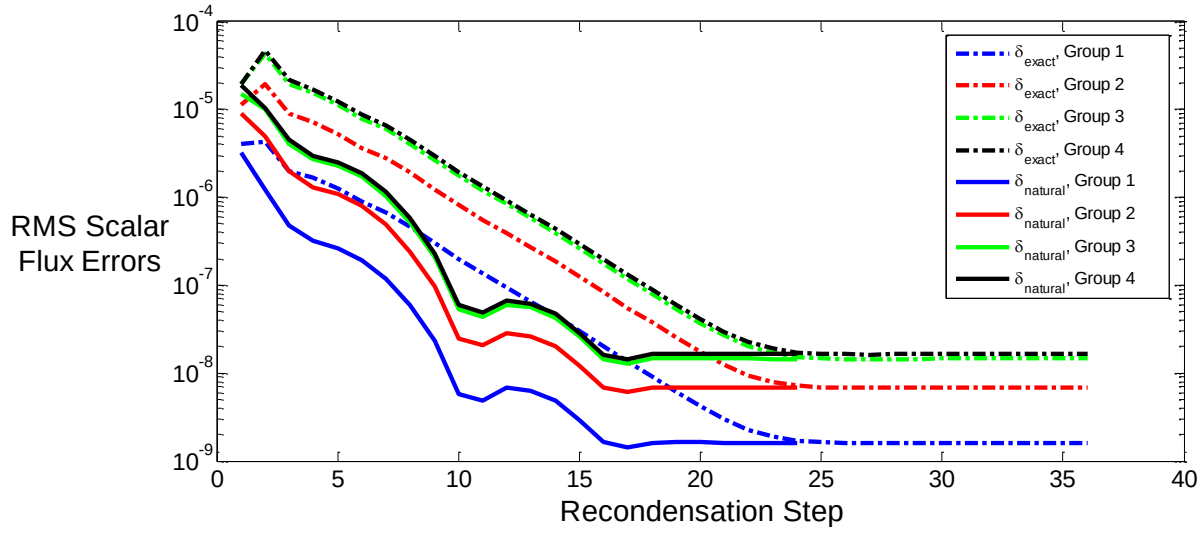

Figure 25 Comparison of the actual RMS scalar flux errors for the first 4 fine group fluxes using the exact and natural del term definitions

When we compare the errors in the fine group fluxes, we find that there is an initial drop in the error when using the natural del term. When we use the previous definition, the error actually increases after the first coarse group eigenvalue solve. Though we only show the results for the fine groups in coarse group 1, the fine group fluxes in the first 3 coarse groups all exhibit similar behavior. It's not until we reach the last 3 coarse groups that the convergence profiles begin to match a little closer to each other. Figure 26 shows the comparison for the 3 fine group fluxes within coarse group 5.

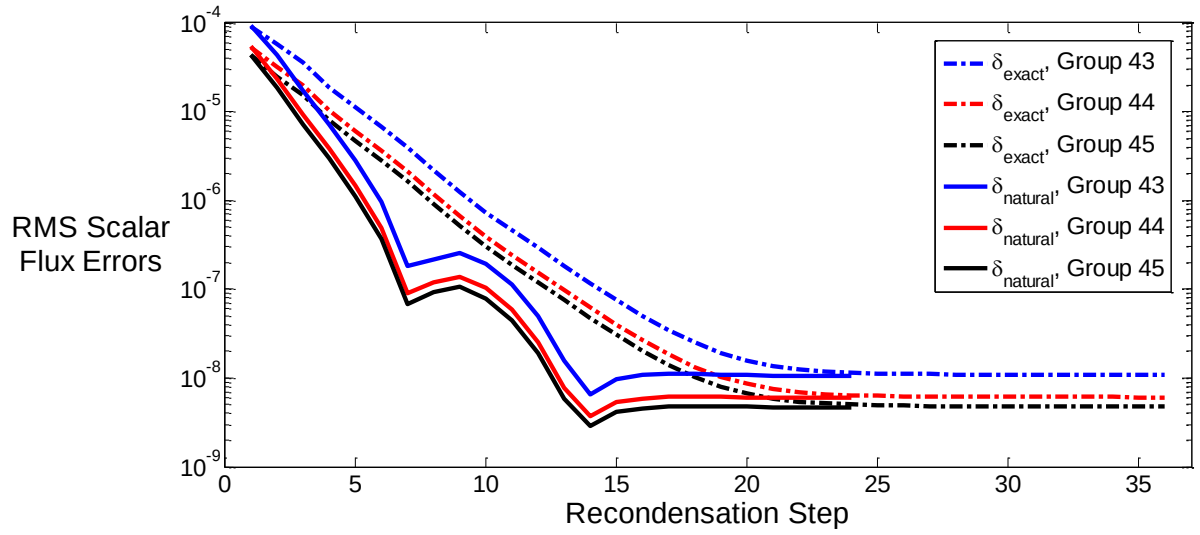

Figure 26 Comparison of the actual RMS scalar flux errors for coarse group 5 using the exact and natural del term definitions

Since this new definition showed so much promise in improving the convergence of recondensation, it seemed logical to evaluate the flexibility of the algorithm developed in Section 4.3.1 by testing it with this approach. Instead of using the fixed convergence criteria, we set the initial criteria to 10^{-5}, ω to 0.50 and β to 10. This choice of parameters provided a good compromise in terms of both minimizing the total number of transport sweeps while also closely matching the convergence profile of the fixed eigenvalue case. We use these parameters with the new del term and compare the results to assess the flexibility of this algorithm.

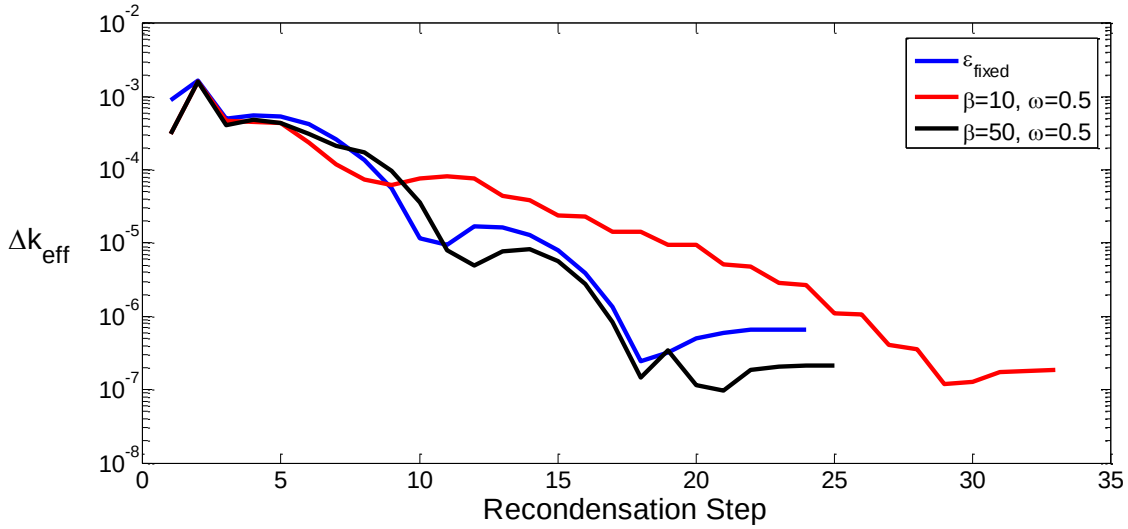

Figure 27 Comparison of the convergence profiles for the fixed and variable criteria cases

We see right away that the previous choice of parameters do not work so well with this new del term. Since the convergence properties of recondensation have changed, β=10 does not allow the convergence criteria to decrease fast enough. This causes the convergence profile to deviate substantially from the fixed criteria case. It also produces the oscillations in the number of power iterations at each recondensation step that we were trying to avoid.

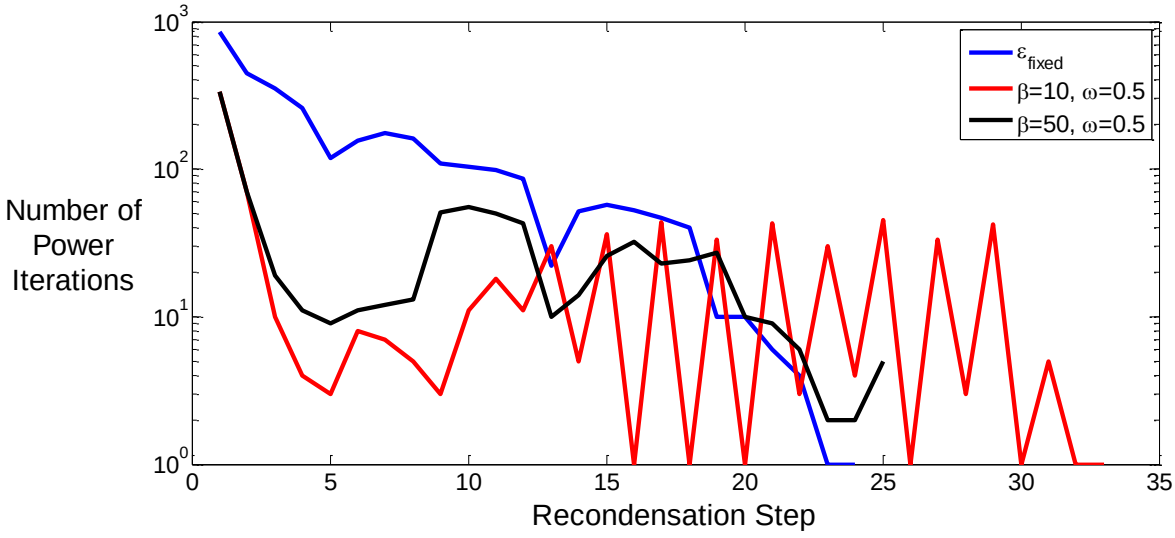

Figure 28 Comparison of the number of power iterations per recondensation step for fixed and variable criteria cases

Therefore, β was increased to 50 to improve the algorithm's performance. This new value was able to adequately dampen out the oscillations in the power iterations and better match the fixed eigenvalue's convergence. For this case, the total overall number of transport sweeps it took to converge was 6,383, a decrease by roughly 2,000 transport sweeps from the variable case in Figure 23.

While the non-monotonic convergence of the recondensation process with this new del term is not well understood at this point, it has been shown that the natural del term can speed up convergence in the outer iteration. From this point on, this definition will be used in DGM. Unfortunately, since the algorithm for choosing the coarse group convergence criteria does not seem as flexible as initially thought, it is hard to know *a priori* what the best choice of β will be for any given case. The best choice of β will likely depend on the methods used and the problem it is applied to. Further discussion of this subject is beyond the focus of this work, therefore, we will continue to set β to 50 when using the natural del term and 10 when using the exact del term.

4.3.3 The Del Term as a Modification to the Streaming Coefficients

We can also change the convergence properties of recondensation by reevaluating how we use the del term into our DGM equations. In Lei Zhu's thesis, he included it into the equations by treating it as an additional source. [41]

The derivation in Section 4.1 took a slightly different approach. When we defined the exact del term, we did so in such a way that it was kept outside of the transport sweep. We could have included it into the sweep process as was done previously, but this would have meant dividing the del term by the streaming coefficients so that we could later multiply them with these same coefficients at the next iteration. This adds an unnecessary number of additional computations that we can avoid by using the DGM equations as defined in Equation (4.8).

One of the unfortunate side effects of this method is that, when solving the coarse group equations, there is an additional term to the right side of our equation that is dependent on the coarse group angular flux values which we are iterating on. This approach poses a problem if we

wish to use nonlinear methods to solve our coarse group eigenvalue problem. Since we would be maintaining an explicit dependence of the residual on the cell-averaged angular fluxes, this requires searching for updates to the angular fluxes at each Newton step. For a typical eigenvalue solve, one only needs to update the scalar fluxes. Therefore, this approach significantly increases the number of unknowns that we need to solve for at each Newton step. In addition to making DGM applicable to nonlinear methods, it would also be far easier to implement DGM in pre-existing codes if this additional angular term could be dropped entirely.

If we analyze how the del term is used in the power iteration, we gain some intuition about how it is being used. We change Equation (4.8) slightly by adding the index *k* of the current iteration to obtain Equation (4.19). Here, we see that the del term is essentially adding a fixed point iteration into each power iteration because the del term itself does not change. Instead, it is the action of the cell-averaged angular flux on the del term which modifies its contribution to our calculation.

$$\begin{aligned} \psi_{out,0,g} &= A_g \psi_{in,0,g} + B_g Q_{0,g}^k + \delta_{0,g}^{io} \bar{\psi}_{0,g}^k \\ \bar{\psi}_{0,g}^{k+1} &= \frac{\mu}{l}\left(B_g \psi_{in,0,g} + C_g Q_{0,g}^k\right) + \delta_{0,g}^{avg} \bar{\psi}_{0,g}^k \end{aligned} \tag{4.19}$$

Therefore, on top of converging the scalar flux and source during the power iteration, we also need to converge the angular source due to the del term at the same time. We do not wish to potentially waste time converging this angular source, though, since the only way we benefit from the recondensation process is by keeping our eigenvalue solve as computationally cheap as possible.

To get around this issue, we make an assumption that simplifies Equation (4.19). A reasonable approximation is to assume that the angular flux at the *k* and *k+1* iterations are the same. This is a good assumption because we know that when the eigenvalue problem is sufficiently converged the angular flux between iterations will be nearly identical. Therefore, as the number of iterations approach infinity, we can simplify the average angular flux equation according to Equation (4.20).

$$\bar{\psi}_{0,g}^{k+1} = \left(\frac{1}{1-\delta_{0,g}^{avg}}\right)\frac{\mu}{l}\left(B_g\psi_{in,0,g} + C_g Q_{0,g}^{k}\right) \tag{4.20}$$

This is equivalent to carrying out a fixed point iteration on the external angular source while keeping the incoming angular flux and source contributions fixed. The convergence properties of the power iteration are now placed solely on the scalar flux and source update as was the case in the original coarse group eigenvalue problem.

We can further simplify the incoming-outgoing equations if we're assuming that the two iterates are the same. Since we assumed that the *k* and *k+1* iterates of the average angular flux were equal, we can substitute Equation (4.20) into the incoming-outgoing relation in Equation (4.19). This removes all dependence on the cell-averaged angular flux and instead uses the del terms to provide corrections to the streaming coefficients.

$$\psi_{out,0,g} = \left(A_g + \left(\frac{\delta_{0,g}^{io}}{1-\delta_{0,g}^{avg}}\right)\frac{\mu}{l}B_g\right)\psi_{in,0,g} + \left(B_g + \left(\frac{\delta_{0,g}^{io}}{1-\delta_{0,g}^{avg}}\right)\frac{\mu}{l}C_g\right)Q_{0,g} \tag{4.21}$$

In theory, these modified equations should allow the coarse group equations to converge without spending unnecessary time converging the angular flux source contribution from the del term. Additionally because this new form doesn't contain any explicit dependence on the cell-averaged angular flux, this opens up the eigenvalue problem for usage with nonlinear eigenvalue solvers other than power iteration and can be implemented into current transport codes more readily.

There is some uncertainty, however, in this new formulation's effect on the stability of the recondensation process. If the approximation that the cell-averaged angular fluxes between power iterations is identical breaks down, then any poorly converged solutions could increase the number of recondensation steps. On the other hand, if this assumption does hold, then any effects should be limited to the power iteration and not affect the convergence of the outer iteration. Therefore, the cases where del is incorporated as a source term according to Equation (4.8) and when incorporated as a modification to the streaming coefficients according to Equations (4.20) and (4.21) are compare in Figure 29.

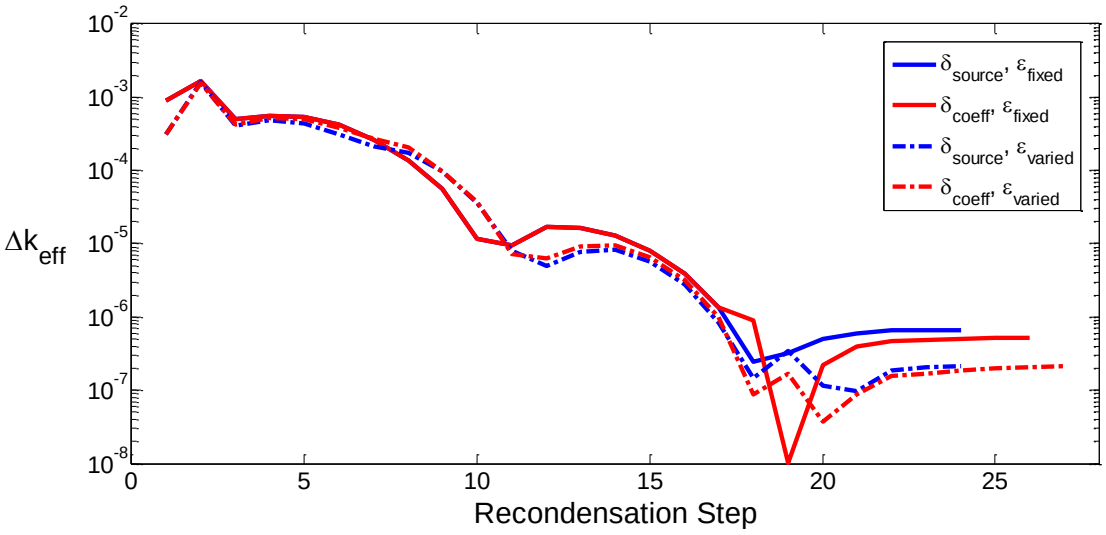

Figure 29 Comparison of convergence profiles when incorporating the del term as a source and as a modification to the streaming coefficients

First, we use the fixed convergence criteria and compare the results when using the del term in its traditional source implementation and as a modification to the streaming coefficients. Both convergence profiles are nearly identical until they approach 10^{-6}. There is some deviation after recondensation step 13, but this difference nearly disappears once the solution is sufficiently converged. We also applied the variable convergence criteria algorithm to the coefficient modification approach. Early on in recondensation, the new approach differs somewhat from the source approach but the profiles still match closely throughout the process.

If we also look at the distribution of power iterations throughout the recondensation process, then we can see the source for these different convergence profiles. For the fixed criteria case, the distribution matches well at the beginning which means that the approximation made to develop the coefficient modification approach holds well. After step 13, fewer power iterations are required to converge, but this means that more work is applied in the following step to compensate for the poorly converged coarse group solution. The normal behavior resumes after this brief oscillation. In Figure 30, the behavior for the variable criteria case seems to match fairly well throughout all the recondensation steps.

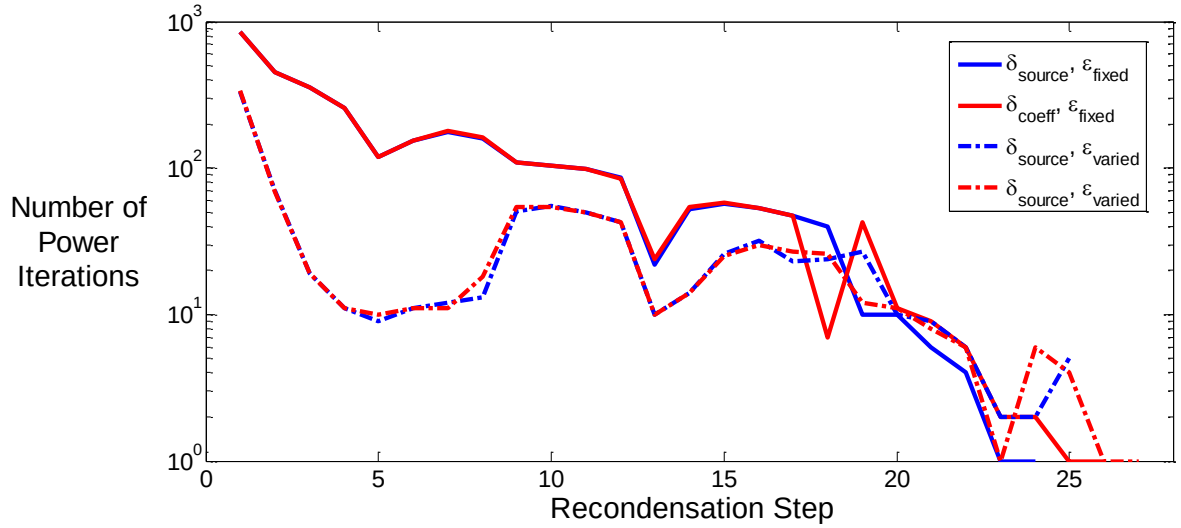

Figure 30 Comparison of the number of power iterations at each recondensation step when using the del term as a source and as a modification to the streaming coefficients

While there are minor differences when using the del term as a correction to the streaming coefficients, these deviations do not significantly affect the total number of transport sweeps applied to converge to the fine group solution. For the fixed criteria case, the number was still around 22,000 total transport sweeps and for the variable criteria case, convergence occurred at around 6,400 total transport sweeps.

Another modification can be made to the del term definition for the cell averaged angular flux equation by dividing with the scalar flux instead of the angular flux.

$$\begin{aligned} \psi_{out,0,g} &= A_g \psi_{in,0,g} + B_g Q^k_{0,g} + \delta^{io}_{0,g} \bar{\psi}^k_{0,g} \\ \bar{\psi}^{k+1}_{0,g} &= \frac{\mu}{l}\left(B_g \psi_{in,0,g} + C_g Q^k_{0,g}\right) + \delta^{avg}_{0,g} \phi^k_{0,g} \end{aligned} \tag{4.22}$$

Since we use the scalar flux, the source term in the average angular flux equation is only angularly dependent because of the del term itself. We can then apply the quadrature summation on the cell-averaged angular flux equation and come up with an equation for the scalar flux. Index *l* is now included to denote the angular dependence of the problem.

$$\sum_l \omega_l \bar{\psi}_{l,0,g}^{k+1} = \sum_l \omega_l \frac{\mu}{l}\left(B_g \psi_{in,l,0,g} + C_g Q_{0,g}^k\right) + \sum_l \omega_l \delta_{l,0,g}^{avg} \phi_{0,g}^k \tag{4.23}$$

In Equation (4.23), it is noticed that the term on the left is the definition of the updated scalar flux to be used in the next power iteration. The first term on the right is the scalar flux that would be calculated if the del term were not included into the coarse group equation. The term on the far right is a summation over the angular del source. Since the del term is the only source of angular dependence, the scalar flux can be removed from the quadrature summation.

$$\phi_{0,g}^{k+1} = \sum_l \omega_l \frac{\mu}{l}\left(B_g \psi_{in,l,0,g} + C_g Q_{0,g}^k\right) + \phi_{0,g}^k \sum_l \omega_l \delta_{l,0,g}^{avg} \tag{4.24}$$

We see that we have another fixed point iteration on the angular del source embedded into the power iteration. To further simplify our equation for the scalar flux, we assume that we conduct this fixed point iteration before actually conducting our sweep through the geometry. Now that the scalar fluxes are equivalent, we can move the far right term to the left and solve for the scalar flux.

$$\phi_{0,g}^{k+1} = \frac{\sum_l \omega_l \frac{\mu}{l}\left(B_g \psi_{in,l,0,g} + C_g Q_{0,g}^k\right)}{1 - \sum_l \omega_l \delta_{l,0,g}^{avg}} \tag{4.25}$$

One can see that we no longer have a dependence on the del term for each angle, but rather the quadrature weighted sum of the del term. This reduces the memory requirements for the coarse group del term since we have completely removed the angular dependence. We still need to find a way to factor this into the incoming-outgoing equation, though, which still depends on the average angular flux. So we expand the fully updated scalar flux in terms of its angular fluxes and find that the corrected angular flux is still multiplied by the same factor as the scalar flux.

$$\phi_{0,g}^{k+1} = \sum_l \omega_l \bar{\psi}_{l,0,g}^{k+1} = \sum_l \omega_l \frac{\frac{\mu}{l}\left(B_g \psi_{in,l,0,g} + C_g Q_{0,g}^k\right)}{1 - \sum_l \omega_l \delta_{l,0,g}^{avg}}$$

$$\bar{\psi}_{l,0,g}^{k+1} = \frac{\frac{\mu}{l}\left(B_g \psi_{in,l,0,g} + C_g Q_{0,g}^k\right)}{1 - \sum_l \omega_l \delta_{l,0,g}^{avg}} \tag{4.26}$$

If we again assume that we have converged the fixed point iteration within the power iteration, then the cell averaged angular flux should be equivalent to the updated angular flux. So we substitute the new angular flux into the incoming-outgoing equations and define a new set of modified streaming coefficients.

$$\psi_{out,l,0,g} = \left(A_g + \left(\frac{\delta_{l,0,g}^{io}}{1 - \sum_l \omega_l \delta_{l,0,g}^{avg}}\right)\frac{\mu}{l} B_g\right)\psi_{in,l,0,g} + \left(B_g + \left(\frac{\delta_{l,0,g}^{io}}{1 - \sum_l \omega_l \delta_{l,0,g}^{avg}}\right)\frac{\mu}{l} C_g\right) Q_{0,g}^k \phi_{0,g}^{k+1} \tag{4.27}$$

Equations (4.25) and (4.27) were then implemented and tested on the 47 group 1D BWR example. The results for this method are compared to the source implementation of the del term in Figure 31.

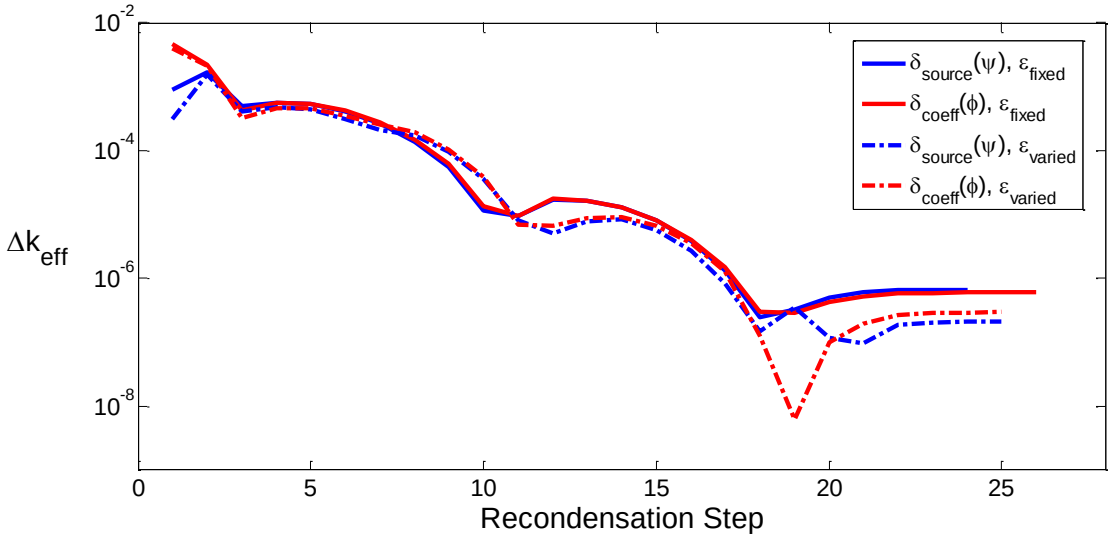

Figure 31 Comparison of convergence profiles when using the del term as a source and as a modification to the streaming coefficients

We observe in Figure 31 that using the scalar flux in the cell-averaged del term does not adversely affect the convergence of recondensation. When using either the fixed or variable convergence criteria, good agreement is shown when compared to the source implementation of the del term. The total number of transport sweep in both cases remains very close to the source implementation as well.

While there is no benefit in terms of accelerating convergence with this different implementation of the del term, it has two main benefits regarding implementation. There is a significant storage reduction since the corrective factor for the cell-averaged flux equation is not dependent on angle. For this approach, one only needs to store this factor for each coarse group and cell for the $i=0$ DGM equation. Also, one no longer needs to build a source completely separate from the fission and scattering to incorporate the del term. With this new approach, we only need to look at the transport method used and add in the del term dependence to the streaming coefficients.

Implementing this same approach to the higher order DGM equations is slightly different. For Equation (4.8), the $i>0$ moments contain a del source that is dependent on the coarse group flux.

At first, it isn't apparent that we can apply any of the algebraic manipulation we conducted for the $i=0$ equations. However, we need to remember that we are no longer in the context of the power iteration but are now in the recondensation procedure. The primary difference in this case is that no successive iteration takes place since this is a fixed source sweep, the sole purpose of which is to update the fission and scattering cross sections and the del term. The fission and scattering cross section update is straight forward since they only require the reconstructed fine group scalar fluxes. The del term must be updated on the fly, though, to avoid storing the incoming angular flux moments on each of the cell's faces. If we do this, then we no longer need to temporarily store any incoming or cell-averaged angular fluxes at each recondensation step. While this approach does not affect the incoming-outgoing equation, it does allow us to simplify the implementation in the cell-averaged angular flux equation. First we apply our quadrature weighted summation in Equation (4.28).

$$\sum_l \omega_l \bar{\psi}_{l,i,g} = \sum_l \omega_l \frac{\mu}{l}\left(B_g \psi_{in,l,i,g} + C_g Q_{i,g}\right) + \sum_l \omega_l \delta_{l,i,g}^{avg} \phi_{0,g} \tag{4.28}$$

When we do so, the left term becomes the new fine group scalar flux moment and the first right term is the scalar flux moment we would have calculated if we weren't using a del term. Since we used the scalar flux to define the del term, we can remove the scalar flux from the summation and leave the del term inside.

$$\phi_{i,g} = \sum_l \omega_l \frac{\mu}{l}\left(B_g \psi_{in,l,i,g} + C_g Q_{i,g}\right) + \phi_{0,g} \sum_l \omega_l \delta_{l,i,g}^{avg} \tag{4.29}$$

This new formulation allows us to first sweep through the *i>0* DGM equations without using a del term and calculate an uncorrected scalar flux moment. Once the scalar flux moments have been evaluated, we can use a corrective term that has no angular dependence. This term is the del term numerically integrated with respect to angle and therefore is a scalar quantity. This same term popped up when we applied this methodology to the coarse group equations and shows that we only need to store the cell-averaged angular flux del term for each cell and DGM moment.

The incoming-outgoing equations still require the full storage of their respective del terms, though. Applying the same averaged angular quantities that we did with the cell-averaged angular flux del term will introduce errors into recondensation because we would no longer be ensuring that the outgoing angular flux is consistent with the fine group problem. Therefore, memory requirements are still the main hindrance to applying DGM to a realistic problem.

4.4 Memory Requirements

While we have proved that these new del terms allow for convergence to the true fine group solution, memory storage has suffered greatly in the process. One thing that has not yet been discussed is the storage required for the coefficient matrices defined in Equation (4.7). Not only do the incoming-outgoing del terms need to be stored, but also a coefficient matrix of size *K*x*K* (*K* being the number of fine groups within a coarse group) for each segment, angle, coarse group

and cell group. Taking the 2D full core example problem from Section 3.8 and assuming an average of 6 segments per azimuthal angle while using 10 coarse groups (20 fine groups per), the total storage required becomes a staggering 18 TB. Therefore, it is extremely important to calculate these coefficient matrices on the fly instead.

First, we look at the definition of the coefficient matrices from our new derivation. We explicitly include the difference between the coarse group and fine group sweeping functions for MOC into this definition. One approach that was considered was to Taylor expand the exponential dependence of these functions. It should be noted that this procedure is only useful if the spatial method applied uses some sort of transcendental function, otherwise it is not necessary.

$$\overline{\overline{\delta A_{g,IJ}}} = a_j \sum_{K \in g} P_i(K,N)\left(e^{-\tau_g(\mu,K)} - e^{-\tau_g(\mu)}\right)P_j(K,N)$$

$$= a_j \sum_{K \in g} P_i(K,N)\left(\sum_{n=1}^{\infty} \frac{1}{n!}\left(\frac{l_k}{\mu}\right)^n \left(\overline{\Sigma_{T,g}}(\mu)^n - \Sigma_{T,g}(K)^n\right)\right)P_j(K,N) \tag{4.30}$$

$$= \sum_{n=1}^{\infty} \frac{1}{n!}\left(\frac{l_k}{\mu}\right)^n \sum_{K \in g} a_j P_i(K,N)\left(\overline{\Sigma_{T,g}}(\mu)^n - \Sigma_{T,g}(K)^n\right)P_j(K,N)$$

By performing the Taylor expansion, the spatial information can be decoupled from the energy information. We move the summation of the Taylor series outside the summation over the total number of fine groups and notice that we can define the coefficient matrix in terms of mass matrices.

$$\overline{\overline{\Sigma_{T,g}{}^n}} = \sum_{K \in g} a_j P_i(K,N)\Sigma_{T,g}(K)^n P_j(K,N) = P\ diag\left(\Sigma_{T,g}(K)^n\right)P^T A \tag{4.31}$$

Only the 1st order mass matrix is really needed because of the following property of the n^{th} order mass matrix.

$$\overline{\overline{\Sigma_{T,g}}}^{n} = P\ diag\left(\Sigma_{T,g}(K)^{n}\right)P^{T}A$$

$$= P\ diag\left(\Sigma_{T,g}(K)\right)P^{T}A\ P\ diag\left(\Sigma_{T,g}(K)\right)P^{T}A\ \cdots P\ diag\left(\Sigma_{T,g}(K)\right)P^{T}A \qquad (4.32)$$

$$= \overline{\overline{\Sigma_{T,g}}}^{n}$$

Instead of having to store a coefficient matrix for every cell, angle and coarse group, one can store all the spatial and angular information separately and the mass matrix for the fine group total cross sections. However, it becomes apparent that we have not really gained anything in terms of performance since we still need to evaluate a matrix exponential to calculate these terms on the fly. Even if a rational function expansion were used to approximate the matrix exponential, this would eventually reduce down to a rational function approximating the exponentials contained inside the coefficient matrix. Therefore, we conclude that evaluating the fine group exponential functions on the fly is the most straight forward approach to calculating the coefficient matrices on the fly.

With this issue resolved, we are now left with the memory burden of the incoming angular fluxes, which would still need to be stored for each segment. The storage requirement is now roughly 500 GB. This is definitely an improvement but still much larger than for the fine group problem. One possibility to solution is to store the del term as angular moments, such as spherical harmonics, instead of as discrete angles. Doing so would incorporate a moments-to-discrete operator, such as that found in Denovo **[**8], that could take the angular moments of the del term and convert them to a value in a discrete direction on the fly. Thus, if we were to store the values of del for 40 discrete angles (for a 2D problem using an S_8 level symmetric quadrature) instead as 3 or 4 moments, we could further reduce storage by an order of magnitude. However, it should be noted that any approximations we make on the incoming-outgoing del term will cause our full consistency to breakdown and again lead recondensation to converge to a different solution.

4.5 Summary

This section provides the first successful attempt at formulating an exact del term, allowing recondensation to converge to the exact fine group solution even when using a higher order method such as step characteristics in 1D. While this new definition was successful, it produced non-monotonic convergence in the recondensation process.

For this reason, a redefinition of the del term was considered which removed all dependence of the del term for a specific moment on that same angular flux. This produced a natural del term with substantially improved convergence relative to the original exact del term definition. Unfortunately, this did not solve the non-monotonic convergence issue.

Instead, the focus was shifted to further optimizing recondensation to reduce the total number of transport sweeps required to converge to the fine group solution. A decent algorithm was produced which used the maximum difference between previous iterates to set the convergence criteria for subsequent coarse group eigenvalue calculations. In preventing the coarse group calculation from converging all the way to the fine group convergence criteria, the total number of transport sweeps was reduced by a factor of 4 in some cases.

Alternatively, all of these tools could be used instead to produce a considerably cheap improvement to a close initial guess of the coarse group flux from an assembly level calculation. Unfortunately, in order to accurately account for the missing spatial information in the DGM moment equations the exact del term definitions established in this section require inordinate amounts of memory for realistic transport problems.

If exact convergence to the fine group solution was not required, future work could investigate ways of approximating this exact del term in an efficient manner through representation with continuous basis functions with respect to angle to provide cheap corrections to a coarse group solution

However, at this point in time, exact recondensation remains limited by the memory requirements of the del term defined using DGM. Therefore, the next section will continue

addressing this issue by relaxing some of the assumption made in the original derivation of DGM to determine if exact recondensation can be made feasible for whole core transport calculations.

5 The Source Equivalence Acceleration Method

Ideally, a single scalar corrective factor would be defined in such a way to drive the recondensation process. Unfortunately, much of the memory issues regarding the current approach to DGM is tied to the incoming-outgoing del terms defined to correct the errors accrued when using the coarse group streaming coefficients in our transport sweep. In order to even make a single scalar quantity feasible, the incoming-outgoing del term needs to be somehow removed from DGM.

One of the main issues with Equation (4.2) is that any previous errors accrued before entering the current cell are not accounted for. So instead of assuming that the true fine group is passed into the current cell, an incoming error term, E_{in}, to the incoming flux is added.

$$\begin{aligned} \psi_{out,g}(K) &= A_g\left(\tilde{\psi}_{in,g}(K) + E_{in}(K)\right) + B_g Q_g(K) + \Delta_g^{io}(K) \\ \overline{\psi_g}(K) &= \frac{\mu}{l}\left(B_g\left(\tilde{\psi}_{in,g}(K) + E_{in}(K)\right) + C_g Q_g(K)\right) + \Delta_g^{avg}(K) \end{aligned} \tag{5.1}$$

In Equation (5.1), $\tilde{\psi}_{in,g}(K)$ refers to the incorrect incoming angular flux for fine group K since the same transport operator is being used across all the fine groups in the current coarse group g. At this point it is assumed that the source for fine group K does not contain any errors. This is a good assumption since the only reason the source would be erroneous is from using inaccurate coarse group streaming coefficients in a previous iteration.

$$\begin{aligned} \psi_{out,g}(K) &= A_g\tilde{\psi}_{in,g}(K) + B_g Q_g(K) + A_g E_{in}(K) + \Delta_g^{io}(K) \\ \overline{\psi_g}(K) &= \frac{\mu}{l}\left(B_g\tilde{\psi}_{in,g}(K) + C_g Q_g(K)\right) + \frac{\mu}{l}B_g E_{in}(K) + \Delta_g^{avg}(K) \end{aligned} \tag{5.2}$$

Separating the incoming errors and error sources in Equation (5.1) provides two different sets of terms in both of the relations in Equation (5.2). The first term represents the transport sweep procedure that produces the incorrect angular fluxes when using the same transport operator across all fine groups within a coarse group. The second term provides a correction which

incorporates all the errors accrued when using the same transport operator across all the fine groups. These two sets of terms can be separated into two different transport sweep procedures such that the incorrect transport sweep can be carried out independently of the error sweep. This separation allows one to use any pre-existing code for the incorrect transport sweep, minimizing the need for significant code rewrite.

The error sweep can then be used to provide the necessary corrections to the incorrect outgoing angular flux and cell averaged angular flux according to Equation (5.3). In these relations, $\tilde{\psi}_{out,g}$ refers to the incorrect outgoing angular flux and $\widetilde{\psi_g}$ the incorrect cell-averaged angular flux. The correct outgoing angular flux, $\psi_{out,g}$, and the correct cell-averaged angular flux, $\overline{\psi_g}$, are produced by adding in the contributions of the incoming error and error sources to the incorrect outgoing angular flux and the incorrect cell-averaged angular flux.

$$\begin{aligned} \psi_{out,g}(K) &= \tilde{\psi}_{out,g}(K) + A_g E_{in}(K) + \Delta_{\mathrm{g}}^{io}(K) \\ \overline{\psi_g}(K) &= \widetilde{\psi_g}(K) + \frac{\mu}{l} B_g E_{in}(K) + \Delta_{\mathrm{g}}^{avg}(K) \end{aligned} \tag{5.3}$$

With these errors separated out, the transformation using our discrete basis functions can be applied to create a DGM-like set of equations.

$$\begin{aligned} \psi_{out,i,g} &= \tilde{\psi}_{out,i,g} + A_g E_{in,i,g} + \Delta_{\mathrm{i,g}}^{io} \\ \overline{\psi_{\imath,g}} &= \widetilde{\psi_{\imath,g}} + \frac{\mu}{l} B_g E_{in,i,g} + \Delta_{\mathrm{i,g}}^{avg} \end{aligned} \tag{5.4}$$

These error terms are useless, however, if there is no way of propagating them throughout the calculation. Further inspection reveals that the entire term added to the incorrect outgoing angular flux is the outgoing error. The error due to the use of coarse group coefficients can now be treated like a separate transport sweep.

$$\begin{aligned} E_{out,i,g} &= A_g E_{in,i,g} + \Delta_{i,g}^{io} \\ E_{avg,i,g} &= \frac{\mu}{l} B_g E_{in,i,g} + \Delta_{i,g}^{avg} \end{aligned} \tag{5.5}$$

In this new sweep, the errors in the average angular flux due to the incoming error and the error within the cell itself, E_{avg}, can now be fully accounted for. Once this error sweep has been completed, only the cell-averaged error is required to correct the cell-averaged angular flux. Since this average error term already incorporates all incoming errors from the sweep, the incoming-outgoing del term can finally be removed.

$$\begin{aligned} \tilde{\psi}_{out,i,g} &= A_g \tilde{\psi}_{in,i,g} + B_g Q_{i,g} \\ \overline{\psi_{\iota,g}} &= \frac{\mu}{l}\left(B_g \tilde{\psi}_{in,i,g} + C_g Q_{i,g}\right) + E_{avg,i,g} \end{aligned} \tag{5.6}$$

Even though an incorrect incoming angular flux is going into the average angular flux calculation, the error term now contains all the information necessary to correct this and reproduce the exact value. If the average error term is divided by the scalar coarse group flux to produce the del term, quadrature summation can then be applied over the average angular flux equations to produce a relation between the correct DGM moment, $\overline{\phi_{\iota,g}}$, the incorrect DGM moment, $\widetilde{\phi_{\iota,g}}$ and the correct coarse group flux, $\overline{\phi_{0,g}}$.

$$\overline{\phi_{\iota,g}} = \widetilde{\phi_{\iota,g}} + \overline{\phi_{0,g}} \sum_{l} \omega_l \delta_{l,i,g}^{avg} \tag{5.7}$$

Since the del term is being multiplied by the coarse group scalar flux, the quadrature sum is only applied to the del term and produces a new del term which is a purely scalar quantity. Therefore all the corrective terms for DGM have been placed into a single scalar quantity for each cell and DGM moment. The coarse group case can be further simplified to produce the following relation.

$$\overline{\phi_{0,g}} = \widetilde{\phi_{0,g}} \Big/ \left(1 - \sum_{l} \omega_l \delta_{l,0,g}^{avg}\right) \tag{5.8}$$

It is important to note that for this factor to fully correct the solution during the coarse group eigenvalue solve, it must be applied to the scalar fluxes after each transport sweep. This makes sense because this factor was derived assuming the errors had been tallied assuming a single

transport sweep, not an entire eigenvalue solve. These errors can be tallied side-by-side with the fixed source sweep in the DGM solve, thereby introducing little extra cost using this approach.

Unfortunately, the devil lies in the details. While all this simplification has produced an incredibly compact corrective factor, it has introduced another complication. For the error sources in Equation (5.5) to be correctly defined, the correct fine group incoming angular fluxes are required. Since these are not known *a priori*, the error sources rely on the accuracy of the DGM solve once each coarse group eigenvalue solve is completed. The accuracy of the DGM solve, though, relies on the accuracy of the del term. So once again the philosophical issue of which comes first, the solution or the simplification, crops up once again. The DGM solve is relying on the del term to be correct in order to improve the flux moments, while the del term is relying on the DGM solve to provide better flux moments so that it too can improve. In its current form, the method still doesn't have the necessary information to converge to the fine group solution. In order to accurately capture all the errors, the error sources used to define the del term must be unfolded.

$$\begin{aligned} \Delta_{\mathrm{i,g}}^{io} &= \overline{\overline{\delta A_{g,\iota J}}}\,\overrightarrow{\psi_{\iota n,J,g}} + \overline{\overline{\delta B_{g,\iota J}}}\,\overrightarrow{Q_{J,g}} \\ \Delta_{\mathrm{i,g}}^{avg} &= \frac{\mu}{l}\left(\overline{\overline{\delta B_{g,\iota J}}}\,\overrightarrow{\psi_{\iota n,J,g}} + \overline{\overline{\delta C_{g,\iota J}}}\,\overrightarrow{Q_{J,g}}\right) \end{aligned} \tag{5.9}$$

Each error source for a given moment relies on the incoming angular fluxes from all the other moments within that same coarse group. Therefore, the moments within each coarse group must be swept simultaneously to recapture the true incoming angular fluxes for use in determining the error source. While this can be done, it is certainly less than ideal and could significantly hinder future attempts at implementation.

On top of this, there is still the issue of computing the coefficient matrices on the fly. In addition to computing the fine group and coarse group exponentials on the fly, this process also introduces $O(K^2)$ computations since a matrix-vector multiplication must be conducted on the incoming angular fluxes. This cost is much greater than the original fine group calculation since each fine group is decoupled at the level of the transport sweep, requiring only $O(K)$ computations. This is an unfortunate side-effect of transforming the original fine group

equations using the discrete basis functions. This unfolding of the moments essentially undoes the transformation that was applied at the beginning which was supposed to lead to less computation. Equation (5.10) highlights this shortcoming.

$$\begin{aligned}\overline{\overline{\delta A_{g,\iota J}}}\,\overrightarrow{\psi_{\iota n,J,g}} &= P\,Diag\left(e^{-\tau_g(\mu,K)} - e^{-\tau_g(\mu)}\right) P^T A\overrightarrow{\psi_{\iota n,J,g}} \\ &= P\,Diag\left(e^{-\tau_g(\mu,K)} - e^{-\tau_g(\mu)}\right)\overrightarrow{\psi_{\iota n,g}}(K)\end{aligned} \tag{5.10}$$

At this point, it is deemed unnecessary to pursue DGM as the sole focus of this research. Simply put, the use of discrete basis functions end up being more trouble than they are worth. Therefore, analysis will be continued on a related method that doesn't use such a transformation, the Subgroup Decomposition Method.

5.1.1 The Subgroup Decomposition Method

At its heart, the Subgroup Decomposition Method (SGD) proposed is essentially a simplification of the earlier DGM formulation. Instead of applying a transformation using discrete basis functions, this method instead conducts the normal coarse group collapse only on the incoming energy group side of the scattering and fission cross sections. This incorporates the fine group flux information while still allowing one to use a coarse group solution to produce a fine group source later.

Recondensation can be used with this approach as well. An initial guess on the fine group flux is used to produce an initial set of coarse group cross sections and coarse group fluxes. What is different with SGD is the production of a second set of scattering cross sections which are only collapsed on the incoming energy group side. The fission cross sections can remain represented in their coarse group form since the outgoing spectrum is solely dictated by the fission spectrum and is typically assumed to have no dependence on the neutron's incoming energy. Both of these cross section are referred to as "subgroup decomposition cross sections". This second set of cross sections allows one to build a fine group source using only a coarse group flux. The next step is a fine group sweep using this new fine group source, which is referred to as the "decomposition sweep". This is slightly different from DGM in that DGM did not require a full

fine group sweep leading into recondensation, only a solve of the DGM equations. After the decomposition sweep, a new fine group flux is obtained which is used to recondense the coarse group cross sections and subgroup decomposition cross sections. This procedure then drives the initial fine group flux guess towards the true fine group solution with minimal computational effort. It has been successfully applied to step difference, linear diamond difference and diffusion in 1D and 2D calculations.[9][38]

The storage of this second set of cross sections is largely unnecessary however, since SGD as well as DGM are essentially both multigrid methods operating in energy. One only needs to look at the calculations behind the update of the coarse group and SGD cross sections to see why this is. The cross sections are collapsed by weighting the fine group cross sections with the fine group flux from the previous iteration. This produces a total reaction rate for neutron production or destruction for that coarse group. This quantity is then divided by the summation of the fine group flux, which serves as the initial guess for the next coarse group solve. This can be shown using the definition of the SGD scattering cross section. The current recondensation step is denoted by index r.

$$\Sigma_{s,K\leftarrow g'} = \frac{\sum_{L\in g'} \Sigma_{s,K\leftarrow L}\phi^{r}_{g'}(L)}{\phi^{r}_{g'}} \tag{5.11}$$

Once the coarse group eigenvalue problem is solved, the SGD cross sections are multiplied by the new coarse group fluxes to construct the new fine group source. The definition of the SGD cross sections is now expanded within the construction of the fine group source.

$$\begin{aligned} Q_g(\mathrm{K}) = \frac{1}{2}\sum_{\mathrm{g}'}\left(\frac{\sum_{\mathrm{L}\in \mathrm{g}'}\Sigma_{\mathrm{s,K\leftarrow L}}\phi^{\mathrm{r}}_{\mathrm{g}'}(\mathrm{L})}{\phi^{\mathrm{r}}_{\mathrm{g}'}}\right)\phi^{\mathrm{r+1/2}}_{\mathrm{g}} \\ + \frac{\chi_{\mathrm{g}}(\mathrm{K})}{2\mathrm{k}}\sum_{\mathrm{g}'}\left(\frac{\sum_{\mathrm{L}\in \mathrm{g}'}\nu\Sigma_{\mathrm{f,g}'}(\mathrm{L})\phi^{\mathrm{r}}_{\mathrm{g}'}(\mathrm{L})}{\phi^{\mathrm{r}}_{\mathrm{g}'}}\right)\phi^{\mathrm{r+1/2}}_{\mathrm{g}} \end{aligned} \tag{5.12}$$

When the source is formed in this manner, a common factor emerges that scales the fine group flux from the previous recondensation step. This scaling factor is nothing more than the ratio of the new coarse group flux to the previous coarse group flux. So instead of defining a separate set of SGD cross sections, a modified fine group flux can be used instead in the next fine group sweep. This is identical to using the SGD cross sections to construct a fine group source but requires much less storage in terms of cross section data.

$$\phi_g^{r+1/2}(K) = \phi_g^r(K)\left(\frac{\phi_g^{r+1/2}}{\phi_g^r}\right) \tag{5.13}$$

The new coarse group and fine group flux are labeled with the r+1/2 index since they take part in an intermediate stage in the recondensation process. The fine group fluxes obtained from the fine group sweep are designated with the r+1 index since these are the fluxes actually used to recondense the cross sections.

As with DGM, SGD has shown promise in terms of providing computationally cheap improvements to a coarse group solution, but so far this has only been conducted in 1D and 2D using step difference, diamond difference and diffusion methods. [8][38][39] Also, the del term or perturbation term defined in SGD is defined as the difference between the fine group and coarse group total cross sections weighted by the fine group scalar fluxes. This approach completely neglects the spatial consistency issue which comes from applying SGD and DGM to high order spatial problems. While this del term may have worked for step difference and finite-difference diffusion, this in no way guarantees accurate convergence of the recondensation process to higher order methods such as nodal diffusion or MOC.

5.1.2 Fully Consistent Coarse Group Collapse for Recondensation

While both DGM and SGD have been limited by certain facets of their approach to solving the fine group problem, they provide key insights into forming a feasible method. DGM provided us with the insight into correcting for spatial inconsistencies when moving to high order spatial methods. Unfortunately, the exact del term in DGM proved to be unwieldy since the original

fine group problem is transformed using discrete basis functions. As a consequence, all of the incoming angular fluxes became coupled to one another through this transformation and added further complexity to the DGM sweep. SGD provided us with a way around this issue since SGD has shown that one can use the original coarse group equations without the need for discrete basis functions and still produce a reconstructed fine group source. By conducting a fine group sweep instead of a DGM sweep, the unpleasantries associated with sweeping various flux moments simultaneously and transforming incoming angular flux moments into and out of their fine group analog can be avoided. SGD did require the production of a second set of scattering cross sections, though, to be stored in addition to the traditional coarse group ones. But, it was found that this could be further simplified by scaling the fine group fluxes from the previous recondensation step using the initial and converged coarse group flux and then use this intermediate flux to construct the fine group source. This approach allows one to use the fine group cross sections directly without the need for additional storage.

All of these aspects can be combined to form a feasible method that still provides cheap improvements to the initial solution and does so in a consistent manner when using any spatial method. We will continue our work with the corrective factor developed in Section 4.3.3 by using a sort of energy scaling approach defined in Equation (5.13) and pass this fine group flux into the fine group equations. Now the error sources can be define using the fine group equations and not the DGM equations, making their computation much more straightforward.

$$\begin{aligned} \Delta_{\mathrm{g}}^{io}(K) &= \left(A_g(K) - A_g\right)\psi_{in,g}(K) + \left(B_g(K) - B_g\right) Q_g(K) \\ \Delta_{\mathrm{g}}^{avg}(K) &= \frac{\mu}{l}\left(\left(B_g(K) - B_g\right)\psi_{in,g}(K) + \left(C_g(K) - C_g\right) Q_g(K)\right) \end{aligned} \tag{5.14}$$

This solves many of the previous issues. The flux moments no longer need to be swept together within the same coarse group, but can instead sweep the fine group fluxes separately. The discrete basis functions no longer have to be applied over and over again, saving unnecessary computation. Best of all, since the error sweep defined in Equation (5.5) is included in the fine group sweep, accurate incoming angular fluxes can be accounted for in the production of our corrective factor. This will allow the recondensation procedure to converge to the true fine group solution.

$$E_{out,g}(K) = A_g E_{in,g}(K) + \Delta_{\mathrm{g}}^{io}(K)$$
$$E_{avg,g}(K) = \frac{\mu}{l} B_g E_{in,g}(K) + \Delta_{\mathrm{g}}^{avg}(K) \tag{5.15}$$

The del term can be tallied in the same manner as the fine group fluxes throughout the sweep and then collapsed for use in the coarse group eigenvalue solve after recondensation.

$$\delta_g^{avg} = \sum_{K \in g} \sum_{l} w_l E_{avg,l,g}(K) / \phi_g \tag{5.16}$$

Upon further reflection, this corrective factor can be further simplified if after looking back to Equation (5.3) and manipulate it as follows. The error in the average angular flux can be defined as the difference between the fine group flux and the fine group flux calculated using the coarse group streaming coefficients. The bar notation denotes the correct angular flux and the tilde notation the incorrect angular flux.

$$E_{avg,g}(K) = \overline{\psi_g}(K) - \widetilde{\psi_g}(K) \tag{5.17}$$

This can then be substituted into Equation (5.16) and the summations applied.

$$\delta_g^{avg} = \sum_{K \in g} \sum_{l} w_l \left(\overline{\psi_g}(K) - \widetilde{\psi_g}(K) \right) / \overline{\phi_g} \tag{5.18}$$

This dramatically simplifies our definition of the del term.

$$\delta_g^{avg} = 1 - \frac{\widetilde{\phi_g}}{\overline{\phi_g}} \tag{5.19}$$

Care must be taken, though, to define these "incorrect" coarse group fluxes accordingly to ensure that the del term still maintains consistency with the fine group problem. In deriving this simplified form, the incorrect outgoing angular flux and cell averaged angular flux were defined according to Equation (5.20).

$$\tilde{\psi}_{out,g}(K) = A_g\tilde{\psi}_{in,g}(K) + B_g Q_g(K)$$
$$\widetilde{\psi_g}(K) = \frac{\mu}{l}\left(B_g\tilde{\psi}_{in,g}(K) + C_g Q_g(K)\right) \tag{5.20}$$

It becomes clear that, although the coarse group coefficients are being used to conduct the sweep, the fine group source produced for the fine group sweep must be used. If any other source is used, then the transport sweeps will not match and the del term will be incorrectly defined.

Next since the streaming coefficients are constant across all the fine groups within a coarse group, a set of coarse group equations can be produced by conducting the multigroup collapse procedure on Equation (5.20).

$$\tilde{\psi}_{out,g} = A_g\tilde{\psi}_{in,g} + B_g\sum_{K\in g} Q_g(K)$$
$$\widetilde{\psi_g} = \frac{\mu}{l}\left(B_g\tilde{\psi}_{in,g} + C_g\sum_{K\in g} Q_g(K)\right) \tag{5.21}$$

Now, instead of requiring an additional set of sweeps across all the fine groups, this new method only requires sweeping across the coarse groups. If the collapsed fine group source is used, then the coarse group sweep will produce a del term that can account for the spatial inconsistencies introduced when moving from a fine group problem to coarse group. Again, care must be taken that our source is the true source used in the fine group sweep. If any inner iterations to the fine group sweep, such as Gauss-Seidel, upscatter or within-group iterations, are included then the converged source from these calculation must be used for the coarse group sweep. Otherwise, the fine group source used in the coarse group sweep is not representative of the fine group reaction rates used in the multigroup collapse procedure.

If this definition is substituted back into the del term defined in Equation (5.8), then the corrective factor applied inside the coarse group eigenvalue problem can defined by Equation (5.22).

$$F_g^r = \frac{\overline{\phi_g^r}}{\widetilde{\phi_g^r}} \tag{5.22}$$

For Equation (5.23), the index for the source iteration is denoted as *k* and that for the recondensation step as *r*. The index value *k*+1/2 denotes an intermediate scalar flux which is the flux calculated after conducting a single, "incorrect" transport sweep. This is then multiplied by this new corrective factor to produce the *k*+1 scalar flux. The corrected scalar flux is then used to update the source for the next sweep, whether that be in the context of a Gauss-Seidel with or without upscatter iterations or in a within-group scatter iteration. Care does need to be taken to make sure the factor is applied after each sweep and not just at the end of the power iteration.

$$\phi_g^{k+1} = F_g^r \tilde{\phi}_g^{k+1/2} \tag{5.23}$$

From now on this factor will be called the Source Equivalence Factor (SEF). This is because the effect of it is to reproduce the collapsed fine group source regardless of how the inconsistent transport operator is defined.

To gain some intuition about how these SEFs accomplish this, their definition can be analyzed using the operator notation of the fine group problem. First, the new source is produced using the previous iterate of the fine group fluxes. Here, $L_g(K)$ represents the fine group streaming and total collision term in the transport equation and the action of $S_g(K)$ on ϕ^k will produce the fine group source using the k^{th} iterate scalar flux.

$$L_g(K)\bar{\phi}_g(K) = S_g(K)\phi^k \tag{5.24}$$

The fine group transport sweep procedure can then be denoted by the action of $L_g(K)^{-1}$ on this source. Once the updated angular fluxes are obtained, they can be summed according to any given angular quadrature. The quadrature summation will be denoted by the operator *M*, are also known as the discrete to moment operator. The correct fine group scalar flux is defined, in operator form, according to Equation (5.25).

$$\bar{\phi}_g(K) = ML_g(K)^{-1}S_g(K)\phi^k \tag{5.25}$$

If an incorrect transport method is used instead, such as that produced after multigroup collapse, to produce the next iterate of the scalar flux, then an incorrect scalar flux is produced. Notice that the use of operator L_g^{-1} represents the application of the coarse group sweep procedure.

$$\tilde{\phi}_g(K) = ML_g{}^{-1}S_g(K)\phi^k \tag{5.26}$$

Also, by the very definition of the coarse group fluxes, it is known that the difference between the correct and incorrect coarse group scalar fluxes is equivalent to summing the differences between the correct and incorrect fine group fluxes within that coarse group.

$$\bar{\phi}_g - \tilde{\phi}_g = \sum_{K\epsilon g}\left(\bar{\phi}_g(K) - \tilde{\phi}_g(K)\right) \tag{5.27}$$

The definitions for the correct and incorrect fine group fluxes in operator notation can be substituted into Equation (5.27). If it is assumed that the same fine group source is used in both definitions, then it is apparent that the difference between the correct and incorrect coarse group flux implicitly calculates the difference between the two transport sweep operators. This provides a measure of the error accumulated due to the multigroup collapse of our total fine group cross sections.

$$\bar{\phi}_g - \tilde{\phi}_g = \sum_{K\epsilon g} M\left(L_g(K)^{-1} - L_g{}^{-1}\right)S_g(K)\phi^k \tag{5.28}$$

This difference is directly related to the del term defined in Equation (5.19). Through another simple manipulation, it is observed that the opposite of the del term is really the relative error in the incorrect coarse group flux.

$$-\delta_g^{avg} = \frac{\tilde{\phi}_g - \bar{\phi}_g}{\bar{\phi}_g} \tag{5.29}$$

Therefore, if the del term is negative, this means that the incorrect sweep has erroneously produced extra neutrons relative to the fine group sweep. Since the incorrect flux is larger than the correct flux, the SEF defined in Equation (5.22) becomes less than one in order to remove these extra neutrons after each incorrect sweep. On the other hand, if the del term is positive, then this means the incorrect sweep has removed neutrons from the calculation which otherwise

would have remained in our correct sweep. In this case, the incorrect flux is smaller than the correct flux, making the SEF greater than 1 in order to replace the neutrons lost after the incorrect sweep. **So, in both cases, the SEF is enforcing neutron balance from the fine group sweep across every coarse group sweep. This ensures that the coarse group problem is consistent with the fine group problem so that the recondensation problem can converge to the true solution.**

Another way to think about this is in terms of net leakage. The SEFs were defined by accounting for all the incoming and outgoing errors at each of the cells' boundaries in Equation (5.5). Summing up the incoming and outgoing angular fluxes across the cell's boundaries gives the net leakage for that cell. Any errors in the incoming and outgoing angular fluxes then produce errors in these net leakages. Therefore, these SEFs can be defined in terms of net leakage according to a simple neutron balance equation to provide further intuition. In this basic neutron balance equation, A_b is the area of surface b enclosing volume V and $\bar{J}_b$ is the correct net surface current normal to surface b, and $\bar{\phi}$ is the correct fine group flux.

$$\Sigma_T \bar{\phi} V = \sum_{b \epsilon S} \bar{J}_b A_b + QV \tag{5.30}$$

In this correct case, neutrons are perfectly conserved because the total number of neutrons leaving the cell and being removed from within equals the net production of neutrons inside the cell. This equation is satisfied when the fine group problem is fully converged. But, when the multigroup collapse is applied, the transport sweep becomes inconsistent for higher spatial order methods and results in incorrect surface currents at the boundaries of the cell. These incorrect net surface currents, $\tilde{J}_b$, then produce incorrect fine group scalar fluxes, $\tilde{\phi}$.

$$\Sigma_T \tilde{\phi} V = \sum_{b \epsilon S} \tilde{J}_b A_b + QV \tag{5.31}$$

If Equations (5.30) and (5.31) are divided, the ratio of these balance equations directly related to the ratio of the correct and incorrect transport sweeps, Equations (5.25) and (5.26). Since neutron

conservation must be upheld in both basic neutron balance and in the transport sweep, the two definitions must be equal to one another. In this case, when the same source is used in the cell, the errors in the surface currents originate solely from the incorrect transport sweep used after conducting the multigroup collapse.

$$\frac{\bar{\phi}}{\tilde{\phi}} = \frac{\sum_{b\epsilon S} \bar{J}_b A_b + QV}{\sum_{b\epsilon S} \tilde{J}_b A_b + QV} = \frac{M\bar{L}^{-1}Q}{M\tilde{L}^{-1}Q} \tag{5.32}$$

It is important to remember that only considering energy condensation is being considered in this case and not homogenization. It is assumed that multigroup collapse is being conducted across the same discretization as the fine group problem. Therefore, the net source of neutrons and the cell boundaries are the same on the numerator and the denominator.

If this simplification weren't enough, these SEFs can actually be folded back into the coarse group source to produce a set of fully consistent coarse group cross sections. This can be done by observing how the SEFs are applied to the previous scalar flux and then incorporated into the production of the coarse group source.

$$\begin{aligned} Q_g^{k+1} = \frac{1}{2}\sum_{g'}\left(\frac{\sum_{K\in g}\sum_{L\in g'}\Sigma_{s,K\leftarrow L}\bar{\phi}_{g'}^r(L)}{\bar{\phi}_{g'}^r}\right)\left(\frac{\bar{\phi}_{g'}^r}{\tilde{\phi}_{g'}^r}\right)\tilde{\phi}_{g'}^{k+1/2} \\ + \frac{\chi_g}{2k}\sum_{g'}\left(\frac{\sum_{L\in g'}\nu\Sigma_{f,g'}(L)\bar{\phi}_{g'}^r(L)}{\bar{\phi}_{g'}^r}\right)\left(\frac{\bar{\phi}_{g'}^r}{\tilde{\phi}_{g'}^r}\right)\tilde{\phi}_g^{k+1/2} \end{aligned} \tag{5.33}$$

The most striking observation is that the denominator of the collapsed cross sections cancels with the numerator of these SEFs. Thus, instead of applying the SEFs after each sweep, the coarse group cross sections can instead be defined such that this correction is already taken into account.

$$\nu\Sigma_{f,g}(\vec{r}) = \sum_{K\epsilon g} \bar{\phi}_g(\vec{r},K)\nu\Sigma_{f,g}(\vec{r},K)/\tilde{\phi}_g(\vec{r})$$
$$\Sigma_{s,g'\to g}(\vec{r}) = \sum_{L\epsilon g'}\sum_{K\epsilon g} \bar{\phi}_{g'}(\vec{r},L)\Sigma_s(\vec{r},L\to K)/\tilde{\phi}_{g'}(\vec{r}) \quad (5.34)$$

This slight variation on the original definition of multigroup collapse shifts complete enforcement of neutron balance directly to the coarse group fission and scattering cross sections. This is a more natural approach since the coarse group cross sections were originally intended to preserve the reaction rates from the fine group problem. **In this modification, though, the reaction rate is divided by the incorrect coarse group flux instead of using the collapsed fine group flux.** Folding the SEFs into these cross sections allows for the simultaneous preservation of cell reaction rates and net leakages without using multiplicative factors after each coarse group transport sweep.

It is important to note that these new cross sections completely replace the original cross sections; therefore no additional storage is added onto the original coarse group problem. Since the original fine group problem was not modified, the fine group cross sections can still be stored per material as well. There is only a slight increase in memory due to storage of the coarse group scalar flux before and after converging the coarse group eigenvalue problem in order to conduct the prolongation of the fine group scalar fluxes. Another slight increase can be attributed to temporary storage of the collapsed fine group source needed to calculate the incorrect coarse group fluxes to produce the modified cross sections. Therefore, a new set of cross sections has been defined that allow the recondensation process to converge to the fine group solution for any spatial method with very little additional storage.

It is important to note that this new approach shares some similarities with the SPH methods since it uses a single multiplicative factor to enforce preservation of reactions and net currents. There are two main differences between SPH factors and SEFs. First, SPH factors are produced to correct for error introduced when spatial homogenization and energy condensation has been applied. SEFs are produced intentionally without including spatial homogenization. Second,

SPH factors are applied to all the cross sections being generated, including the total cross section. Since the total cross section is used in the transport calculation, an iterative procedure is used to converge the SPH factors before application to the coarse group eigenvalue problem. An additional constraint is also required for this approach, typically some sort of volume-flux normalization equation. For this application, SPH factors would require an iterative process to converge these factors for each recondensation iteration. On the other hand, SEFs are produced assuming that the total cross section isn't modified, therefore all necessary corrections are being applied only to the right-hand side cross sections. Since the total coarse group cross section remains unaffected by the application of these SEFs, only a single coarse group transport sweep is required at each recondensation iteration. [14]

These new cross section definitions were applied to the 47 group 1D BWR problem and used in conjunction with step difference, diamond difference and step characteristics methods. For each method, the errors in the eigenvalue were calculated relative to the eigenvalue of the fine group problem using like methods. Therefore, the recondensation solution using step difference is being compared to the fine group solution using step difference, for example.

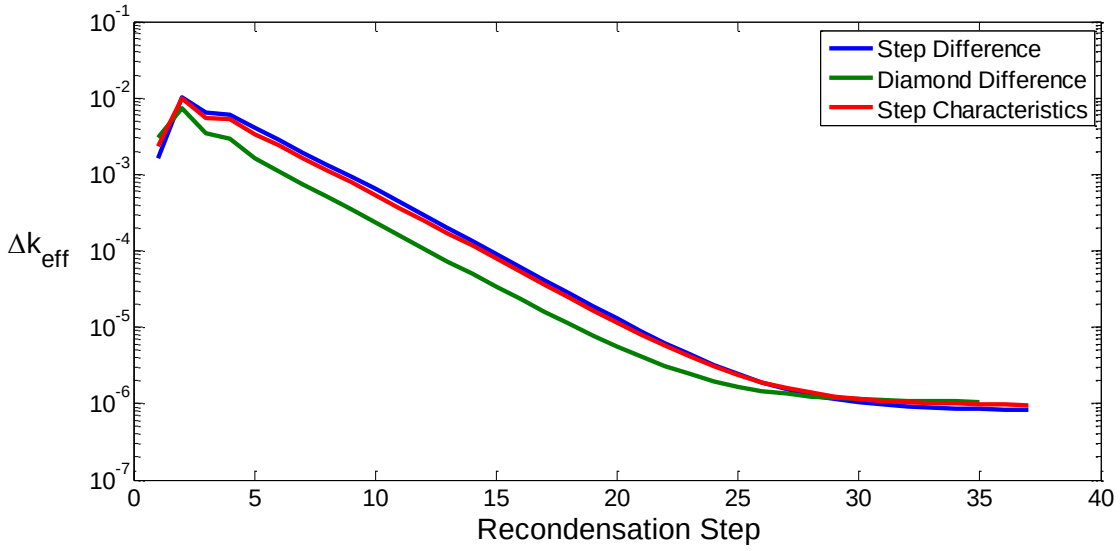

Figure 32 Comparison of eigenvalue convergence when applying the consistent cross sections to various 1D Transport methods for the 47 Group BWR problem.

In Figure 32, it is seen that across all these different transport methods, recondensation with these new cross sections definitions converges within 0.1 pcm of the reference eigenvalue when the convergence criteria is set to 10^{-7}. Since the error is roughly on order with the convergence criteria, this means that this approach is consistent and can be applied to any spatial method.

5.1.3 Applying the Coarse Group Sweep

One of the great benefits of this new approach is that it provides much more flexibility in how recondensation is applied. First, we look at the impact of separating the condensation of our coarse group total cross section from the fission and scattering cross sections.

In all the previous methods, only the previous iterate of the total coarse group cross section could be used in the DGM sweep. Because the fine group fluxes could only be reconstructed after the DGM sweep, the del term was correcting for the use of the previous coarse group total cross section and not the total cross sections that would be used in the next coarse group eigenvalue problem. With these consistent multigroup cross sections, the calculation no longer needs to be held back by a lagging del term.

Instead, the coarse group total cross sections can be collapsed once the fine group sweep has been conducted and then use the newly collapsed coarse group total cross section in the fixed source coarse group sweep. This takes into account all the errors accrued using the current iterate of the coarse group total cross sections as opposed to using the previous iterate. Instead of defining our recondensation procedure according to Figure 2, the new procedure depicted in Figure 33, called the Source Equivalence Acceleration Method (SEAM), allows the collapse of the total cross section to be separated from the collapse of the fission and scattering cross sections.

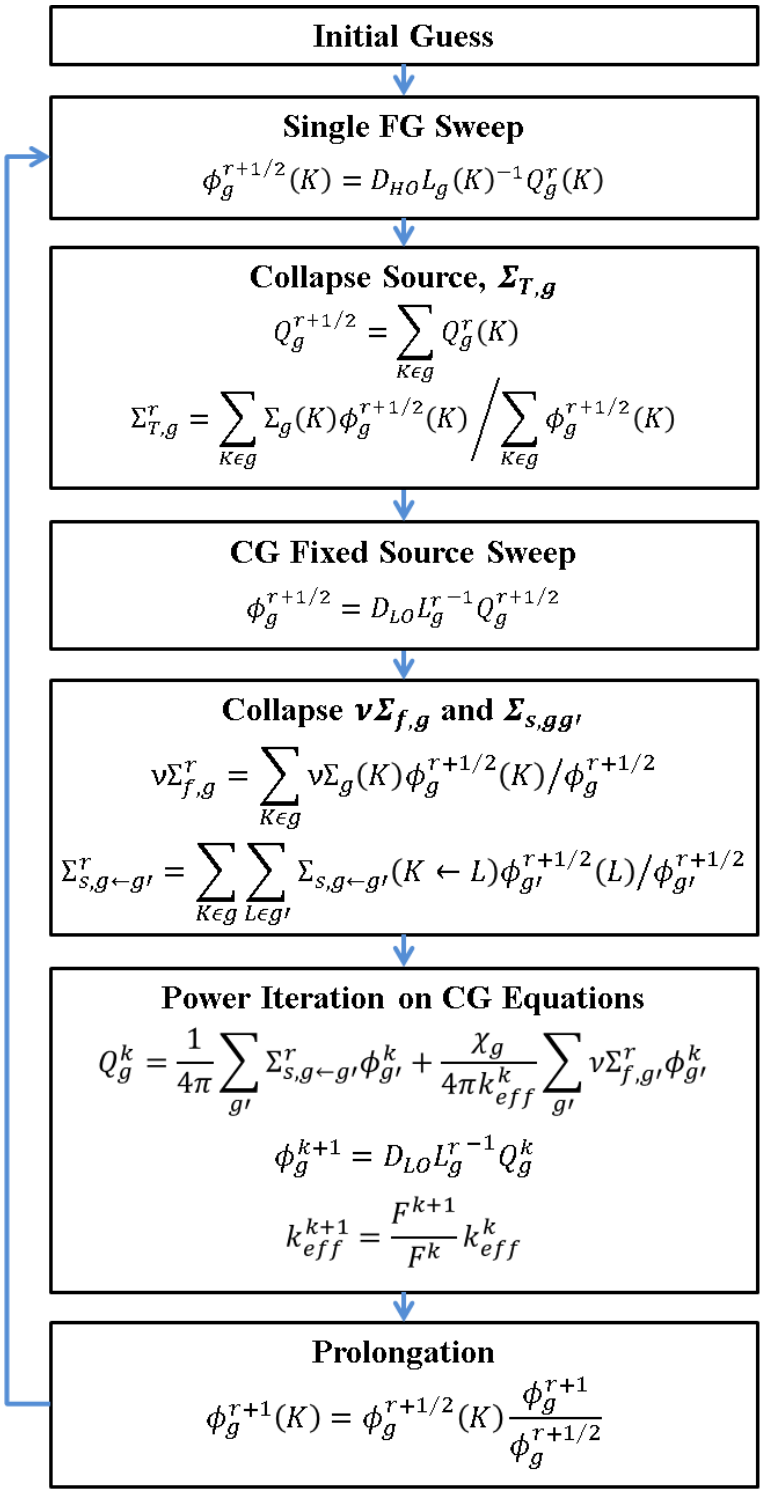

Figure 33 New flow chart for recondensation using the Source Equivalence Acceleration Method (SEAM)

This new procedure was tested on the 47 group 1D BWR problem and the convergence compared to the old recondensation procedure. A flat flux both in energy and space was used to initialize the recondensation process.

One important thing to note is that the Krasnoselskii iteration is no longer being used. In Figure 34, even though the second coarse group eigenvalue jumps up to an extremely large eigenvalue, recondensation still manages to recover and eventually converge. This shows that the fully updated cross sections can be used in both cases without having to use any dampening factors. Even though using the previous total cross section for the coarse group sweep is equivalent to the use of the del term as defined in Equation (4.8), these new cross section definitions have stabilized the recondensation procedure.

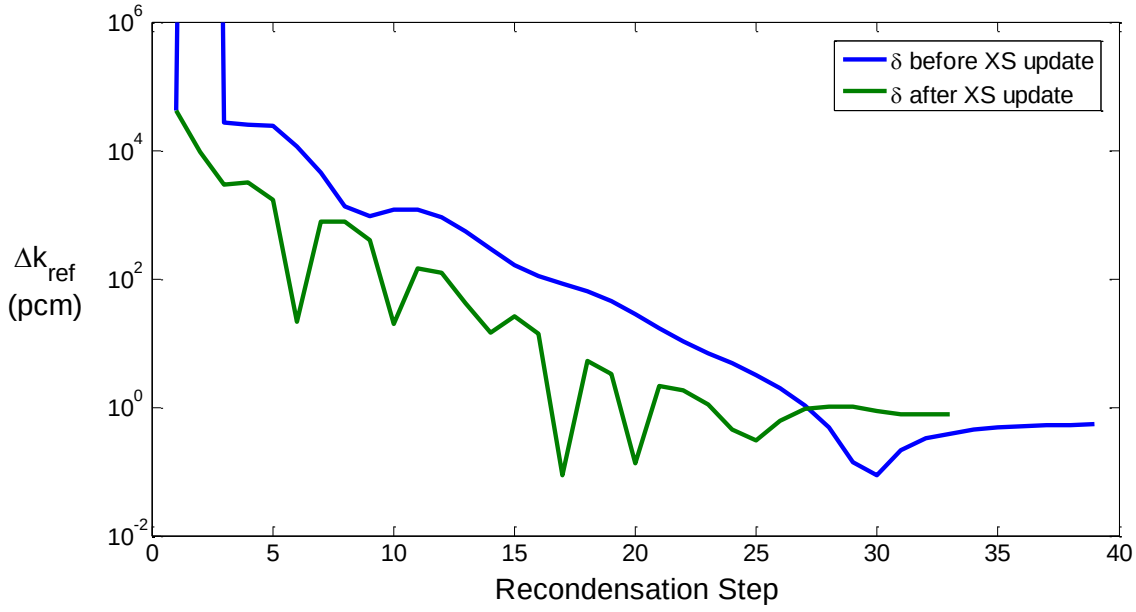

Figure 34 Comparing convergence profiles when conducting the coarse group sweep before and after the new coarse group total cross section is recondensed.

In addition to providing better stability, collapsing the coarse group total cross sections before the coarse group sweep helps reduce the number of recondensation steps required to converge the eigenvalue. Unfortunately, this improvement brings with it the oscillatory behavior that was first observed in Figure 24 but in this case the effect is even more pronounced. Clearly, there is

something else holding the recondensation process back outside of how the recondensation procedure is ordered.

5.1.4 Removing the Non-Monotonic Convergence of the Recondensation Process

Since SEAM's cross sections now provide full consistency with the fine group problem, there must be some other source which is driving the non-monotonic convergence. Taking another look at the recondensation procedure is therefore necessary to try and identify the source of the oscillations observed in Figure 34. One possible issue could be with the coarse group eigenvalue problem itself, but this can't be the source of the non-monotonic convergence since the fluxes are being tightly converged in this calculation. Recondensation of these cross sections can't be the source either since the new definitions were already shown to be fully consistent and actually improve convergence. The fine group sweep, though, is where the least amount of effort has been placed in terms of converging the solution and is the only other possible source of this behavior.

One of the primary assumptions made when GEC and DGM were first derived was that a single fixed sourced sweep using the DGM equation or fine group equations would be sufficient to improve the coarse group solution in the next recondensation step. While the eventual convergence of the recondensation process does show that this is the case, it was theorized that the oscillations may be because the convergence of the coarse group solution may be outpacing that of the fine group problem. Because of this, the coarse group problem can only converge as quickly as the fine group sweep allows. To test this, many different approaches to better converge the fine group problem were tested. These methods were applied to both the coarse group eigenvalue solve and the single fine group sweep or source iteration and compared relative to the reference solution converged using those same methods.

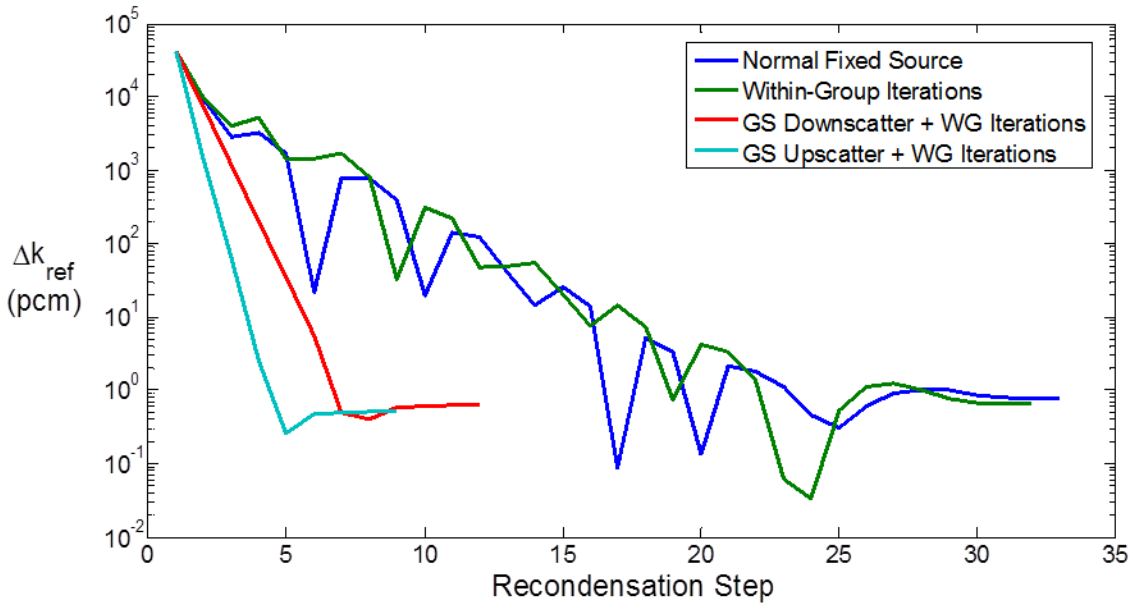

Figure 35 Comparing the convergence of the eigenvalue relative to the 47 group 1D BWR fine group solution when using various energy convergence methods.

To test this idea the following four calculations were considered:

- Normal power iteration – This provided a direct reference by which the alternatives could be compared and their relative performances ascertained
- Within-group iterations – Instead of producing a fixed source, the within-group portion of the source is converged by a fixed point calculation. First, the source is built ignoring the contributions of within-group scattering. The within-group scattering is then added to this source and the transport sweep is conducted. After sweeping through all the cells and angles, the flux is updated, a new within-group source is built and another sweep is conducted. This iterative process completes when the source has converged to some set tolerance or has reached the maximum allowable number of iterations. In this case, the same tolerance set for the coarse group eigenvalue solve is used and the max number of iterations is set to 10. However, the results in Figure 35 show very little improvement.
- Gauss-Seidel with downscatter – The within-group iterations begin with the fastest group and the converged flux used to produce the source for the next slowest group. The

upscatter sources were calculated using the previous iterate of the scalar fluxes and therefore remained fixed throughout the calculation. When this procedure was applied to the source iteration on our fine group problem, it removed most, if not all, of the oscillations. As suspected, the oscillations in recondensation were a direct result of a poorly converged fine group solution.

- Gauss-Seidel with upscatter – This approach applies the previously defined downscatter loop until reaching the section of the scattering matrix containing upscatter cross section. When Gauss-Seidel first reaches the first group with upscattering cross sections, it initially treats them as a fixed source and proceeds as if it were a purely downscatter problem. When the final energy group has been swept, all the fluxes in the upscatter section are used to update the upscatter sources. All of the fluxes in the upscatter section are then swept again using the new upscatter sources. This process repeats until fluxes converge to the set tolerance. As seen in Figure 35, convergence improves even more if upscatter iterations are included into the fine group source iteration.

The tolerances for the within-group and upscatter iteration conducted at each power iteration are set using residual from the previous power iteration divided by 100 or the convergence criteria for the power iteration, whichever is larger. This ensures that the update of the k_{eff} is not being held back by a lack of convergence in space and energy between power iterations.

While this work definitively shows that these oscillations were in fact due to the fixed source fine group sweep, it is not readily apparent in the 47 group problem tested that converging the fine group flux at each step is actually beneficial in terms of computational effort. Therefore, these approaches were also tested on the 361 group 1D BWR problem to be defined in Section 5.2. This model was thought to be significantly harder to solve since the resonances become much more pronounced and lead to significant variations in the shape of the flux across the neutron spectrum. Also, the number of upscattering groups increases to 89, further making the fine group problem more difficult to converge with a single fixed source sweep.

This time, though, a slight modification is made to the original fine group convergence scheme by first applying Gauss-Seidel without and with upscatter to the coarse group problem and then

apply them to the fine group source iteration as well. The results were compared to test if the assumption holds that the coarse group problem was already sufficiently converged in the eigenvalue solve. Figure 36 shows that applying Gauss-Seidel only to the coarse group problem does absolutely nothing to remove the oscillations, both the green and blue lines lay directly on top of one another. Only when Gauss-Seidel is applied to the fine group problem is there nearly a complete removal of the non-monotonic convergence in the recondensation process.

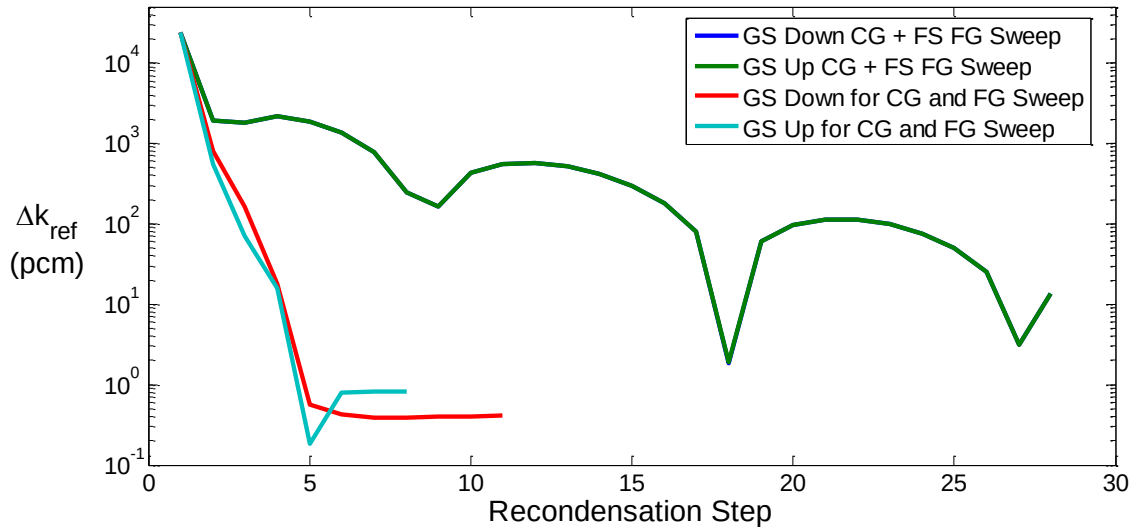

Figure 36 Effect of using energy convergence methods on the convergence of k during recondensation for the 361 group 1D BWR problem using the 34 coarse group mapping

These results also reveal how dependent the convergence of recondensation is on the problem at hand. For this 361 group problem, the oscillations still occur and the overall convergence occurs much more slowly when only a fixed source sweep is applied at each recondensation step. After 30 recondensation steps, the eigenvalue hasn't even come close to converging. When Gauss-Seidel with and without upscatter is applied, the eigenvalue converges to below 1 pcm using only 5 recondensation steps. This further emphasizes the importance of converging the fine group source at each recondensation step. While this does significantly decrease the number of recondensation steps, it is necessary to look at how much computational effort is being spent on converging the coarse group eigenvalue problem and the fine group source.

First, we look at the total number of transport sweeps that has been done for each coarse group. In Figure 37, it is seen that many more sweeps have been conducted across all coarse groups when the fixed source solve of the fine group problem is used. The number of sweeps decreases significantly when the fine group source is being converged. Because the tolerance of the fine group source iteration are being set close to that of the coarse group eigenvalue solve, the next coarse group solve uses cross sections containing improved fine group reactions rates. However, this drop in coarse group sweeps is still roughly proportional to the number of recondensation steps. Therefore the total number of fine group sweeps also an important measure of the additional effort being placed on converging the fine group source.

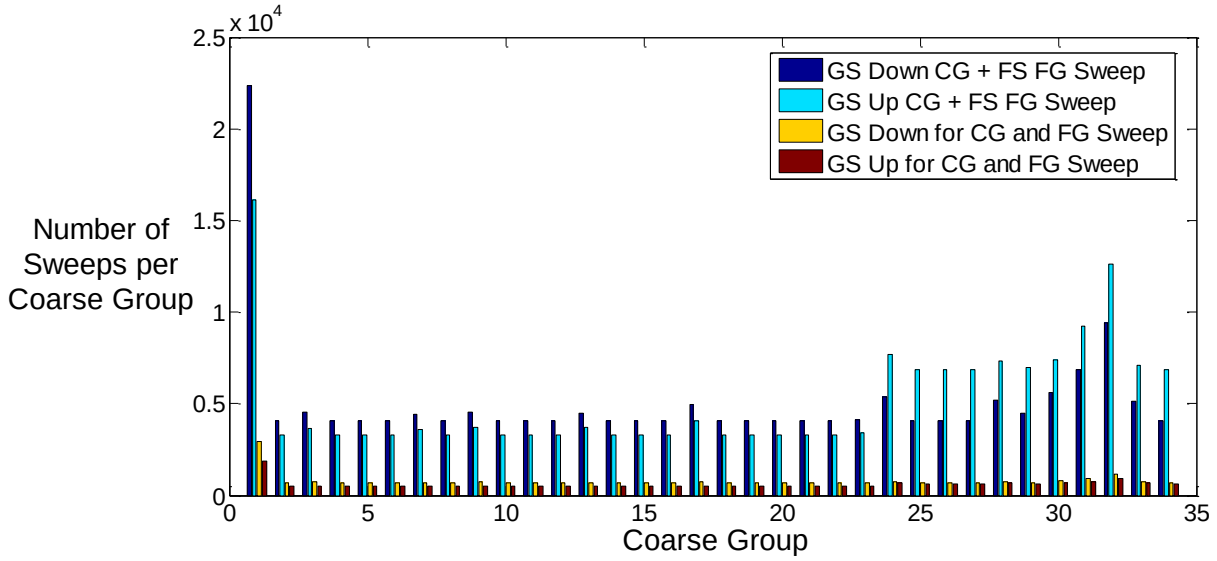

Figure 37 The total number of sweeps conducted for each coarse group over the recondensation process for the 361 group 1D BWR problem

In Figure 38 the use of Gauss-Seidel on the coarse group eigenvalue problem and a fixed source sweep on the fine group problem results in a flat distribution of the total number of fine group sweeps applied. This is to be expected since all the fine groups are being swept the same number of times with the fixed source approach. This also reaffirms that not converging the energy variable in our coarse group problem results in little to no acceleration.

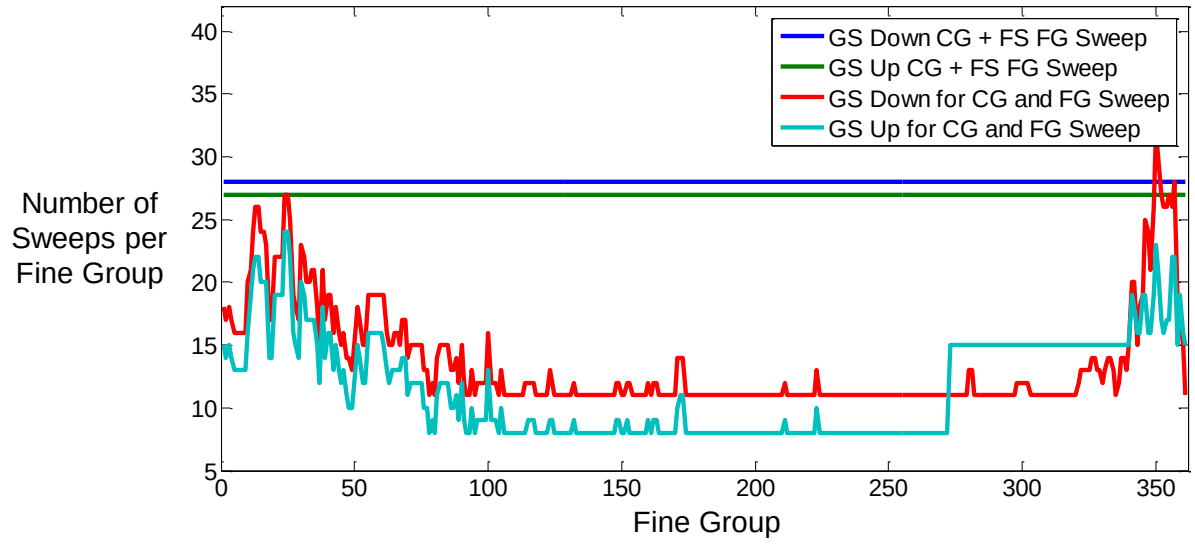

Figure 38 The total number of sweeps conducted for each fine group over the recondensation process for the 361 group 1D BWR problem

On the other hand, if Gauss-Seidel is applied to the fine group problem, the "additional" effort leads to a significant reduction in the number of fine group sweeps. This reduction is due to the additional emphasis placed on the fluxes which are harder to converge when within-group iterations and upscattering is included. Because both sets of iterations better distribute the computational effort across the fine group fluxes, sweeps are not wasted on groups which are easier to convergence. This is highlighted in Figure 38 by the number of sweeps applied to fine groups 100 to 275. For these groups, only one or two transport sweeps per recondensation step are required to converge the flux to the necessary tolerance.

5.2 Application to the 1D HTR and 1D MOX Fine Group Problems

A set of two distinct 1D reactor models were used to test the 1D application of SEAM to evaluate this new method's ability to accelerate fine group transport calculations. Each reactor type represents significantly type of physics dominating the calculation.

Lei Zhu developed a 1D 295 group HTR model to test his initial work with DGM and recondensation. [41] A simple diagram of his simplified HTR core is provided in Figure 39. The

reflective boundary condition placed on the left edge and the vacuum boundary on the right edge approximately simulates an annular graphite core.

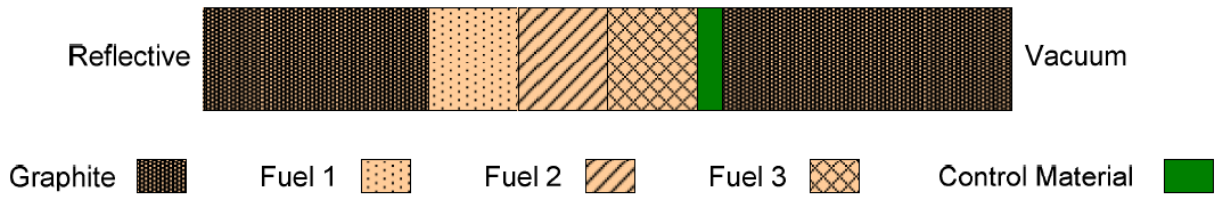

Figure 39 Simple description of the 1D 295 Group HTR model

While scattering is the dominant interaction for this problem due to the presence of graphite, relatively few of the resonances are captured in the cross section data set provided for this problem. The homogenization of the fuel region effectively self-shields a significant portion of the resonances to begin with, thereby making the resonances less pronounced. A high absorbing medium was added to this model, simulating the placement of the control rod at the outer edge of the fuel region. The thermal absorption cross section for this region is extremely high and results in a significant flux depression and spectral shift. Another difficult aspect of graphite reactor physics, which is captured in this 295 group model, is the large influx of thermal neutrons at the boundary of the fuel region and the graphite reflector/moderator. A full description of the geometry used in this model can be found in Table 5.

Table 5 Description of the geometry for the 1D 295 Group HTR Model

Region	1	2	3	4	5	6
Material	Graphite	Fuel 1	Fuel 2	Fuel 3	CR	Graphite
Width (cm)	90	36	36	36	10	116

The 1D 361 group BWR model, developed by Nathan Gibson, uses 361 group cross sections produced according to the SHEM-361 fine group structure. [12] Though, the cross sections produced are largely unrealistic since no energy self-shielding model was applied in their creation, this core still captures many of the difficult physics involved with the analysis of MOX loaded LWR cores. Key to analyzing this core is being able to capture the swapping of different spectra between neighboring MOX and UO_2 assemblies. While there are slight differences

associated with the fission spectra of both assembly types, the true differences come from the presence of different resonances and the larger thermal absorption in the MOX fuel. This 361 group model is able to accurately characterize these spectral shifts which may be lost with a few group model.

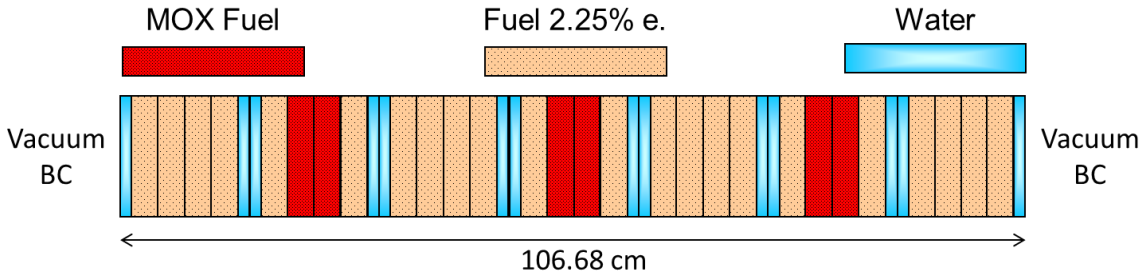

Figure 40 Simple description of the 1D 361 Group BWR model tested

Figure 40 presents a simple model of the 361 group MOX loaded core used for this analysis. This core consists of an alternating array of 4 low enriched UO2 assemblies and 3 MOX loaded assemblies. These assemblies are 15.24 cm in width and each contain water spacings separating neighboring assemblies. Vacuum boundary conditions are applied to both sides of the core.

Both of these cores were tested in a 1D code built using MATLAB. While MATLAB has the nice advantage of many built-in functions to do fast matrix-vector products and matrix inversion, since it is not a pre-compiled compute language but rather an interpreted language, MATLAB executes instructions directly and translating them on the fly. This leads to very poor performance on transport calculations and is why performance comparison for this section are measured in terms of fine group, coarse group and total number of transport sweeps.

5.2.1 Effect of Coarse Group Structure

In reactor analysis, choice of coarse group structure is dictated by the ability of said structure to produce a solution to within some accuracy relative to the fine group problem. If solution accuracy is not attained, the coarse group structure is tossed since some key physics are likely being neglected. However, in some cases it is not guaranteed that an accurate eigenvalue is truly

representative of the accuracy of the whole solution. Depending on the structure analyzed, error cancellations may occur in core and lead to an erroneously accurate eigenvalue.

Therefore, recondensation using SEAM is applied to the BWR and HTR problems described in Section 5.2. For each reactor type, two different coarse group structures are chosen for the coarse group problem by roughly applying the methodology proposed by Gibson and Forget. This produced a 34 group structure for the BWR problem and a 23 group structure for the HTR problem. The performances of these structures were compared to a simple 2 group structure for both reactor types. These 2 group structures were chosen such that the second group contained all groups containing upscatter and the first group comprised of the remaining groups. The eigenvalue convergence relative to the reference fine group calculation for each case is compared in Figure 41. For all cases, 1D step characteristics was applied using an S_{32} angular quadrature.

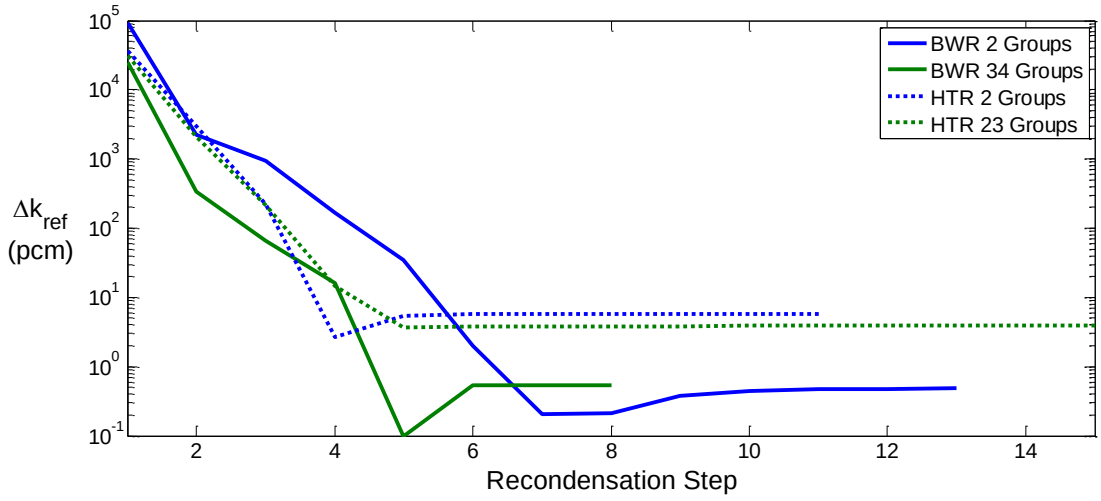

Figure 41 Convergence profiles for the 1D BWR using 2 and 34 groups and for the 1D HTR using 2 and 23 groups

In the BWR case, the eigenvalues converge to nearly the tolerance applied to the fine group eigenvalue problem. The difference between the recondensation and fine group eigenvalue asymptotically approaches a value below 1 pcm. This further shows that the new cross section formulation still maintains full consistency with the fine group problem. In addition, the 34 group structure is used instead of the 2 group, the recondensation process converges faster, at least in terms of recondensation steps. This is to be expected since the 34 group structure is

more likely to capture the physics of the fine group problem while the 2 group structure neglects much of these details. However, since error cancellations can hide the true error of the solution, the error in neutron fluxes for the 34 group and 2 group solutions are calculated and compared relative to the reference fine group solution.

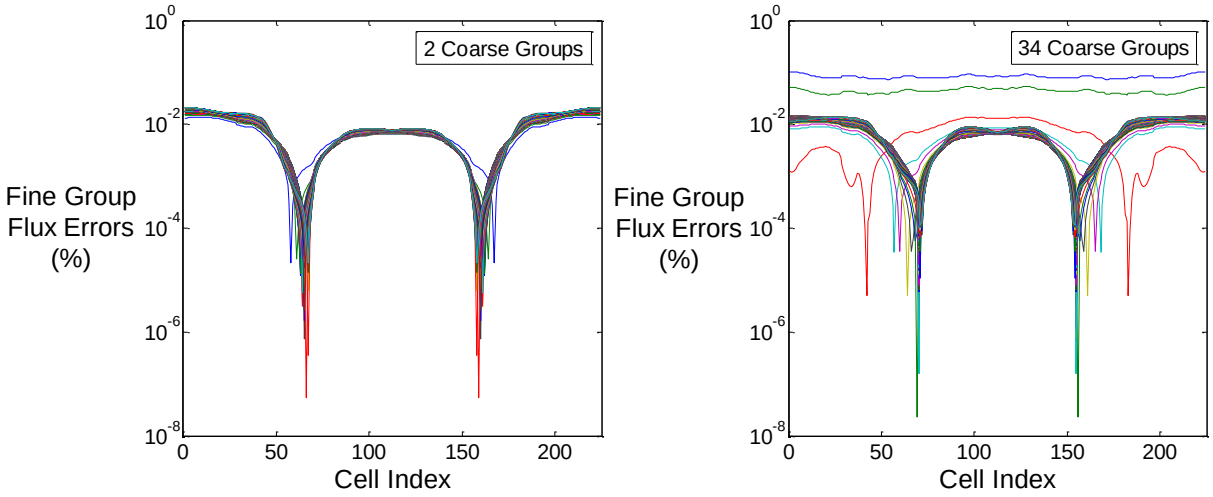

Figure 42 Relative error in the 361 group fluxes when using a 2 group and 34 group mapping for the 361 group BWR problem

For both structures, nearly all 361 group fluxes are near or below 0.01% of the reference solution. Therefore, Figure 42 definitively shows that SEAM is fully consistent and adequately resolves errors such that error cancellation is no longer a concern in our neutronics analyses. The errors in most of the fluxes are so low that the error takes the shape of the second harmonic. This also explains why most of the errors dip so low near the same spatial points because this is where the magnitude of the second harmonic vanishes, leaving a near exact fundamental mode at those points. This has big implications since any coarse group structure may be chosen and consistency with the fine group problem will always be maintained.

It should be noted, though, that errors in 3 of the fine group fluxes for the 34 group structure problem deviate significantly from the other groups. The primary cause is that the values of these 3 fluxes are on the order of the convergence criteria. Therefore any proportionally large changes between power iterations in these fluxes will not affect satisfaction of the convergence

criteria since the L_∞ norm defines the flux residual. Because the L_∞ norm used in 34 group recondensation problem required fewer fine group sweeps to satisfy the convergence criteria, the errors in these fluxes remained large relative to the other converged fluxes. While the 1D results in this section continue to use this definition of the residual to ensure an apples to apples comparison, the 2D calculations of Section 6 will employ the L_2 norm of the total source residual to avoid this issue.

With accuracy shown in the BWR case, it is important to determine how much if any computational effort has been spared using this method. For the 2 group and 34 group cases, the total number of coarse group sweeps and fine group sweeps applied are compared in Table 6. These numbers do not take into account the total number of angles being swept over. The number of sweeps applied per fine group is also calculated, as well as the total number of transport sweep which does take into account the total number of angles.

Table 6 Comparison of transport sweeps for 361 group 1D BWR using different coarse group structures

CG Structure	CG Sweeps	FG Sweeps	Sweeps/FG	Tot. Sweeps*
2 Group	2,874	10,129	28.06	208,048
34 Group	22,415	4,269	11.83	426,944
Reference	--	128,381	355.62	2,054,096

* Total sweeps includes number of angles swept – (CG + FG) x Number of Angles

Although the number of coarse group sweeps increases when using the 34 group structure, the number of fine group sweeps is reduced. The result is convergence of the fine group problem, though only roughly 12 full fine group sweeps have been conducted. For the 2 group problem, 28 sets of fine group sweeps are required to converge. However, the additional effort placed on converging the 34 group problem unfortunately outweighs the reduction in fine group sweeps, leading to poorer performance relative to the 2 group problem when looking at the total number of transport sweeps. In this category, the 2 group problem reduces the number of sweeps by an order of magnitude, while the 34 group problem is only able to reduce this by a factor of 5.

A more detailed representation of the recondensation process is also provided in Table 7. This table compares the total number of transport sweeps conducted at each recondensation step and the difference between in eigenvalue relative to the reference fine group solution.

Table 7 Comparison of convergence for 361 group 1D BWR using different coarse group structures

Rec. Step	2 G – Sweeps	Δk (pcm)	34 G – Sweeps	Δk (pcm)
1	1632	93,064.9	30368	24,463.5
2	40656	2,263.5	129936	340.4
3	62256	928.1	169872	66.9
4	91632	165.7	278864	15.9
5	125040	34.6	388784	0.1
6	156160	2.0	406896	0.5

All the solutions begin very far away from the solution since a flat flux initial guess is used to seed the recondensation process. This provides very poor coarse group cross sections initially since flux dips at the resonance peaks are not being captured. The initial guess used here is the worst case scenario which means that performance could be further enhanced even using the simplest of self-shielding models.

While the 34 group case is able to drive the eigenvalue to within 1 pcm of the reference in only 5 recondensation steps, this requires roughly double the number of total transport sweeps relative to the 2 group case. Although the 2 group problem works better in reducing the total number of transport sweeps, it should be noted that 34 group recondensation reduces the number of fine group transport sweeps by over half that used for the 2 group problem. Therefore, if the 34 group problem were made cheaper, then this would lead to significantly better performance.

While the BWR problem was shown to converge within 1 pcm in Figure 41, recondensation only converges to within 10 pcm for the 2 group and 23 group structures in the HTR problem. This is very interesting since it has been shown that the simplifications made in Section 5.1.2 still allow SEAM to be fully consistent with the fine group problem. To verify that consistency is being maintained, the relative errors in the fine group scalar fluxes are compared in Figure 43.

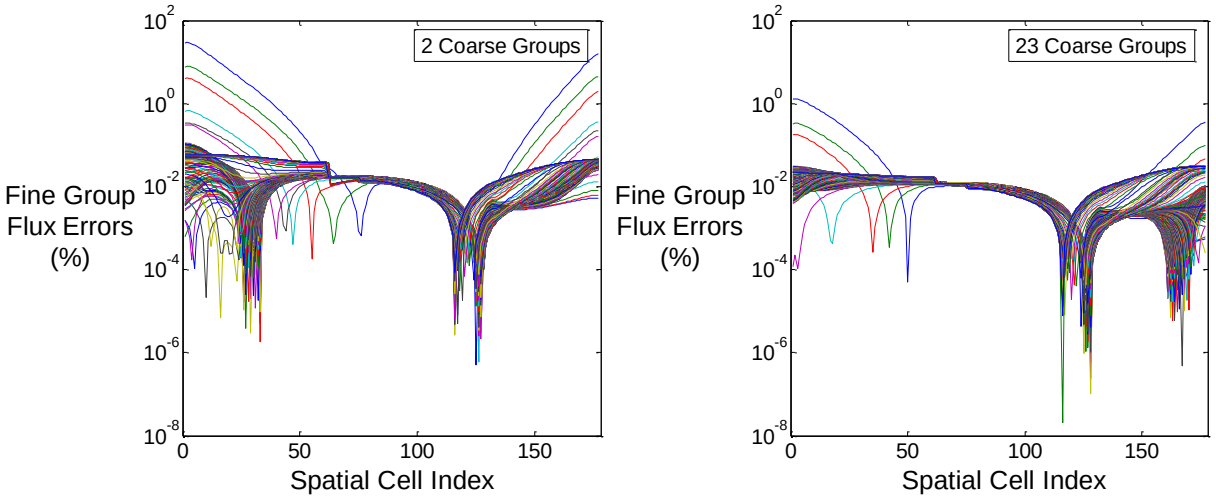

Figure 43 Relative errors in the fine group fluxes for the 295 group 1D HTR problem using 2 and 23 coarse groups.

In both the 2 group and 23 group case, although the relative errors in the fuel region (spatial indices 79 to 132) still remain near or below 0.01%, the errors in the reflectors reach 0.1% and, in the 2 group case, can reach up to 20%. This is likely due to the high scattering ratio present in the HTR problem. While the higher scattering ratio of graphite in the reflector makes the reactor much more efficient in terms of its neutron economy, this also makes the within-group more difficult to converge. Therefore, the maximum number of within-group iterations was increased to allow additional effort to be spent on converging the within-group scalar fluxes. This number of iterations was increased from 10 to 50 to 200 and finally to 3200, but using S_2 so as to converge without using more total sweeps than 200 with S_{32}.

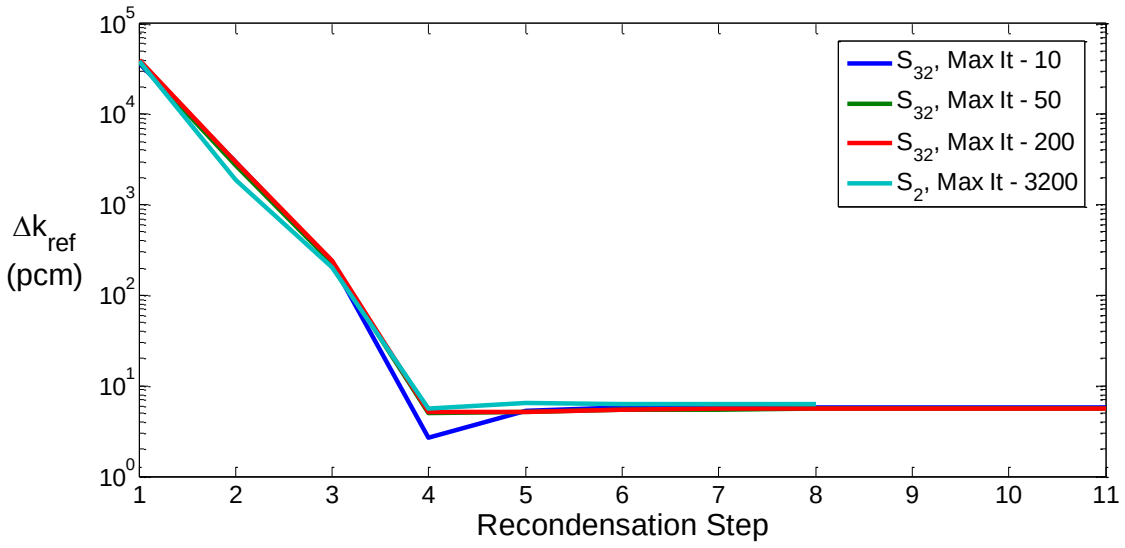

Figure 44 Comparison of the eigenvalue convergence when increasing the maximum iterations allowed in the within-group iteration and upscatter iteration.

Although increasing effort was placed on converging the within-group fluxes, no significant improvements were seen. While there is some variation in the number of recondensation steps required to converge, the error in the eigenvalue remains near 10 pcm.

Instead of focusing on possible issues with SEAM, the focus was switched to the fine group problem. If scattering is the issue, then the reference fine group solution itself may not be fully converged even though the same tolerances are used to obtain the solution in both cases. The tolerance used in obtaining the reference solution was further tightened from 10^{-7} to 10^{-8} and compared to the recondensation solution in Figure 45. The tolerances in recondensation were still kept at 10^{-7}.

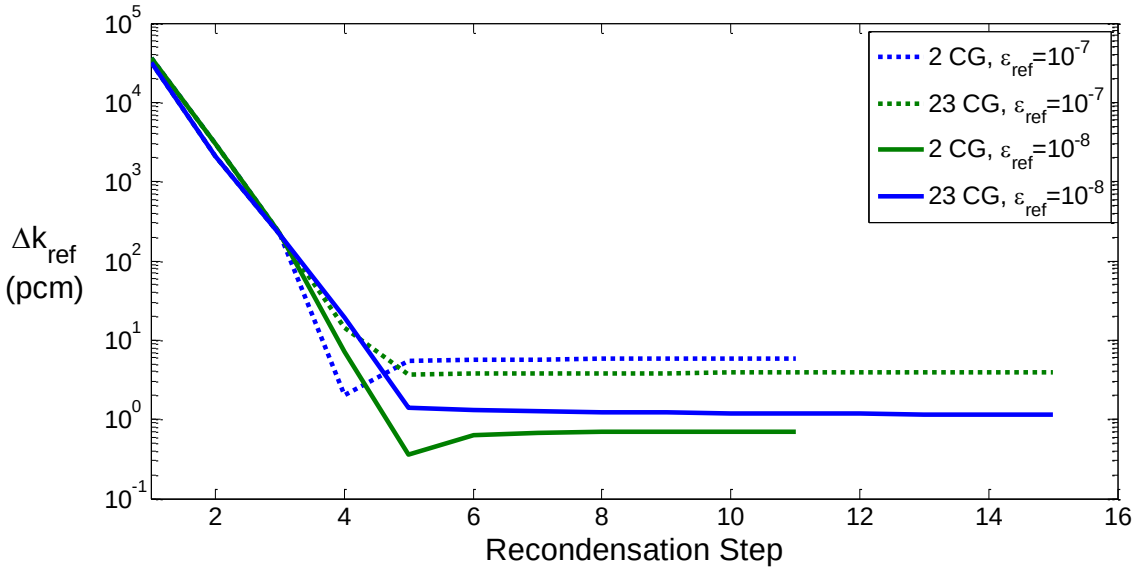

Figure 45 Comparison of the eigenvalue relative to the reference solution converged using a tolerance of 10-7 and 10-8

The results from this hunch prove to be very fruitful. When the reference eigenvalue is calculated using a tolerance of 10^{-7}, the recondensation eigenvalue is on the order of 10 pcm off from this reference solution. This is very similar to the behavior shown in Figure 41. When the reference case uses a tolerance of 10^{-8}, the recondensation eigenvalue actually matches to within 1 pcm. This suggests that the reference fine group solution was insufficiently converged to begin with, even though it uses Gauss-Seidel with upscatter and the same tolerance recondensation used. The scalar fluxes must be analyzed, though, to ensure that error cancellation has not given us falsely accurate eigenvalue.

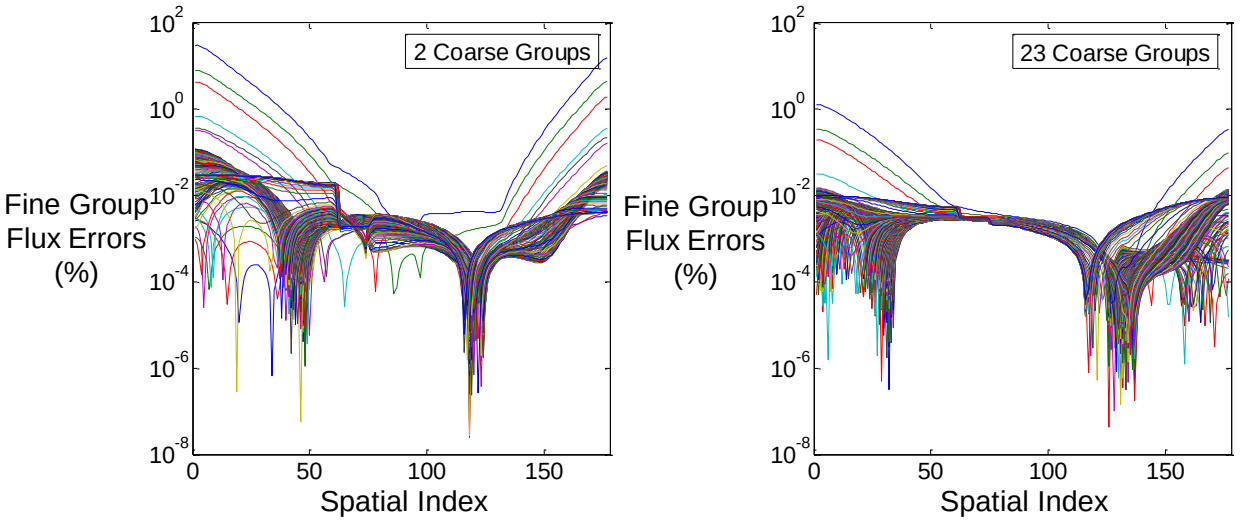

Figure 46 Fine group flux errors in the converged recondensation solution relative to the fine group problem converged with a tolerance of 10^{-8}.

In Figure 46 scalar fluxes in the 2 group and 23 group solutions are, for the most part, driven to or below 0.01% in the fuel regions and in the reflector regions immediately adjacent to the fuel. While much of the larger errors further out into the reflector still remain, it becomes apparent that this does not have a significant impact on the global solution. This also suggests that the source of the convergence issues in the reference fine group problem is being caused by the difficulty converging the scattering source. This appears to be eased when using recondensation.

To see why this is the case, the coarse group and fine group sweeps required to achieve convergence are compared in Table 8. For the 2 group case, the number of sweeps per coarse group is roughly 10,000 and in the 23 group case, this number is closer to 6,000. In both cases, this is much greater than the 2,000 sweeps applied per fine group in the reference solution. This means that the within-group problem is converged to a higher degree in recondensation, resulting in a significantly better spatial solution than would be obtained in the reference fine group solution.

Table 8 Comparison of transport sweeps for 295 group 1D HTR using different coarse group structures

CG Structure	CG Sweeps	FG Sweeps	Sweeps/FG	Total Sweeps
2 Group	21,808	11,947	40.5	540,080
23 Group	149,267	7,961	27.0	2,515,648
Reference (10^{-7})	--	622,793	2,111.1	9,964,688
Reference (10^{-8})	--	844,487	2,339.2	13,511,792

While the 23 group problem shows only modest improvements relative to the reference, a reduction by a factor of 6, 2 group recondensation is able to cut transport sweeps by a factor of over 20. The large number of transport sweeps added to the 23 group problem come from the need to sweep across 10 times as many group as the 2 group problem. Therefore, even though the 23 group sweeps are able to reduce the total number of fine group sweeps to 27, this is the done at the expense of more coarse group transport sweeps. Again, if the 23 group solution could be made significantly cheaper than the fine group transport sweeps, then this would become the better option of the two.

As expected, SEAM is able to reproduce the eigenvalue to within 1 pcm of the reference solution, as can be seen in Table 9. However, the 23 group SEAM calculation appears to converge a bit slower than the 2 group calculation once it drops below 100 pcm. A likely explanation is that the 2 group convergence dropped temporarily close to the reference eigenvalue and is actually converging slower.

Table 9 Comparison of convergence for 295 group 1D HTR using different coarse group structures

Rec. Step	2 G – Sweeps*	Δk (pcm)	23 G – Sweeps*	Δk (pcm)
1	27,568	36,066.7	29,056	31,400.2
2	154,896	3,031.9	519,568	2,090.3
3	261,968	228.9	1,134,304	221.3
4	386,096	7.1	1,775,808	19.5
5	480,768	0.3	2,410,784	1.4
6	509,136	0.6	2,450,144	1.3

* Total number of sweeps including angles

The HTR problem is distinct in this respect, since choosing the coarse group structure solely based off the fine group total cross sections does not produce significant improvements in the convergence of recondensation. This does make sense because, outside of the control material,

there are no significant changes in the total cross sections. This means that the shape of the scalar flux is defined primarily by the scattering and fission sources. While a different approach to choosing coarse group structures should be developed for optically thin problems, this issue remains outside the scope of this thesis.

5.2.2 Coupling to Different Polar Angle Quadratures

While these new multigroup cross sections were originally derived to maintain equivalence between fine group and coarse group problems, they allow for much more flexibility than just application for multigroup collapse. Since the errors initially corrected originated from inconsistencies in the streaming operator, it is also possible to take into account other approximations which give rise to similar errors. If the coarse group sweep is only dependent on the collapsed fine group source to achieve consistency, then it should be the case that any angular quadrature could be used in the coarse group problem without sacrificing exact equivalence. For this purpose, the consistent cross section can be defined according to the high order fine group scalar flux and the low order coarse group scalar flux according to Equation (5.35).

$$\begin{aligned} \nu\Sigma_{f,g}(\vec{r}) &= \sum_{K\epsilon g} \bar{\phi}_{g,HO}(\vec{r},K)\nu\Sigma_{f,g}(\vec{r},K)/\tilde{\phi}_{g,LO}(\vec{r}) \\ \Sigma_{s,g\prime\to g}(\vec{r}) &= \sum_{L\epsilon g\prime}\sum_{K\epsilon g} \bar{\phi}_{g\prime,\mathrm{HO}}(\vec{r},L)\Sigma_s(\vec{r},L\to K)/\tilde{\phi}_{g\prime,LO}(\vec{r}) \end{aligned} \tag{5.35}$$

To test this idea, recondensation was applied to the 361 group BWR problem using the 2 group and 34 group structures. This time, however, the coarse group problem was calculated using the angular quadrature in the reference problem, S_{16}, and a lower order quadrature, S_2. The results of this simplification in the coarse group problem can be found in Figure 47.

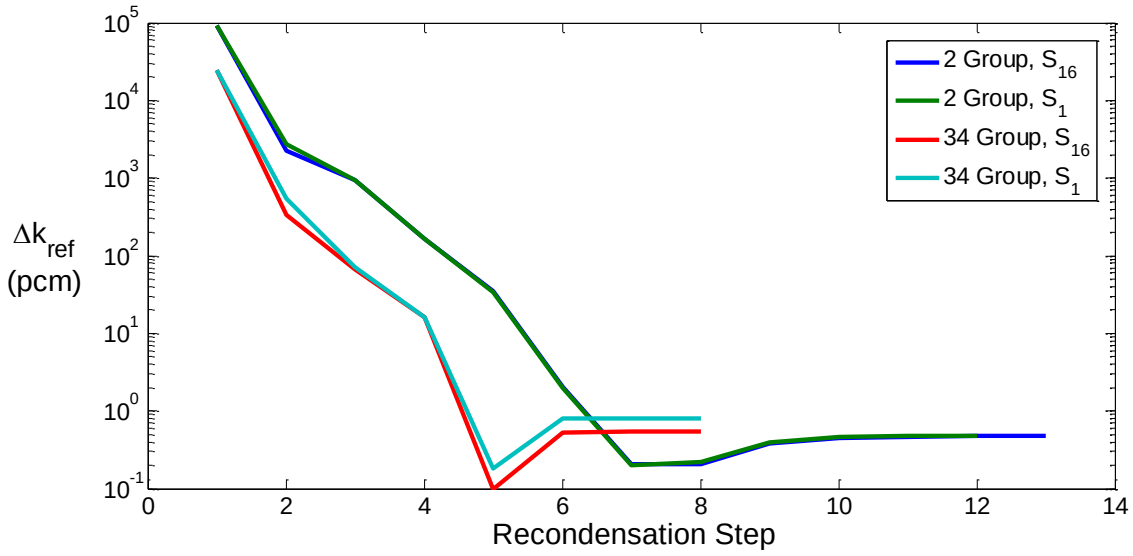

Figure 47 Comparison of convergence in k for the 2 and 34 group mappings, as well as S_{32} and S_2 couplings, relative to the reference 361 group BWR calculation

As expected, consistency is still achieved even when using a different angular quadrature in the coarse group problem. Since the convergence profile remains relatively unchanged when switching to a lower order angular problem, this suggests, at least in 1D calculations, that the angular dependence does not contribute significantly to the spatial problem. The distribution of sweeps across the coarse group solution also remains unaffected by the switch, as is highlighted in Figure 48.

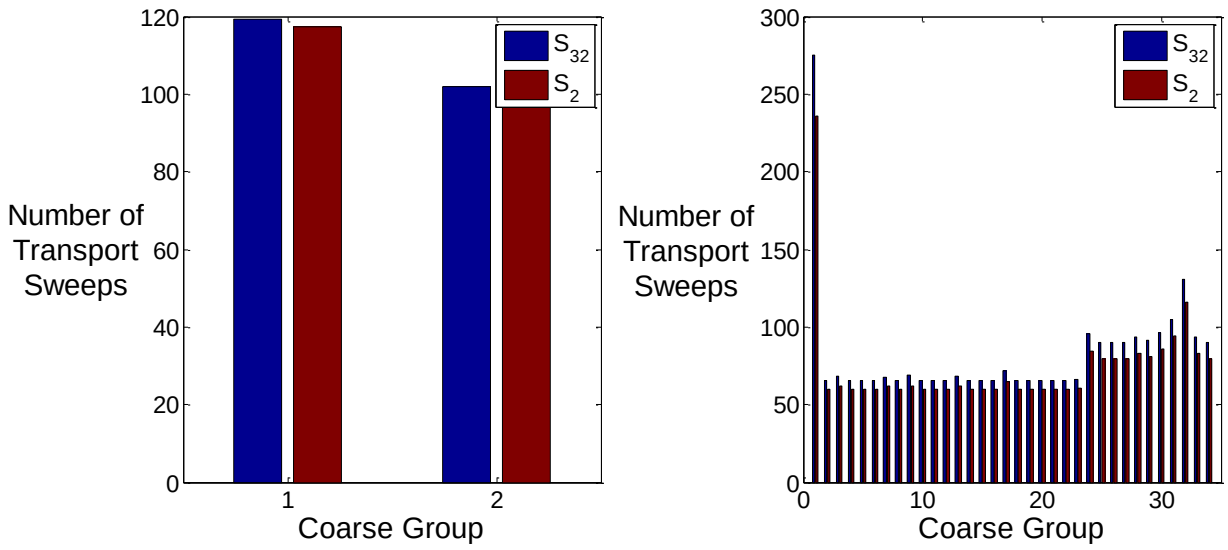

Figure 48 Distribution of the number of group sweeps per recondensation step for the 2 and 34 group mappings using the S_{32} and S_2 couplings

Little variation in the distribution of fine group sweeps is also observed in Figure 49. For the fine group fluxes above the upscatter cutoff, the number of sweeps per fine group remains the same regardless of the choice of angular quadrature and coarse group structure. For the fine groups in the upscatter regime, there is a slight deviation from the higher order angular calculation but this difference does not appear to be significant.

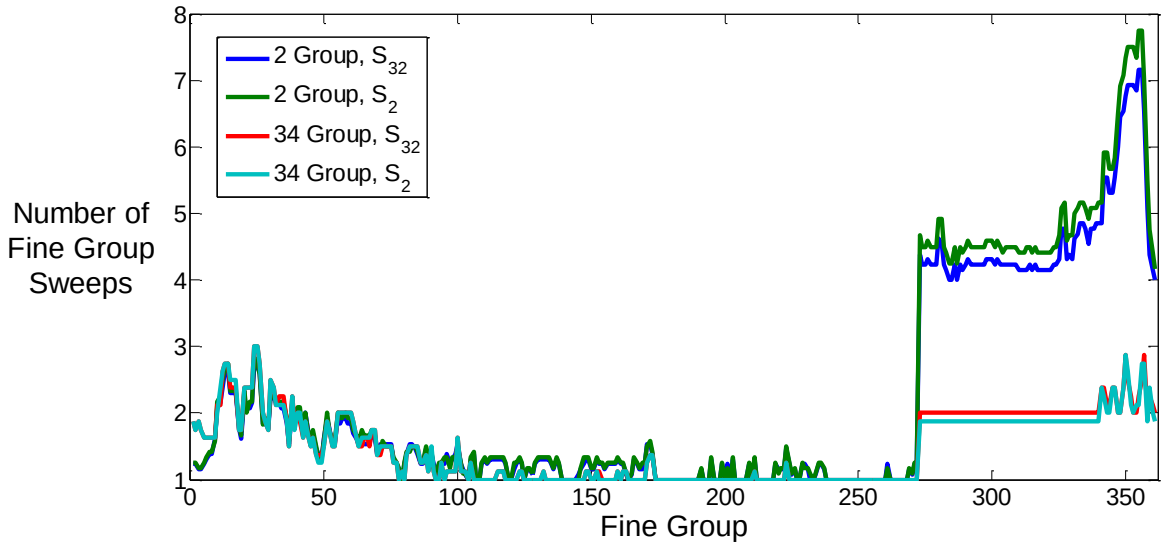

Figure 49 Distribution of the number of fine group sweeps per recondensation step for the 2 and 34 group mappings using the S_{32} and S_2 couplings

Now that consistency has been determined when applying differing quadrature in conjunction with multigroup collapse, the number of coarse group and fine group sweeps used during recondensation is compared in Table 10. The number of sweeps per fine group is also shown to highlight the total reduction in effort placed on the fine group problem.

Table 10 Total number of transport sweeps for the 361 Group 1D MOX BWR Benchmark

Method	CG Sweeps	FG Sweeps	Sweeps/FG	Total Sweeps
2 Groups, S_{32}	2,874	10,129	28.06	208,048
2 Groups, S_2	2,623	9,746	26.99	158,559
34 Groups, S_{32}	22,415	4,269	11.82	426,944
34 Groups, S_2	20,081	4,199	11.63	87,265
Reference	--	128,381	355.62	2,054,096

In the two group case, although the number of angles used in the coarse group sweep has been reduced by a factor of 16, the total number of sweeps is only reduced by roughly 25%. This is explained by the fact that computational effort in the two group problem is weighted heavily towards the single fine group source iteration. Only about 20% of transport sweeps are accounted for in the coarse group solve when the S_{32} quadrature is used. Since the fine group problem must be solved according to the desired quadrature, any further reduction in the angular

dependence of the coarse group problem will not result in significant reductions in the total number of transport sweeps.

However, the 34 group case is markedly different. When using an S_{32} quadrature, the coarse group problem makes up roughly 84% of the total number of transport sweeps. Therefore, the behavior observed when using the S_2 quadrature in the coarse group problem makes sense. Since computational effort is weighted significantly towards the coarse group problem, reducing the number of angles in the transport sweep significantly reduces the total number of transport sweeps.

This approach was also applied to the 1D HTR model to ensure this performance is not reactor dependent. When the low order angular approximation is applied to the coarse group problem, the convergence properties stay relatively the same as their S_{32} counterparts, as seen in Figure 50.

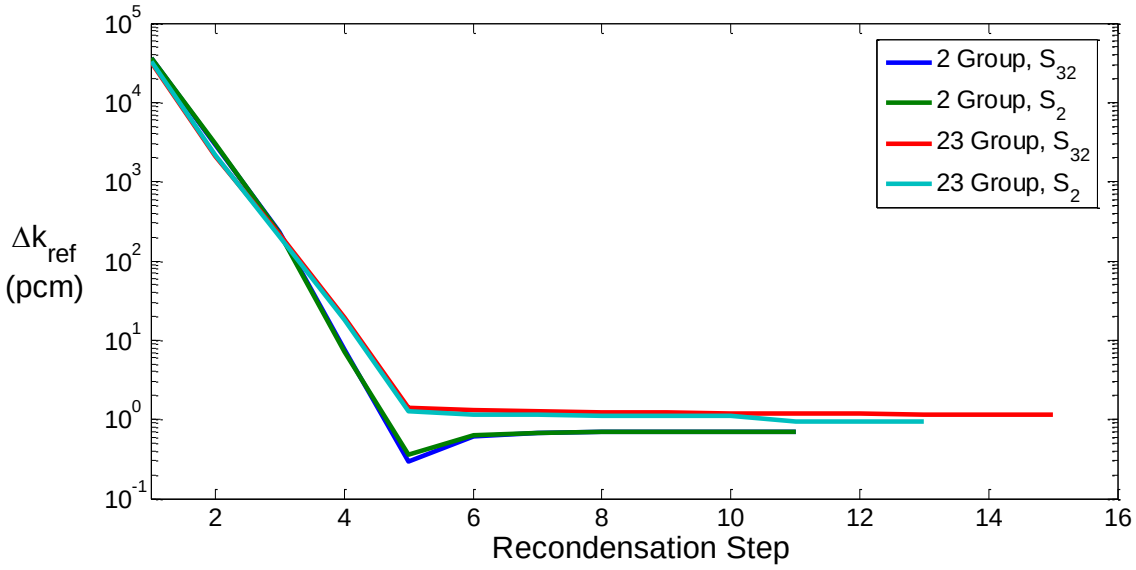

Figure 50 Comparison of convergence in k for the 2 and 23 group mappings, as well as S32 and S2 couplings relative to the reference 295 Group HTR calculation

The eigenvalues in both the 2 and 23 group cases still remain roughly 1 pcm away from the reference fine group calculation (using the 10^{-8} tolerance) conducted when the S_2 quadrature is

applied to the coarse group calculations instead of S_{32}. This shows that consistency is still maintained even when applied to a problem significantly different from the BWR model.

The number of coarse group and fine group sweeps is then compared in Figure 10 to identify gains in computational performance. Similar to the BWR calculation, the reduction in the total number of transport sweeps is much more apparent in the 23 group case than the 2 group case.

Table 11 Total number of transport sweeps for the 295 Group 1D HTR Calculations

Method	CG Sweeps	FG Sweeps	Sweeps/FG	Total Sweeps
2 Groups, S_{32}	21,808	11,947	40.50	540,080
2 Groups, S_2	22,253	12,041	40.82	214,909
23 Groups, S_{32}	149,267	7,961	26.97	2,515,648
23 Groups, S_2	153,331	7,356	24.94	271,027
Reference	--	844,487	2,339.2	13,511,792

The reduction in the 2 group case is slightly more pronounced because of the extra effort required to converge the within-group fluxes during Gauss-Seidel. When the S_{32} quadrature is used, the coarse group problem takes up roughly 64% of the number transport sweeps. In this case, since the number of sweeps in the coarse group problem is proportional to the number of angles, the maximum possible reduction is roughly half the total when using S_{32}.

When using 23 groups, the coarse group sweeps represent 94% of the total number of transport sweeps. One would expect that if the S_2 quadrature were used in place of S_{32} then the total effort to drop significantly more than in the 2 group case. Indeed, this is what happens. The total number of transport sweeps is reduced by almost a factor of 10.

It should be noted that for these cross sections to be fully consistent, the coarse group and fine group problems must stay on the same mesh. Otherwise, the collapse of the fine group source for use in the coarse group sweep may not yield a consistent coarse group flux for use in the cross section definitions.

5.2.3 Coupling to Different Spatial/Angular Methods

On top of applying a low order angular method, the argument made in Section 5.2.2 can be taken one step further. If the collapse of the fine group source is conducted in a consistent manner,

then the coarse group problem should be able to use a dissimilar spatial method from the fine group problem. This means that a fine group problem using step characteristics and an S_{32} angular quadrature could be coupled to a coarse group problem using step difference with an S_2 angular quadrature or even diffusion. All that would need to be done is collapse the fine group source and apply the coarse group sweep using the differing spatial method. Then the cross section can be defined according to Equation (5.36). In these equations *X* refers to the spatial method applied to the coarse group problem.

$$\begin{aligned} \nu\Sigma_{f,g}(\vec{r}) &= \sum_{K\epsilon g} \bar{\phi}_{g,SC,HO}(\vec{r},K)\nu\Sigma_{f,g}(\vec{r},K)/\tilde{\phi}_{g,X}(\vec{r}) \\ \Sigma_{s,g'\to g}(\vec{r}) &= \sum_{L\epsilon g'}\sum_{K\epsilon g} \bar{\phi}_{g',\mathrm{SC,HO}}(\vec{r},L)\Sigma_s(\vec{r},L\to K)/\tilde{\phi}_{g',X}(\vec{r}) \end{aligned} \tag{5.36}$$

To verify these claims, this approach is tested on the 1D BWR example. The convergence of the eigenvalue throughout the recondensation process is compared when using the following spatial methods for the coarse group calculations: step characteristics, step difference, diamond difference and finite-difference diffusion. For all the transport methods, an S2 quadrature is used for the low order angular problem. In the case of finite-difference diffusion, the coefficients are calculated without regard to any transport correction methodologies to truly show the universal application of this methodology. The diffusion coefficients for the coarse group problem are defined according to Equation (5.37).

$$D_g(\vec{r}) = \sum_{K\epsilon g} \bar{\phi}_{g,SC,HO}(\vec{r},K)\frac{1}{3\Sigma_{T,g}(\vec{r},K)}/\bar{\phi}_{g,SC,HO}(\vec{r}) \tag{5.37}$$

Regardless of the type of spatial method applied to the coarse group problem, the eigenvalue still converges to within 1 pcm of the reference case. Just as well, when transport methods are applied, the convergence plots match very closely to the original step characteristics approach.

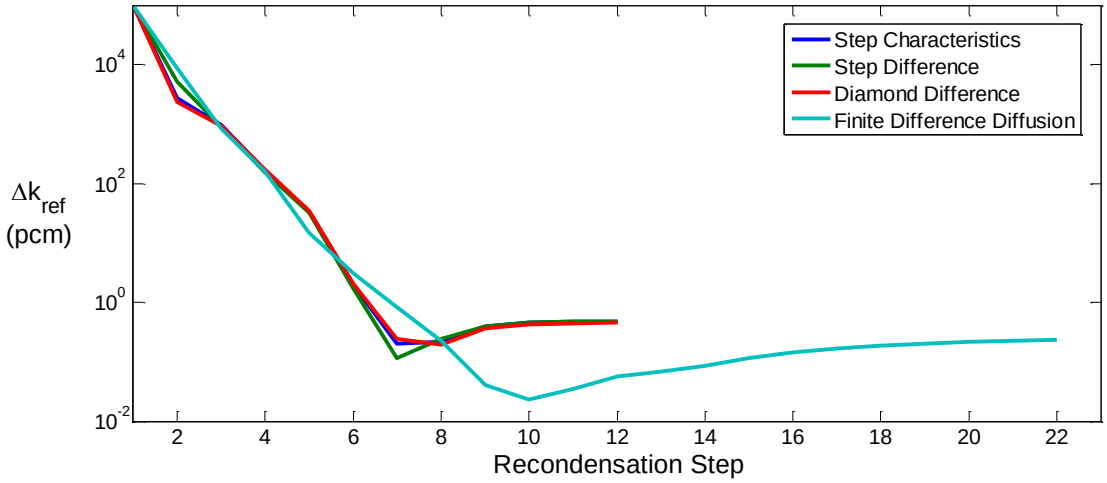

Figure 51 Comparison of convergence profiles when using dissimilar spatial methods for the coarse group eigenvalue problem. Comparison is relative to the reference 361 group BWR calculation.

However, while the diffusion case did eventually converge to within 1 pcm of the reference solution, a significant deviation occurs. In this case, recondensation requires nearly twice as many steps before the eigenvalue is converged. This has to do primarily with how the diffusion problem is solved in this case. For this comparison, diffusion is treated in a similar fashion. No scattering is included in the diffusion matrix when each coarse group flux is solved. This is obviously hindering the diffusion problem since the within-group scattering component of the source can be moved to the LHS, removing the need to converge a within-group flux. When this is done, the convergence profile matches those observed for all the transport cases.

While showing consistency is important, if using different spatial methods causes the number of fine group sweeps to increase relative to using the original spatial method, then this largely defeats the purpose of replacing an expensive computation with a cheap one. Therefore, in Table 12, the number of coarse group and fine group sweeps are compared when using the different spatial methods. One of the 34 group cases is also provided to show that finite diffusion can be used effectively when upscatter is included in the coarse group calculation.

Table 12 Comparison of transport sweeps for 361 group BWR model using different spatial methods

Method	CG Sweeps	FG Sweeps	Sweeps/FG	Total Sweeps
2 G SC	2,623	9,746	27.0	158,559
2 G SD	2,655	10,049	27.8	163,439
2 G DD	3,009	9,776	27.1	159,425
2 G FD Diff.	7,064	20,828	57.7	340,312
2 G FD Diff.*	470	10,205	28.27	163,750
34 G FD Diff.**	6,807	4,649	12.88	81,191
Reference	--	128,381	355.6	2,054,096

When the different transport methods are used, the number of coarse group and fine group sweeps does not deviate much from when step characteristics was applied. While this hasn't reduced the number of transport sweeps, it is important to note that both step difference and diamond difference do not require the evaluation of exponentials on the fly for every single sweeps. This makes this approach very appealing but it should be noted that an interpolation table can be built to evaluate exponentials without the overhead of intrinsic functions. [2]

However, if finite-difference diffusion is used instead, the number of transport sweeps doubles due to the increase in the number of recondensation steps. This increase is because the within-group convergence is being treated the same as it was in the transport sweeps. This is not truly representative of how a diffusion problem is actually solved, since the within-group scattering can be moved into the diffusion matrix. When the 2 group case is treated in this fashion, since no upscattering is present in this case, the number of coarse group solves reduces down to the number of power iterations. For the 34 group case, when within-group scattering is moved to the diffusion matrix, the total number of transport sweeps/diffusion solves is slightly improved with respect to step characteristics using the S_2 quadrature. One could move the entire scattering matrix to the diffusion matrix, but this is unnecessary for this example since the coarse group solve only account for less than 8% of the total number of sweeps/solves. Therefore, any further reductions in computational effort due to fewer coarse group solves would be small and, at least in this case, not worth the effort.

Again, to ensure that this performance is matched for different reactor types, this same methodology is applied to the 1D HTR example. The fine group problem is still solved using step characteristics S_{32}, but is coupled via the consistent cross sections with coarse group calculations employing step characteristics, step difference, diamond difference and finite-difference diffusion. For the transport methods, the S_2 angular quadrature was used and Equation (5.37) was used to define the diffusion coefficients for the diffusion methods.

The convergence of the eigenvalue with respect to the reference solution was compared for each of these cases for the 2 group problem in Figure 52. The finite difference approach was also applied after moving the within-group scattering source to the diffusion matrix to remove the within-group iterations. Although not shown in this figure, diffusion was also applied to the 23 group problem to show its use in the context of Gauss-Seidel with upscatter.

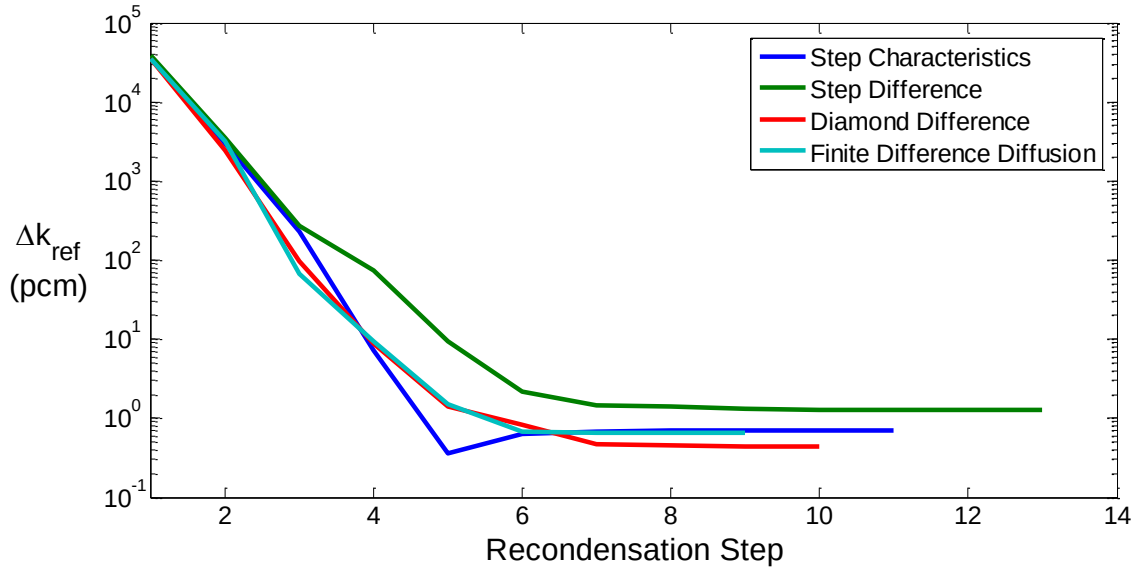

Figure 52 Comparison of eigenvalue convergence when coupling various spatial methods used on the coarse group eigenvalue solve to the step characters method used on the 295 group HTR model.

While diamond difference and finite-difference diffusion match the convergence profile for step characteristics closely, step difference requires a couple more recondensation steps to converge. This could be due to the presence of the strong absorber at the edge of the outer fuel region. It is not well-understood as to why this is the case, but the result is a significant increase in the total number of transport sweeps required to converge in recondensation. The finite-difference

approach also behaves differently. In the HTR example, applying the within-group iterations to the diffusion problem does not lead to an increase in recondensation step. In fact, diffusion matches the transport methods very well in this case. This behavior can be seen in Table 13.

Table 13 Comparison of transport sweeps for 295 group HTR model using different spatial methods

Method	CG Sweeps	FG Sweeps	Sweeps/FG	Total Sweeps
2 G SC	22,253	12,041	40.82	214,909
2 G SD	27,033	15,525	52.63	275,433
2 G DD	22,747	12,343	41.84	220,235
2 G FD Diff.	21,399	11,668	39.55	208,087
2 G FD Diff.*	310	13,451	45.60	215,526
23 G FD Diff.**	56,232	5,661	19.19	146,808
Reference (10^{-8})	--	844,487	2,339.2	13,511,792

* Scattering reduced Power Iteration

** Scattering reduced Gauss Seidel with Upscatter

Another interesting behavior is observed when within-group scattering is moved into the diffusion matrix; the number of coarse group solves decreases substantially while the number of fine group sweeps increases. This is unintuitive since one would expect better convergence of the within-group problem to positively impact the convergence of the fine group problem.

When diffusion is applied to the 23 group structure instead, diffusion's performance relative to transport actually improves, reducing the number of sweeps per fine group from about 25 to 19. This case produces the best improvement upon the reference fine group problem observed thus far. The total number of transport sweeps and the total number of fine group sweeps are each reduced by nearly a factor of 100. Further improvement could be made with the diffusion problem in this case if the full scatter matrix were included in this case, since the coarse group solves represent roughly 1/3 of the total number of sweeps and solves. Again, since the diffusion calculation is significantly cheaper than a transport sweep, the improvement this would have may be negligible.

5.2.4 Application to Linear Anisotropic Scattering in 1D HTR

So far, all of these examples have assumed that the sources in the core are solely isotropic. However, this approximation is not always satisfied in a real reactor. Therefore, the consistent multigroup cross sections are applied to the 1D HTR example when linearly anisotropic scattering cross sections are included into the model. While initially it would seem that consistency could not be achieved since the source is now coupled to angle, it is really only coupled indirectly. The linearly anistropic cross sections only produce a "moment" source upon which an evaluated basis function must be multiplied to actually produce the contribution to the full angular dependent source.

$$Q_g(\mu) = \sum_{K\epsilon g} Q_{0,g}(K) + \mu Q_{1,g}(K) = \sum_{K\epsilon g} Q_{0,g}(K) + \mu \sum_{K\epsilon g} Q_{1,g}(K) \tag{5.38}$$

Since these moment sources are decoupled from the angular dependence, direct summation of the fine group anisotropic source on the same mesh should provide the subsequent coarse group fixed source sweep with the correct information to produce the consistent multigroup cross sections.

However, producing a new collapsing procedure for the linearly anisotropic scattering cross sections is not straight forward. Although the first order flux moments are used to produce the anisotropic source, these fluxes are not guaranteed to be positive. This means that if we normalize the anisotropic scattering reaction rates with the sum of these fluxes across multiple fine groups, cancellations may occur and drive the denominator to or close to 0. To avoid this issue, the isotropic scalar flux is used instead to normalize the reaction rate. Therefore, the consistent formulation for the linearly anisotropic scattering cross sections is defined according to Equation (5.39).

$$\Sigma_{s,1,g\prime\to g}(\vec{r}) = \sum_{L\epsilon g\prime} \sum_{K\epsilon g} \bar{\phi}_{g\prime,1,\mathrm{SC,HO}}(\vec{r}, L)\Sigma_{s,1}(\vec{r}, L \to K)/\tilde{\phi}_{g\prime,0,X}(\vec{r}) \tag{5.39}$$

This equivalent cross section for the anisotropic HTR model was tested using the S_{32} and S_2 angular quadratures applied to the 2 group problem. S_2 was applied in this case since the Gauss-

Legendre quadrature is able to exactly integrate functions up to order 1. Therefore, this is the bare minimum number of discrete angles which can be used to accurately calculate the first order fluxes. If quadratic or cubic anisotropy were included an S_4 quadrature would be required to calculate the second and third order fluxes. The use of such a low order operator also highlights how flexible this method can be.

The convergence profiles for both sets of angular quadratures is calculated relative to the solution using tolerances of 10^{-7} and 10^{-8} as was initially compared in Section 5.2.1. This comparison is necessary since we originally showed in the isotropic case that the solution provided by recondensation actually matched the reference fine group problem better when the tolerance was set 10^{-8}. Since most of the physics from the isotropic case still applied to the anisotropic one, it is expected that the recondensation solution will match the 10^{-8} tolerance fine group solution better.

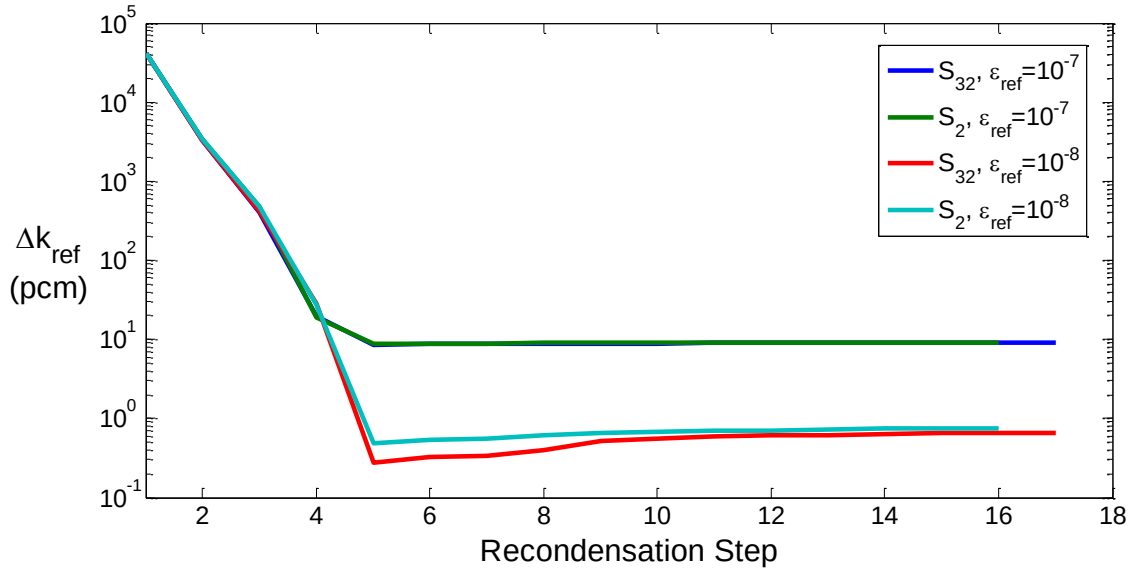

Figure 53 Convergence of the eigenvalue relative to the Linear Anisotropic 295 group HTR calculation converged to tolerances of 10^{-7} and 10^{-8}

As expected, the recondensation process is shown to match closer to the 10^{-8} reference solution in Figure 53. In this case, the eigenvalue matches the reference to within 1 pcm while, in the 10^{-7} case, the eigenvalue is off by about 10 pcm. This is impressive considering equivalence is still

maintained between the 295 group S_{32} problem and the 2 group S_2 problem when linear anisotropy is included.

The transport sweeps conducted for the coarse group eigenvalue solve and fine group source iterations are compared in Table 14 for the recondensation process.

Table 14 Comparison of transport sweeps for 295 group HTR model using linear anisotropic source

Method	CG Sweeps	FG Sweeps	Sweeps/FG	Total Sweeps
S_{32}	16,456	12,973	43.98	470,864
S_2	17,109	12,677	42.98	219,941
Reference (10^{-7})	--	760,281	2,577.2	12,164,496
Reference (10^{-8})	--	960,345	2,660.2	15,365,520

The number of sweeps for both the high order and low order angular coarse group problems roughly match those observed in the isotropic source case, even though the reference fine group problem requires significantly more transport sweeps to converge the solution. Therefore, even more improvements are realized when applying angular coupling to anisotropic scattering problem.

While application of this methodology has been shown for a linearly anisotropic problem, it should be noted that conducting anisotropic neutronics analyses in 1D is very straightforward. While it is expected that this same approach should work in 2D just as well as it does in 1D, the practical application will likely be much more difficult. This is because it requires use of the full set of spherical harmonics to represent the angular dependent source instead of simple Legendre polynomials. The approach taken in this section could be simplified if an isotropic source calculation were used in the coarse group problem instead and the angular moments were scaled using the scalar fluxes instead of the higher order angular moments.

5.3 Summary

In this section, a new method called SEAM was derived from the original work from the first attempts at making DGM fully consistent for 1D step characteristics. The del term was further simplified from its initial definition in Section 4.1 using intuition gained from previous work with DGM and also from SGD developed at Georgia Tech. These simplifications allowed a

drastic reduction in the storage required to produce a fully consistent transport calculation despite the inherent differences in the transport operator when moving from the fine group to coarse group problem. Instead of storing a corrective term for each cell, fine group, and discrete angle, this new factor only needs to be stored for each cell and coarse group.

Not only has this new method reduced memory requirements, but it also comes at little extra computational cost to the normal problem. All that is required is a fixed source coarse group sweep using the collapsed fine group source from the previous fine group sweep. Initially, the collapsed fine group flux from the fine group sweeps was divided by this intermediate coarse group flux to produce a new equivalence factor. This factor is multiplied with the coarse group scalar flux after each coarse group sweep to produce a corrected coarse group flux. Further inspection revealed that this procedure was identical to using the fixed source coarse group flux to weight the fine group reaction rates in the group collapse procedure instead of the collapsed fine group flux.

In both cases, the main effect is a recreation of the fine group source in the coarse group sweep even though the coarse group transport operator isn't consistent with the fine group. Thus it was coined the Source Equivalence Acceleration Method (SEAM). This method was found to be more general in its application, though, than just correcting for errors accrued in the group collapse process. By definition, the formulation of the Source Equivalence Factors allow for one to use any inconsistent method, whether it be a lower angle approximation or an entirely different method altogether from the one used in the fine group calculation. This provides SEAM with a simplicity and flexibility that is astounding.

SEAM was then applied to two 1D test problems, a 361 group BWR core and a 295 group HTR core to highlight. Although the physics behind each of the cores are markedly different, a sizeable reduction in fine group transport sweeps was observed across the board. On top of correcting the errors accrued in the typical group collapse with step characteristics, the flexibility of SEAM was tested by coupling high-order to low-order angular approximations, step characteristics to step difference, diamond difference and even finite-difference diffusion. These examples culminated in the successful coupling of a 295 group, S_{32} calculation using step

characteristics to a 2 group, S_2 calculation using step difference while including linearly anisotropic scattering.

Up to this point, though, SEAM's favorable characteristics have only been demonstrated using a 1D transport code developed in MATLAB. While MATLAB served as an ideal language to test out these ideas and concepts, no time comparison were provided since it suffers from the same performance deficiencies as other interpreted programming languages. Even relatively small 1D calculations took on the order of minutes to converge. Therefore, a new code was chosen for further application into 2D transport. Since the original intent of this thesis was focused on spatial consistency using higher order methods, OpenMOC was chosen for its use of the Method of Characteristics (MOC) in 2D and its potential for large-scale parallelization. Since it is largely writing in C++ and includes the option of applying Coarse Mesh Finite Difference (CMFD) acceleration, it serves as the best candidate by which to accurately assess the performance of SEAM and compare it with CMFD.

6 2D Testing of SEAM using OpenMOC

While the derivation and testing of SEAM conducted in 1D transport calculations produced very favorable results, this method is of little use if it does not produce the same level of acceleration or better than the methods currently used to accelerate 2D transport calculations. CMFD, for example, has shown incredible flexibility in its ability to accelerate transport calculations. In OpenMOC, for example, CMFD has been implemented and shown to reduce the number of transport sweeps by nearly a factor of 40. [2] The objective of this section, therefore, is to determine SEAM's ability to accelerate fine group 2D transport calculations by implementing it in OpenMOC and then compare its performance to that of CMFD acceleration.

6.1 A Brief Introduction to OpenMOC

OpenMOC makes use of the transport method called the Method of Characteristics (MOC). Instead of solving for the average angular flux across an entire cell one at a time, MOC tallies angular fluxes which are solved across many tracks superimposed upon the problem. One of the primary simplifications in standard MOC is that the source in each cell is spatially flat. Therefore, each cell in an MOC calculation is a flat source region.

Along each track, the angular flux is calculated using the same process as 1D step characteristics. The incoming angular flux and the source along the direction of the track are used to calculate the outgoing angular flux and the segment average angular flux. The segment is defined as a distinct portion of a single track spanning across a single cell. In order to calculate the full cell-averaged scalar flux, these segment angular fluxes must be integrated across the angular and spatial domain. Calculating this integral exactly is computationally intractable, so multiple tracks are generated across a set of discrete azimuthal angles to carry out this integration. These tracks are also generated in the transverse direction in order to cover the full spatial domain of the problem. OpenMOC generates these tracks according to a cyclical characteristics condition. The key advantage of using cyclic characteristics is that reflective boundary conditions can be accurately described. This is done by propagating a track at a specific azimuthal angle and

allowing it to reflect at the boundaries of the problem. Since this is the exact path a neutron would take when hitting a reflective boundary condition, this guarantees that the boundary condition is being accurately characterized. In order to conserve memory, however, the track must eventually terminate either where the propagation was initialized or at the beginning of some other track. The result is a discrete set of allowable angles generated that approximately match the number of azimuthal angles and track spacing specified by the user.

Each track is then split into segments corresponding to the section of the track passing through a single cell. Once the average angular fluxes across all the segments of each track are calculated, then summation according to Equation (6.1) is used to produce a scalar flux for each cell. In this equation, *i* denotes the cell, *m* refers to the azimuthal angle, *k* designates the segment and λ is the index for the polar angle. Therefore, $\bar{\psi}$ is the average angular flux for each azimuthal angle, segment and polar angle. The cell *i*'s volume is V_i and the mapping from segment *k* to cell *i* is labeled using the Kroenecker delta function $\delta_{i\,N_k}$.

$$\phi_i = \frac{1}{4\pi V_i} \sum_m \omega_m \sum_k \delta_{i\,N_k} \sum_\lambda W_\lambda \bar{\psi}_{m,k,\lambda} \tag{6.1}$$

Here, the azimuthal weights, ω_m, are divided according to the spacing between each nearest angle, $\Delta\varphi$, effectively creating a Riemann sum to conduct the azimuthal integration. This weight is also multiplied by the effective spacing, $\Delta s_\perp$, for that angle to include transverse integration in space.

$$\omega_m = \Delta\varphi \Delta s_\perp \tag{6.2}$$

The polar angle and weights, W_λ, is set using a predetermined angular quadrature. OpenMOC allows the user to specify if one would like to use Tabuchi-Yamamoto (TY) or Gauss-Legendre (GL) quadrature. TY is most popular in use since it was derived to best approximate the Bickley-Naylor function used in the Collision Probability Method (CPM) calculations. Use of the TY quadrature therefore allows MOC to nearly achieve the accuracy of CPM without having to invert a dense matrix at each power iteration or sacrifice accurate characterization of reflective boundaries. [37]

Of course, OpenMOC doesn't store the average angular flux for every single segment before conducting the summation. The segment angular fluxes are tallied on the fly according to Equation (6.3) as the transport sweep is conducted.

$$\phi_i = \frac{Q_i}{4\pi\Sigma_i} + \frac{1}{4\pi V_i}\sum_m \omega_m \sum_k \delta_{i\,N_k} \sum_\lambda W_\lambda \left(\psi_{in,m,k,\lambda} - Q_i\right) \frac{(1 - e^{-\tau})}{\Sigma_i} \tag{6.3}$$

In order to minimize memory requirements, the exponentials used in the scalar flux tally must be calculated on the fly, making this the most expensive aspect of the transport sweep. In OpenMOC, this cost is minimized by use of linear interpolation of the exponential function based on the maximum optical length appearing in the problem. This produces a cost efficient way of accurately calculating exponentials on the fly.

Currently, OpenMOC relies on power iteration to solve the eigenvalue problem. Thus, after all transport sweeps are conducted and the scalar fluxes have been updated, k_{eff} is updated for the next iteration. This is calculated directly using the Rayleigh Quotient, which determines the current estimate of the eigenvalue problem through the ratio of the volume integrated fission rate to the integrated absorption rate and net leakage from the system.

$$k_{eff} = \frac{F}{A + L} \tag{6.4}$$

Once k_{eff} is updated, the scalar fluxes are normalized using the volume integrated fission rate before constructing a new source from the updated fluxes for the next set of transport sweeps.

6.2 New Features Implemented to Test SEAM

While many useful features, such as CMFD, have been recently added to OpenMOC, some of the methods and approaches specific to SEAM did not exist in the code. Therefore, before the flexibility and feasibility of SEAM in 2D could be tested, the following functionality was incorporated.

6.2.1 Group Collapse for Transport Problems

Previously, OpenMOC had no support for cross section condensation other than to one group. What made inclusion of SEAM less intuitive was the fact that group collapse required information to be passed between two different transport solver objects. Although homogenization had already been implemented in the CMFD class in OpenMOC, the CMFD method itself is more self-contained than SEAM. A completely separate CMFD class was created in order to incorporate homogenization and diffusion solver into the code. Therefore, a whole new approach to incorporating condensed cross sections was required since both the coarse group and fine group problem operate on the same unstructured mesh. Creating a whole new class of solver was avoided by incorporating much of SEAM's functions, including group collapse, into the Geometry class. Since both solvers rely on the same spatial mesh, and the track information is kept separate from Geometry object, the functions added to the Geometry object would have access to the scalar fluxes in both solvers. In addition to that, the separation of the track generator from the geometry allows one to use two completely different tracks in both solvers without needing a second Geometry object. This setup seemed ideal at the time but there will likely be plenty of room for optimization in the near future.

6.2.2 Gauss-Seidel with Upscatter Option

This addition to the code was deemed necessary due to the enhanced performance observed in the 1D calculations when Gauss-Seidel was used in the fine group problem in place of the normal power iteration. This proved difficult because the original transportSweep() function was originally developed to sweep across all energy groups simultaneously to take advantage of SIMD vectorization. Therefore, the groupTransportSweep(int g) function was added into the ThreadPrivateSolver class to sweep across energy groups individually. Power iteration using Gauss-Seidel was implemented as a separate function, convergeUsingGS(), in the Solver class.

While Gauss-Seidel is fairly straightforward to implement in diffusion calculations, incorporation into transport calculations require a bit more effort. Since the full scattering matrix can't be moved explicitly to the LHS of the transport equation before the transport sweep is

applied, fixed point iterations, called within-group iterations, are applied to converge the within-group scattering source. If the scattering matrix is completely lower triangular, meaning only downscattering exists, then scattering source can converge using only *G* sets of within-group iterations. However, if upscattering exists, another iteration must be nested within the downscatter loop in order to sufficiently converge the upscatter source.

6.2.3 Step Difference Transport Kernel

While MOC uses the characteristic definition of the angular flux according to Equation (4.1), in all practicality any method that provides an average angular flux can be used in tallying the scalar flux according to Equation (6.1). Therefore, a step difference scalar flux tally function was added into OpenMOC to modify the outgoing and average angular flux calculations across a given track segment.

$$\bar{\psi} = \frac{\mu}{\mu + \tau}\psi_{in} + \frac{l}{\mu + \tau}Q \tag{6.5}$$

At first, it may seem odd that such a low order method was added into OpenMOC in the first place, but this functionality was included for two reasons. First, step difference was included to show that SEAM can still couple a high order transport problem to any other method. Theoretically, SEAM should be able to couple a fine group calculation using MOC to a less expensive transport calculation such as step difference. This should make the coarse group transport calculation even cheaper. Second, the coarse group total cross sections change after each recondensation step. Since the exponential interpolation table is built using the maximum exponential term and a set precision level, the table itself might need to be updated if the values of the coarse group total cross sections increase substantially between recondensation iterations. Using SEAM to create an equivalent step difference problem sidesteps both issues by removing the need to calculate exponentials altogether in the coarse group transport sweep. Since coupling step characteristics with step difference in 1D did not adversely affect the convergence of recondensation, this appears to be a good option for making the coarse group problem cheaper to solve.

6.2.4 Fine Group Volumes and Leakage Scaling for the Rayleigh Quotient

One primary difference between OpenMOC and the 1D code tested in MATLAB is in how the eigenvalue is updated. In the 1D code, the current eigenvalue was estimated using the ratio of the total fission rate of the current iteration to the total fission rate of the previous iterate. Since the corrective factors used in SEAM effectively create a coarse group source equivalent to the fine group source, this allows the eigenvalue to be updated in a manner consistent with the fine group problem. However, in OpenMOC, the eigenvalue is updated through the Raleigh Quotient defined in Equation (6.4). While the fine group fission and absorption sources can be equivalently represented in the coarse group problem, the net leakage is not guaranteed to be equivalent. This differs from CMFD in which the net current across each surface of a cell is preserved. Since preservation occurs across each surface, then the net leakage across the boundary of the full boundary can be preserved as well. On the other hand, SEFs only provide one additional degree of freedom per spatial cell, allowing only the net leakage for each cell to be preserved. Therefore, an additional factor, the Leakage Scaling Factor, is defined to enforce an equivalent fine group net leakage in the coarse group problem.

This factor is defined by the ratio of the fine group leakage from the previous fine group sweep to the coarse group leakage from the fixed source coarse group calculation. This factor is an integral quantity defined according to Equation (6.6) specifically for a 2D problem. In 2D, the net leakage from the boundary of the problem is defined by a line integral across the boundary, s_b, of the geometry and an integral over the angular half-space $\vec{\Omega} \cdot d\vec{s} < 0$ (all directions leaving the boundary surface at a given position).

$$f_{leakage} = \frac{\sum_K \int_{\vec{\Omega}\cdot d\vec{s}<0} \oint_{s_b} \psi(r,\Omega,K)\vec{\Omega} \cdot d\vec{s}}{\sum_g \int_{\vec{\Omega}\cdot d\vec{s}<0} \oint_{s_b} \psi_g(r,\Omega)\vec{\Omega} \cdot d\vec{s}} \tag{6.6}$$

Also, since it is not guaranteed that tracks for every angle will cross every cell on the boundary, the net leakage for these cells might be inconsistent with the fine group problem. Assuming the fine group problem uses a fine angular discretization which can conduct the full angular

integration in Equation (6.6), the Leakage Scaling Factor will also be able to ensure that net leakage is preserved in the coarse group problem when using a coarse angular discretization.

Since the Rayleigh Quotient is defined using volume integrated reaction rates, the volumes used in the fine group calculation must also be used in the coarse group calculation, even if the coarse group calculation uses a coarser angular discretization or track spacing. Incorporating all of these factors leads to a modified form of the Rayleigh Quotient for the coarse group problem.

$$k_{eff} = \frac{F_{CG}(V_{FG})}{A_{CG}(V_{FG}) + f_{leakage} L_{CG}} \tag{6.7}$$

It should be noted that all the reaction rates and net leakage used to update k_{eff} in Equation (6.7) are calculated at each power iteration in the eigenvalue solve for the coarse group problem.

6.3 Application of SEAM to the C5G7 Benchmark Core

The C5G7 benchmark problem is an extremely well documented and analyzed MOX-loaded core commonly analyzed when developing new methods. This benchmark originated primarily as a means to test the ability of codes to analyze non-homogenized problems. This quarter core model is comprised of 4 square assemblies, each 21.42 cm wide. Each assembly contains a 17 by 17 lattice of square pin cells, defined using a pitch of 1.26 cm. A 21.42 cm wide water moderator region is also included and surrounds the inner 2 by 2 array of fuel assemblies.

The core itself is comprised of two assembly types. One is an assembly containing solely UO2 fuel, while the other contains a mix of three different enrichments (4.3%, 7.0% and 8.7%) of MOX loaded fuel. Two other materials are included which describe the properties of the guide tubes and fission chambers in the fuel assemblies. The benchmark includes a 7 group cross section data set for each of these 7 materials in the core. [25]

In addition to including many of the basic physics of light water reactors, the unstructured nature of this benchmark is ideal for analysis with MOC. All of the C5G7 calculations used in this thesis, including the normal power iteration, SEAM and CMFD were run using a single node

containing two six-core processors on the Nuclear Science and Engineering educational cluster. A total of twelve threads were used on each of the C5G7 runs.

6.3.1 Energy Acceleration of the C5G7 Benchmark Using SEAM

The C5G7 benchmark problem was first tested using a track spacing of .14 cm and only 4 total azimuthal angles. A whole core version of the C5G7 core was used in place of the typical quarter core model. This is the result of a slight inconsistency in how SEAM is applied and how OpenMOC conducts its transport sweep of the core. As it was built for parallelization, the transport sweep is initialized using the boundary angular flux from the previous iteration and is propagated down multiple tracks at once until reach their respective ends. Those same tracks are swept backwards using the boundary angular flux at the opposite end. To enforce consistency using SEAM with this setup, the fine group boundary angular fluxes must also be collapsed at the boundaries in addition to the sources within the flat source regions. Unless the same track spacing and number of azimuthal angles are used, there isn't a clear cut way to ensure that the fine group boundary angular fluxes are collapsed in a consistent manner along, say, a reflective boundary condition. Since a whole core uses vacuum boundaries to enforce the same condition across the entire core periphery, this was the ideal approach to test SEAM in OpenMOC.

SEAM was then used to accelerate the normal 7 group problem by creating an equivalent 1 group problem at each recondensation step. The same geometry and tracks were used in the 1 group problem as the 7 group problem. Instead of using MOC, though, the 1 group problem used step difference with only 1 polar angle. The converged flux and coarse group fission cross section from this calculation are shown in Figure 54, along with a plot of converged fission rate distributions calculated using the converged solution.

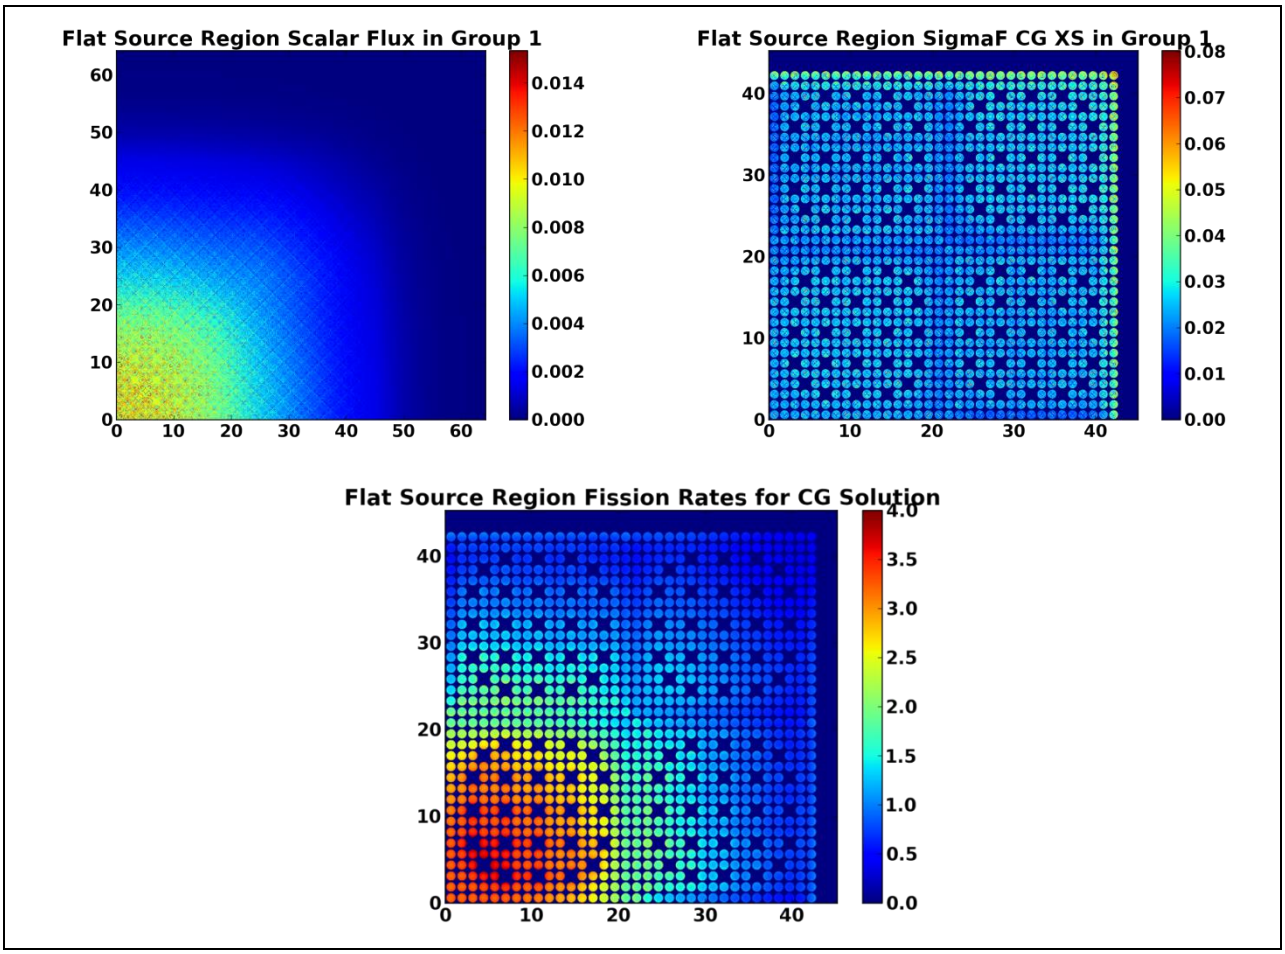

Figure 54 Comparison of the scalar flux and the fission cross section for the 1 group problem. Note that while both the fluxes (top left) and cross sections (top right) look rather non-physical, the product of the two produces the correct fission rates (bottom center) across the C5G7 core.

One of the interesting things to note from the scalar flux and fission cross section is that, while the global solution looks reasonable, the local solution looks unphysical. While the flux still peaks in the center (top left image in Figure 54) and has a roughly cosine shape to it, the ray effects in the flux almost appear to be amplified; a square grid structure appears. The fission cross sections (top right image in Figure 54) also peaks at the core periphery, which is to be expected because of the influx of thermal neutrons at the assembly edges, but the variation of the cross section within the assemblies are non-intuitive. It should be noted that although the fission cross sections are higher at the periphery, this doesn't necessarily mean the fission rates will be

highest. On the contrary, while the spectrum does shift at the periphery, the fluxes at the core edge are significantly lower than those in the center.

Since SEAM does not homogenize the spatial problem, there is a one-to-one mapping between the fine group and coarse group geometries. Therefore, a detailed analysis of the fission error distribution can be made down past the fuel-pin level into the individual flat source regions.

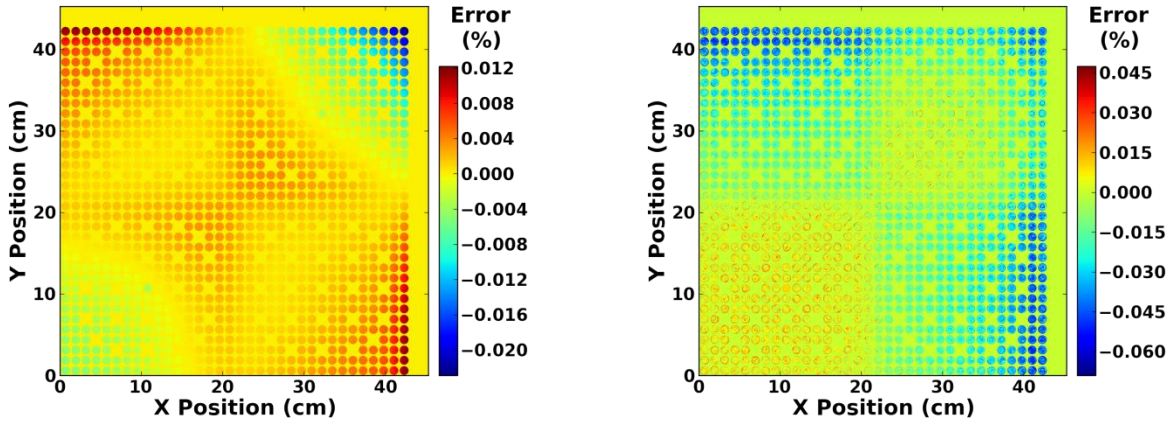

Figure 55 Percentage errors in SEAM 1 Group fission rates relative to the solution using normal power iteration converged with a 10^{-5} tolerance (left) and a 10^{-6} tolerance (right)

Figure 55 shows that, while the flux and cross sections may look non-physical due to apparent artifacts in both distributions, the product of these values form a near perfect match to the flat source regions fission rates calculated using the normal power iteration on the fine group problem. The errors in the fission source relative to the normal 7 group calculation remain below 0.1%. This is impressive considering the SEAM mostly used power iteration of the 1 group problem using step difference and only 1 polar angle to converge the 7 group problem.

However, this calculation did use a source convergence criterion of 10^{-5}, so it was initially expected that the fission rate errors would drop well below .01%. In order to verify that equivalence has been maintained, the differences in the fission rates of the unaccelerated fine group solution using a source convergence criterion of 10^{-5} and 10^{-6} were calculated. If the errors in the SEAM calculation are on par with these errors, then this provides a good indication that equivalence has been maintained.

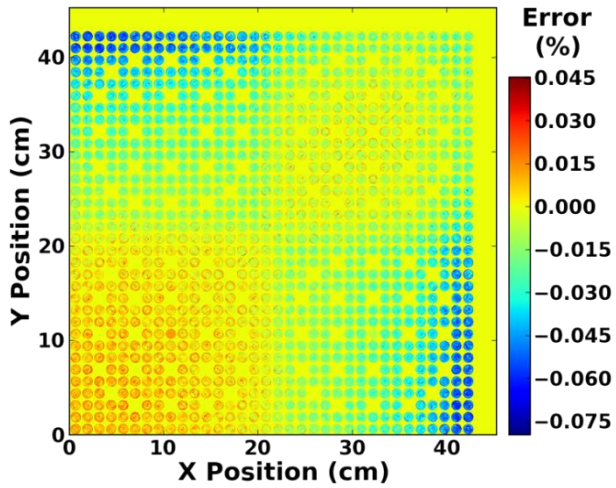

Figure 56 Errors in fission rates using normal power iteration when using a 10^{-5} tolerance as opposed to 10^{-6} for a 4 azimuthal angle, 0.14 cm track spacing calculation

The fission rate errors between the two unaccelerated solutions in Figure 56 match very closely to the error distribution on the right side of Figure 55. While the errors are certainly greater than initially expected, the comparison of the unaccelerated and accelerated solutions relative to the more precise solution indicates that equivalence has been maintained in the 1 group calculation using SEAM.

After equivalence was shown, 1, 2 and 4 group SEAM calculations were conducted to determine the relative impact of coarse group structure on acceleration. This was done by comparing the convergence profiles for all 3 calculations. Again, in each SEAM calculation, the primary problem solved used the same number of azimuthal angles and track spacing as the 7 group problem, but conducted the transport sweep using step difference and 1 polar angle. A stability issue did arise in subsequent SEAM calculations and for this reason two successive power iterations are used between recondensation steps. This issue will be explained in detail in Section 6.3.2.

Since CMFD represents one of the most widely used methods to accelerate transport calculations, and as such, represents the best choice to compare with in terms of pre-existing methods. Therefore, the results using SEAM were compared to a 7 Group CMFD accelerated

calculation using pin cell homogenization. This was deemed to be the best possible acceleration achievable by CMFD for the C5G7 problem since the "coarse" mesh provides a detailed representation of the averaged pin cell fluxes and the lack of group collapse allows the spectral effects to be fully accelerated. For this application, the relaxation factor was set to 0.6 to stabilize CMFD acceleration.

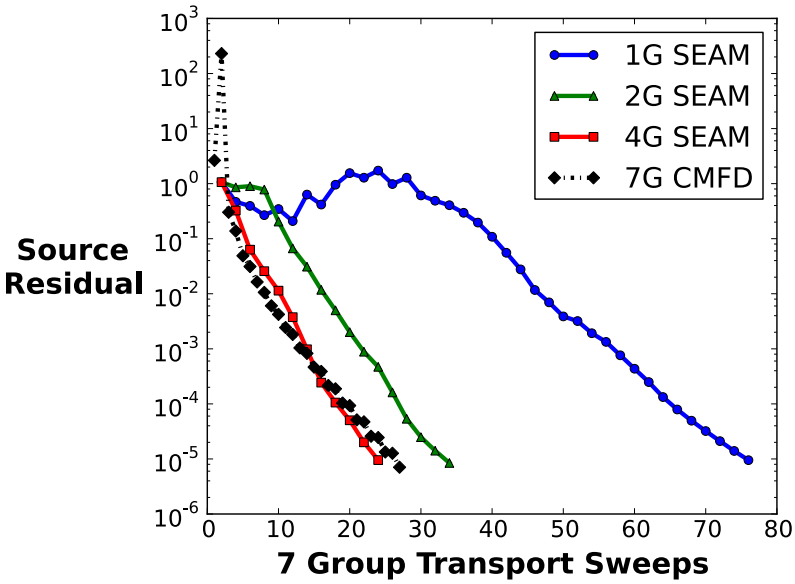

Figure 57 Convergence profiles for SEAM using different coarse group structure and 7 Group CMFD using pin cell homogenization

As expected, the ability of SEAM to accelerate improves as the number of coarse groups increases. Stability also seems to improve when using more groups. In the 1 group calculation, oscillations occur at first before finally converging and for the 2 group calculation there is an initial stagnation in source residual before convergence begins. When SEAM uses 4 groups, though, convergence is stable and appears to be nearly monotonic. In this case, SEAM is able to converge the solution while only using 24 7 group transport sweeps. This is actually a slight reduction from the 27 sweeps required by CMFD. Therefore, even when the cross sections are collapsed to 4 groups, the detailed acceleration of the FSR fluxes is able to match the performance of pin cell homogenized CMFD.

However, the convergence profiles do not provide the best indication of true performance with respect to time. Although SEAM was able to match the performance of CMFD, it is extremely important to compare the time required to conduct the power iterations since the transport sweeps used in SEAM computationally expensive relative to the diffusion calculations used in CMFD.

Table 15 Comparison of performance between PI, SEAM and CMFD for the 4 az, 0.14 cm C5G7 problem

Method	FG Sweeps	FG Time	CG Sweeps	CG Time	Total Time
Normal PI	768	502.6	N/A	N/A	502.6
1 Group	76	51.8	6519	1930.2	2015.4
2 Group	34	21.5	1507	563.4	601.8
4 Group	24	15.6	939	441.1	468.4
CMFD	26	19.7	26*	52.4	72.1

* The number of diffusion eigenvalue solves used in CMFD

In fact, in Table 15, the diffusion solve in CMFD is able to beat the SEAM transport sweeps by nearly a factor of 10. The relatively cheap diffusion calculations in CMFD are key to its success in accelerating transport calculations. Since CMFD is able to create a cheap yet equivalent diffusion problem, SEAM's equivalent transport problem can't compete in terms of total time.

However, it is important to remember that SEAM is not being used to its full ability in this test case. The reduction in the number of groups is unimpressive and the same number of azimuthal angles and track spacing are being used as in the normal 7 group. It is no surprise that SEAM, at best, barely beats the performance of the normal power iteration. Since SEAM is theoretically able to correct for any inconsistency in the transport sweep operator, it should be possible to incorporate both energy and angular acceleration into the same calculations. If this is the case, then the transport sweeps used in SEAM may be cheap enough relative to a most costly transport problem that the performance could match CMFD.

6.3.2 Simultaneous Energy-Angle Acceleration

In all previous SEAM calculations, equivalence with the respect to the transport operator was maintained when using different number of polar angles while sweeping across the same tracks.

In MOC, though, the accuracy of the transport sweep is dictated not solely by the number of polar angles but also by the number azimuthal angles and the track spacing. Since all of these quantities tie in directly with the transport operator, SEAM should be able to create an equivalent transport problem even when using different number of azimuthal angles and track spacings.

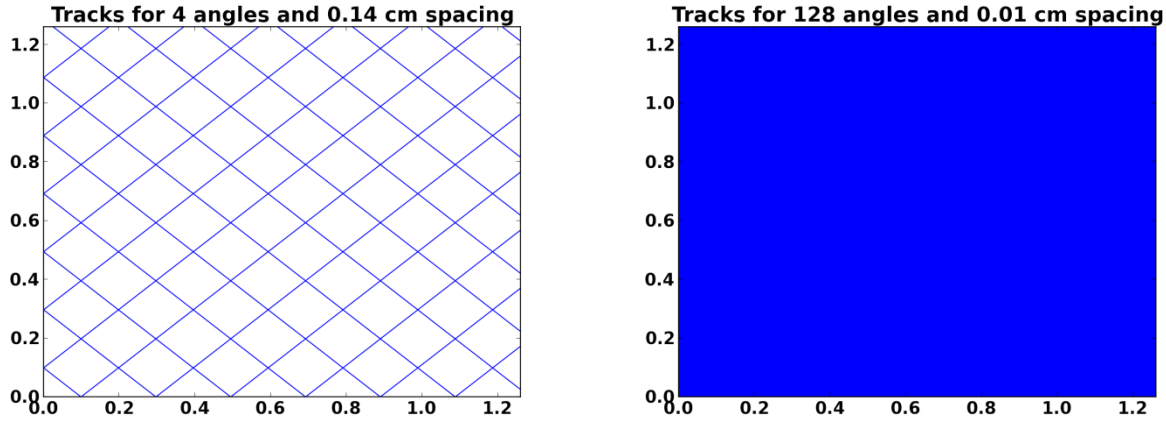

Figure 58 Equivalence between a lower order and high order angular discretization can be maintained using SEAM in addition to group collapse. In theory, a very coarse track distribution (left) can be used to accelerate a very fine track distribution (right) to accelerate high fidelity calculations.

For example, Figure 58 compares the tracks generated across the area of a single pin cell using 4 azimuthal angles and 0.14 cm track spacing to 128 azimuthal angles and 0.01 cm track spacing. **As long as the coarser tracks pass through every cell in the problem, then an equivalent transport problem can be formed using SEAM.** If this can be achieved, then the transport sweeps used in SEAM become significantly cheaper than the full transport sweep and improve its performance relative to CMFD.

Initially, this was tested on the C5G7 model with the full transport problem using 7 groups, 32 azimuthal angles, 0.02 cm track spacing, 3 polar angles and MOC. With SEAM, an equivalent problem was formed at each recondensation step using 4 groups, 4 azimuthal angles, 0.14 cm track spacing, 1 polar angle and step difference. The results using 1 and 2 power iterations per recondensation step were compared when using both the 4 and 32 azimuthal angle full transport sweeps

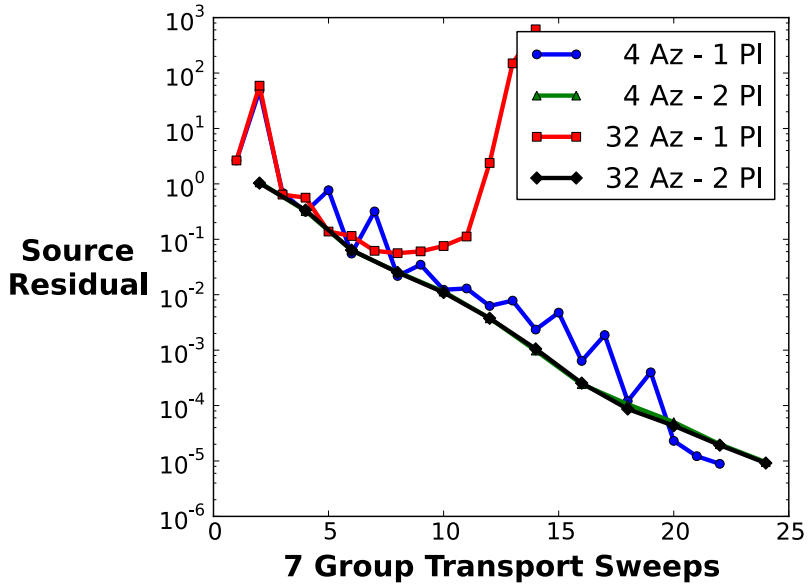

Figure 59 Convergence profiles when using 1 and 2 power iterations between recondensation steps when coupling between similar and dissimilar track layouts in MOC

While using only 1 power iteration was stable for the 4 azimuthal problem, the equivalence formed with the 32 azimuthal angle problem caused recondensation to diverge. At this time, the behavior is not well understood and further investigation is required to understand the mechanism behind the divergence in this case. According to Figure 59, using 2 power iterations per recondensation step stabilizes the 32 azimuthal angle problem and dampens out the oscillations without significantly increasing the number of full transport sweeps required to converge. This also shows that convergence using SEAM is largely independent of the angular discretization and track spacing of the full transport problem.

However, it still remains to be seen whether the cheaper transport problem SEAM forms is equivalent to the full transport problem when sufficiently converged. Therefore, the flat source region fission rates of the SEAM coarse group solution were compared to the full transport solution using power iteration. This comparison was made with the power iteration solution using both a 10^{-5} and 10^{-6} tolerance on the source residual.

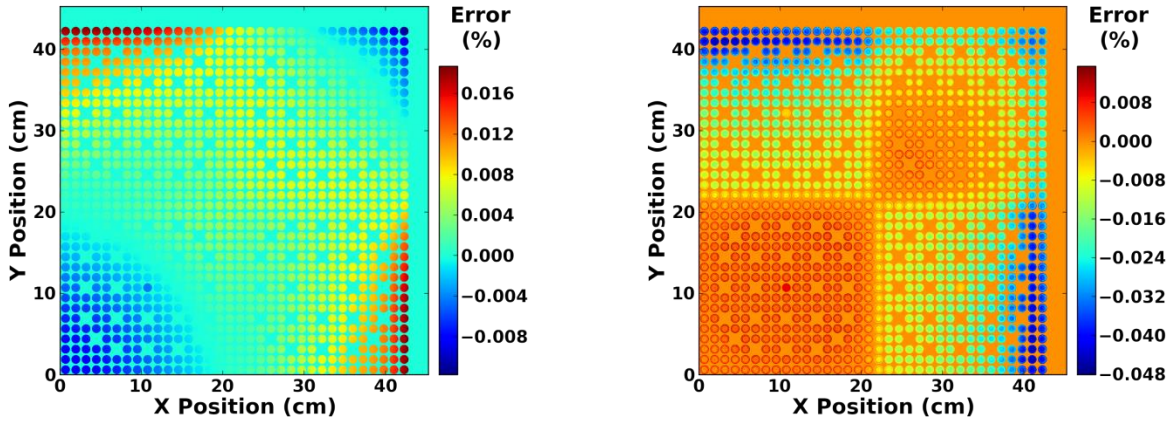

Figure 60 Fission errors of the converged solution using SEAM to accelerate a 32 az, 0.02 cm C5G7 model using an equivalent 4 az, 0.14 cm, 4 group transport problem at each recondensation step. The comparison is made with the full solution converged using 10^{-5} (left) and 10^{-6} (right) as the tolerance for the source residual.

The magnitude of the errors shown in Figure 60 match well with the errors observed in Figure 55. All the local fission rate errors relative to the 10^{-5} tolerance calculation remain well below .1%. SEAM's fission errors relative to the 10^{-6} tolerance also fall within the range of the errors between the 10^{-5} and 10^{-6} tolerance calculation plotted in Figure 61.

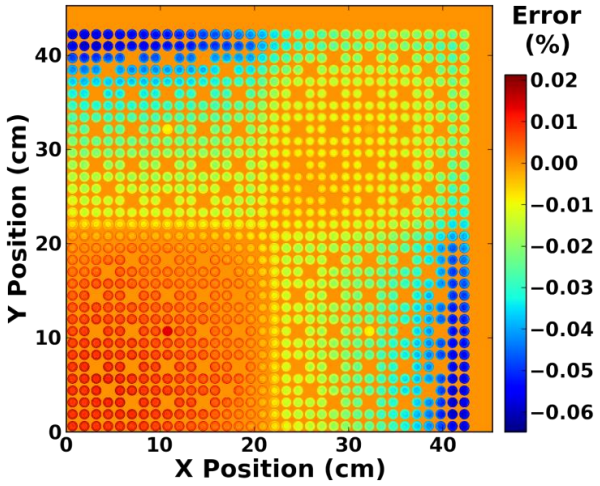

Figure 61 Errors in fission rates using normal power iteration when using a 10^{-5} tolerance as opposed to 10^{-6} for a 32 azimuthal angle, 0.02 cm track spacing calculation

The comparison of these errors proves that SEAM has effectively maintained equivalence using fewer azimuthal angles and wider track spacings than the full transport problem and doing so

without homogenization. While macro-track transport-based acceleration of MOC calculations was developed and used in CASMO-4E provide to a similar effect, these macro-tracks are formed using a subset of the finely spaced tracks already being used and therefore require use of the same angular discretization. **On the other hand, SEAM builds the low order problem independently of the high order problem, allowing a completely different angular discretization and arbitrary track spacing to be used in the low order transport problem.** This represents a significant step forward in showcasing SEAM's flexibility. [32]

The simultaneous group and energy acceleration of SEAM was then tested using different coarse group structures. The same coarse angular discretization and track spacing was used for 1, 2 and 4 group collapse of the C5G7 core. As discussed previously, 2 power iterations are used prior to each recondensation step to ensure stability and near monotonic convergence of SEAM. The results of SEAM are also compared to 7 group CMFD using pin cell homogenization.

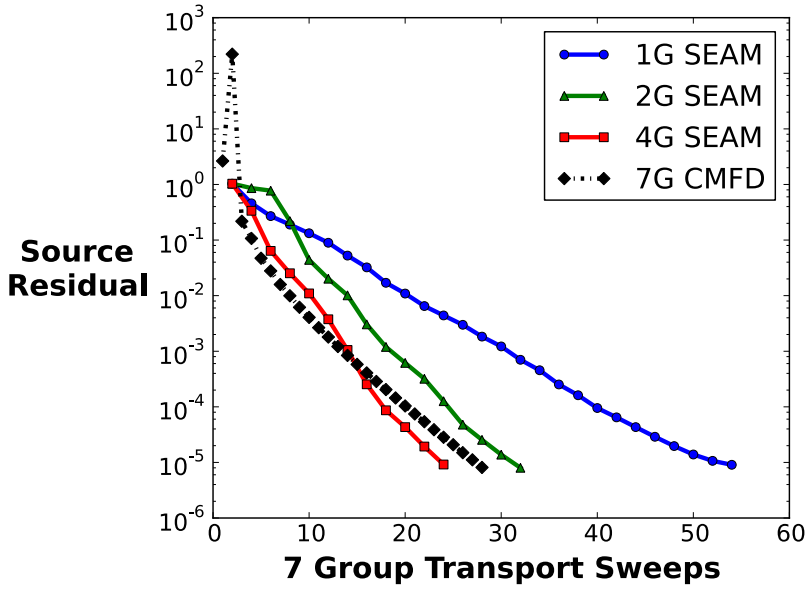

Figure 62 Comparison of convergence profiles for SEAM using 4 azimuthal angles, 0.14 cm track spacing and for 7 group CMFD using pin cell homogenization for a 32 azimuthal angle, 0.02 cm track spacing discretization of the C5G7 problem.

A subtle difference in the convergence profiles of the 1 group problem is observed in Figure 62. Whereas 1 group convergence in the 4 azimuthal angle C5G7 problem was at best oscillatory, even when using 2 power iterations per recondensation step, here the 1 group convergence is nearly monotonic from the beginning. This allows SEAM to converge after 54 full transport sweeps with 1 group collapse instead of using 76. SEAM using 2 and 4 group collapse retains similar convergence properties as the 4 azimuthal angle problem, as does 7 group CMFD.

While the number of full transport sweeps remained relatively the same, SEAM's overall performance improved significantly relative to the 4 azimuthal angle problem. This is to be expected since the full transport sweep becomes the most intensive part of the calculation. As long as SEAM's transport sweeps are much cheaper than the full sweeps, then performance will improve as the full sweep becomes more expensive.

Table 16 Comparison of performance between PI, SEAM and CMFD for the 32 az, 0.02 cm C5G7 problem

Method	FG Sweeps	FG Time	CG Sweeps	CG Time	Total Time
Normal PI	782	5,227.4	N/A		5,227.4
1 Group	54	342.9	2138	627.4	993.5
2 Group	32	206.1	1170	437.7	658.0
4 Group	24	152.2	982	439.7	603.4
CMFD	27	275.5	27*	55.1	330.6

* The number of diffusion eigenvalue solves used in CMFD

The data in Table 16 show that all the SEAM calculations, even for the 1 group problem, accelerate the normal power iteration calculation. SEAM's best performance is achieved in the 4 group calculation, reducing the total time to convergence by roughly a factor of 9. This is still more expensive than 7 group CMFD, which is able to converge the problem in roughly half the time of the 4 group SEAM calculation. The cheapest possible transport calculation still can't compete with the diffusion solves used in CMFD.

However, it is clear that additional overhead is accumulating in the full transport calculations since the time spent conducting the full transport sweep isn't scaling with the SEAM calculation. While SEAM averages roughly about 6.4 seconds per full transport sweep regardless of the number of groups, 7 group CMFD spends on average 10.2 seconds. This overhead comes from the added cost of tallying the partial currents at each of the pin cell boundaries during the full sweep. This overhead doesn't accumulate in SEAM since this new method only requires the scalar fluxes from the full sweep. While SEAM does require the additional fixed source sweep at each recondensation step, this is an extremely cheap calculation relative to the full transport sweep and only requires the scalar flux to be tallied. In total, the fixed source sweep along with all the group collapse functions only add 21.5 seconds to the 4 group calculation.

The equivalence possible with SEAM doesn't end here, though. This can be taken to an extreme degree by creating equivalence between a 4 azimuthal angle, 0.14 cm spacing, 1 polar angle step difference sweep and a 128 azimuthal angle, 0.01 spacing, 3 polar angle MOC sweep. From Figure 58, it is seen that the 128 azimuthal angle tracks provide extremely good coverage of a single cell while the 4 azimuthal angle tracks barely cover it. This equivalence was tested with

SEAM also adding in coarse group collapse to 1, 2 and 4 groups. SEAM's results are compared to that of 7 group CMFD using pin cell homogenization in Figure 63.

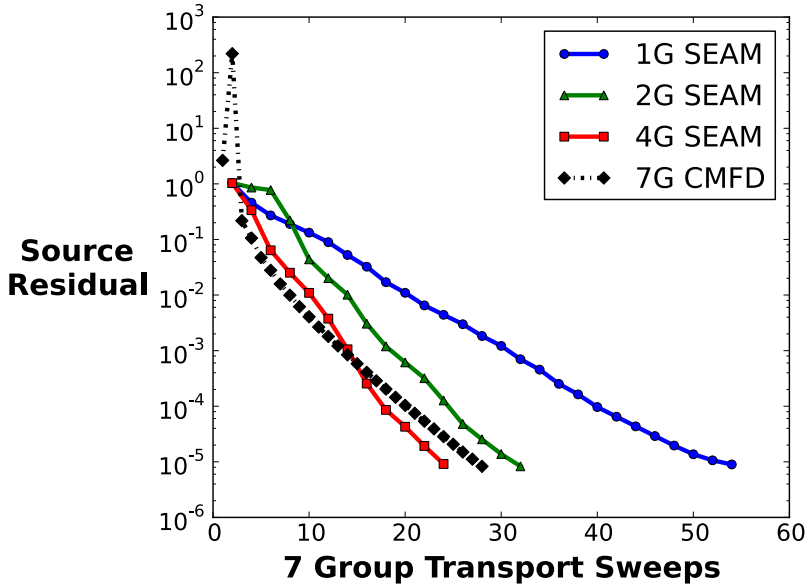

Figure 63 Comparison of convergence profiles for SEAM using 4 azimuthal angles, 0.14 cm track spacing and for 7 group CMFD using pin cell homogenization for a 128 azimuthal angle, 0.01 cm track spacing discretization of the C5G7 problem.

The convergence profiles produced using SEAM on the 128 azimuthal angle C5G7 problem are very similar to the 32 azimuthal angle equivalence. The 1, 2 and 4 group SEAM problems and the 7 group CMFD calculation all converge in using the same number of full transport sweeps as was needed for the 32 azimuthal angle problem. The performance of each, measured in Figure 16, show that while the convergence profiles remain the same, the time to convergence continues to improve for SEAM relative to both the normal power iteration and 7 Group CMFD.

Table 17 Comparison of performance between PI, SEAM and CMFD for the 128 az, 0.01 cm C5G7 problem

Method	FG Sweeps	FG Time	CG Sweeps	CG Time	Total Time
Normal PI	782	45,984.0			45,984.0
1 Group	54	2,254.2	2137	637.2	2917.8
2 Group	32	1342.2	1170	493.6	1850.6
4 Group	24	1011.4	982	479.7	1503.0
CMFD	27	1853.7	27*	55.4	1909.1

* The number of diffusion eigenvalue solves used in CMFD

The best performance is still achieved when using the 4 group collapse in conjunction with angular acceleration in SEAM. While 7 group CMFD still speeds up the calculation by a factor of 24, the best performance of SEAM is able to achieve a speed up of over 30 relative to the normal power iteration. While the diffusion eigenvalue problem solved in CMFD still requires nearly 10 times less time than the cheap transport eigenvalue problem used in SEAM, the overhead of the partial current tallies used in CMFD continues to scale proportionally with the number of azimuthal angles and inversely to the track spacing. Therefore, even though the diffusion problem beat SEAM's transport sweeps by 424 seconds, the additional tallies in CMFD accrue an extra 842 seconds. **And while CMFD performance could be improved through further optimization of the surface tallying function in OpenMOC, it should be noted that SEAM requires no new code to be added into the transport sweep and therefore is much less intrusive to pre-existing codes.**

It is apparent from these results that energy acceleration doesn't have a profound impact on the C5G7 benchmark. Since much of the difficult spectral information has been collapsed down into a few group problem, there isn't a significant amount of acceleration which can be achieved through energy acceleration by itself. Therefore, while this thesis has primarily been focused on energy acceleration, the flexibility and utility of SEAM have been shown to extend far beyond what it was originally intended for. Therefore, it is worth further exploring the possibility of purely angular acceleration for few group problems, meaning the group collapse procedure is not involved in the SEAM calculation.

6.3.3 Angular Acceleration for Few Group Problems

The original idea behind the development of SEAM was to be able to create a few group transport method that can equivalently represent the original fine group problem. Ideally, the bulk of the work would be placed on converging the few group problem and reduce the number of fine group transport sweeps. However, in problems already using few groups, there is little to gain from further reducing the number of groups using SEAM. Instead, the results from Section 6.3.2 point towards angular acceleration being the most beneficial for few group problems. Therefore, angular acceleration will be considered when using SEAM with 7 groups.

While the cross section can continue to be defined differently for each flat source region as was required in group collapse. The modification of the coarse group cross sections using the corrective factors defined in Section 5.1.2 added no additional storage since the coarse group cross sections already needed to be stored for each cell. When using the same number of groups, the corrective factors can instead be applied to the scalar flux after each "incorrect" sweep to produce the equivalent fine group source for the next calculation instead of storing new cross sections for each cell. In this case, only the factors need to be stored for each cell and fine group. Since the storage required for cross sections increases as G^2, this can significantly reduce storage requirements if only angular acceleration is desirable. However, if one is collapsing the fine group cross sections and storing coarse group cross sections for each cell then using the new cross section definitions in Equation (5.34) makes more sense.

This idea of purely angular acceleration was tested on the C5G7 using SEAM without group collapse but still creating an equivalent transport problem using 4 azimuthal angles, 0.14 cm track spacing, 1 polar angles and step difference. This equivalence was formed for the 4, 32 and 128 azimuthal angle C5G7 problems used in Sections 6.3.1 and 6.3.2. Figure 64 shows the results of the SEAM calculations and compares them with the results of 7 group CMFD acceleration using pin cell homogenization.

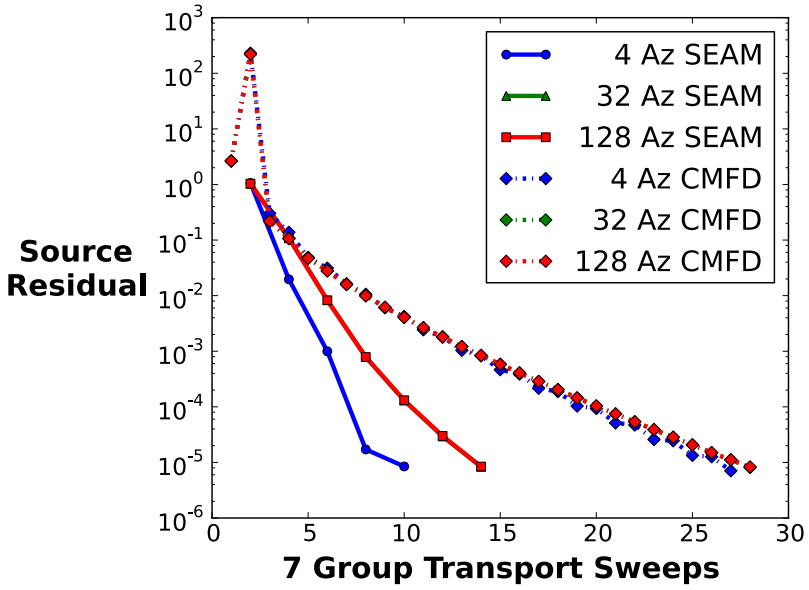

Figure 64 Comparison of convergence profiles for SEAM coupling 4 azimuthal angles, 0.14 cm track spacing to finer track spacings and for CMFD using pin cell homogenization, all using 7 groups. All the CMFD runs have the same convergence profile, as does the SEAM calculation for 32 and 128 azimuthal angles.

For C5G7, 7 group SEAM is able to reduce the number of full transport sweeps by almost half relative to 7 group CMFD. The 4 azimuthal angle problem only requires 10 full sweeps, while both the 32 and 128 azimuthal angle problem only require 14. The stellar performance in the 4 azimuthal angle problem is to be expected, however, since fewer approximations are being made. In this case, the same tracks are being used in the full and equivalent transport sweeps, the only difference is that the SEAM transport sweep uses 1 polar angle and step difference. The true performance measured in time shows that this is still just as time consuming as the normal power iteration. The run time does not significantly improve because the polar angles are included into the SIMD vectorization used in OpenMOC. This means that the sweep process across an individual segment can be carried out simultaneously across all the polar angles at once.

Table 18 Performance of Normal PI, 7 Group SEAM and 7 Group CMFD for the C5G7 Problem

Method	4 Az. Angles		32 Az. Angles		128 Az. Angles	
	Sweeps	**Time**	**Sweeps**	**Time**	**Sweeps**	**Time**
Normal PI	768	502.6	782	5,227.4	782	45,984.0
SEAM	10/756	440.1	14/845	670.3	14/845	1,110.5
CMFD	26	72.1	27	330.6	27	1,909.1

Another interesting behavior is observed in Table 18 for the 32 azimuthal angle problem. In this case, 4 group SEAM is able to outperform 7 group SEAM by over a minute. Even though 7 group SEAM significantly reduces the number of full transport sweeps, the 4 group problem is able to provide a cheaper means of converging the equivalent transport problem. Since the time required to converge the equivalent problem is still larger than time spent sweeping the full transport problem, 4 group SEAM still wins out.

The true benefit of angular acceleration with SEAM is seen in the results of the 128 azimuthal angle C5G7 analysis. For this problem, conducting even a few of the full transport sweeps becomes much more expensive than converging the equivalent transport problem. Since 7 group SEAM is able to reduce the number of full transport sweeps by half relative to the CMFD problem, SEAM substantially decreases the time required to converge the solution. For 7 group SEAM, the result is a speed up of over 45 relative to the normal power iteration and a speed up of 1.7 relative to CMFD.

The primary reason SEAM is able to reduce the number of full transport sweeps by so much relative to CMFD is because SEAM does not make use of spatial homogenization. While CMFD can provide very good acceleration of the global problem, the local problem within each pin cell can't be effectively accelerated once spatial homogenization takes place. **SEAM, on the other hand, performs no spatial homogenization and therefore can accelerate each flat source region individually.** In this case, there isn't even an approximation used with respect to energy, therefore the group fluxes for each flat source region can be accelerated individually too. This results in a substantial acceleration of the 128 azimuthal angle problem relative to CMFD.

6.4 Application of SEAM to the "C5G361" Core

While the C5G7 benchmark served as the ideal problem to initially test SEAM in 2D transport, this problem was not able to provide a good basis for determining the performance of SEAM in a truly fine group problem. Therefore, in order to establish SEAM's ability to accelerate a fine group transport calculation, a modified benchmark problem is proposed that incorporates the 361 group microscopic cross sections used in the 1D BWR analysis of Section 5.2 into the geometry of the original C5G7 benchmark. From here on, this problem will be referred to as the C5G361 problem.

6.4.1 Description of the "C5G361" Core

The C5G361 core retains the exact same spatial description of the original C5G7 problem except for a few minor differences. Since the compositions of the guide tubes and fission chambers could not be accurately ascertained from the definition of the C5G7 benchmark description, these materials were replaced with water to simplify the model. Also, due the presence of large resonances with this very detailed fine group data, the number of rings used to describe each fuel pin was increased from 3 to 7 in order to adequately resolve the sharp flux gradients at the resonances.

The exact compositions of the various Pu and MA isotopes were also not defined in the C5G7 description, so the following composition was used in order to be representative of LWR MOX fuel at the 4.3%, 7.0% and 8.7% enrichments.

Table 19 Plutonium and Minor Actinide Compositions Assumed in Macroscopic Cross Section Production for the MOX fuel in the C5G361 problem (Uranium not included here)

Pu Composition (At. %)		MA Composition (At. %)	
Pu-238	1.60	Np-237	51.1
Pu-239	60.2	Am-241	32.3
Pu-240	24.7	Am-243	16.6
Pu-241	8.60		
Pu-242	4.90		
Total Pu	89.4	Total MA	10.6

This 361 group 2D core is the first of its kind to be tested in OpenMOC and should prove difficult enough to test the limits of both SEAM and CMFD in their ability to accelerate fine group transport calculations.

The whole core version of this problem was first analyzed using 4 azimuthal angles, 0.14 cm track spacing and 3 polar angles. As stated before, the fuel rings were divided into 7 concentric, equal area rings and the same moderator discretization as used in the C5G7 core was used again in this case. However, the guide tubes and fission chambers were replaced with additional moderator using a fine square mesh cell. OpenMOC was used to converge the source residual to a tolerance of 10^{-5}. The converged value of k_{eff} was found to be 1.180338. The total run time to converge this problem ended up being 86,983 seconds, or just over 24 hours. While the angular discretization used for this problem was very poor, moving to a calculation using 32 angles and a 0.02 cm spacing would have required roughly 10 days to fully converge. However, the analysis from Section 6.3 indicated that the convergence profiles of both SEAM and CMFD were mostly independent of the angular discretization used in the full transport problem. Therefore, using 4 azimuthal angles should provide a good indication of performance using a more resolved angular discretization.

6.4.2 C5G361 Fine Group Transport Acceleration with SEAM and CMFD

For this 361 group 2D problem, two coarse group structures were tested using SEAM. The first is a 9 group structure that largely glosses over many of the prominent resonances except for the 6.67 eV resonance of U-238. The 23 group structure separates out the rest of the resolved resonances to determine if this would significantly improve the convergence of SEAM. The 9 group SEAM calculation was conducted twice, using 1 and 2 power iterations to determine if the behavior observed in Figure 59 was specific to the C5G7 problem or if it's inherent to SEAM. Each SEAM calculation used the same 4 azimuthal angles and 0.14 cm track spacing as the original problem, but continued to use step difference with 1 polar angle for the transport sweep.

These results were then compared to runs using CMFD acceleration. First, pin cell homogenization was used with two different relaxation factors, 0.45 and 0.6 to determine if the cheaper diffusion calculation could provide quick acceleration of the problem. For each relaxation factor, the same coarse group structure as the SEAM calculations were used to conduct 9 group and 23 group calculations with CMFD. The same coarse group structure and the finest possible mesh size were thought to provide a good basis by which to compare SEAM.

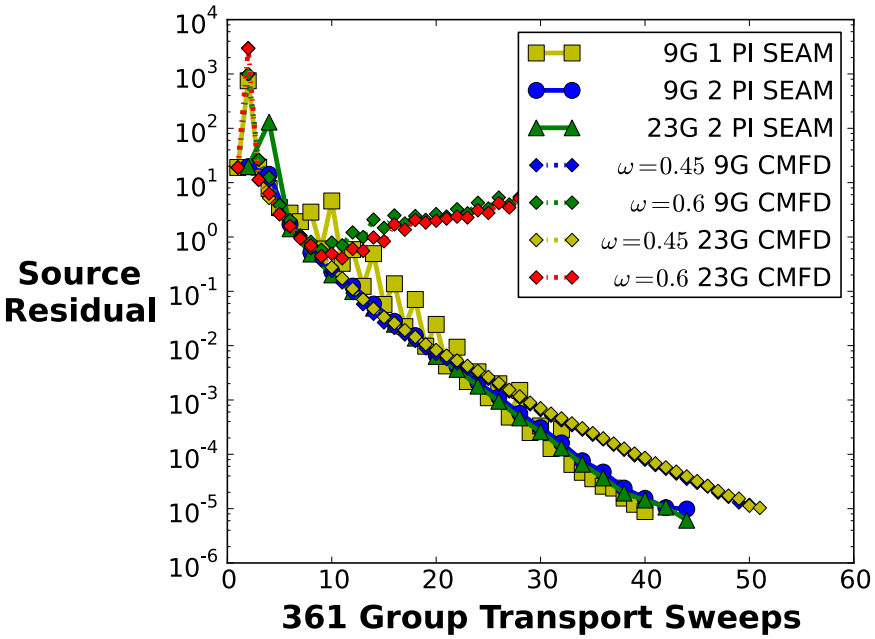

Figure 65 Comparison of convergence profiles for 9 and 23 group SEAM and CMFD (pin cell homogenization) acceleration of the C5G361 problem

Even though a very different problem is being considered in this case, oscillations in the convergence of SEAM persist when using 1 power iteration per recondensation step. Since this behavior seems to be problem independent, it is advisable to always conduct 2 power iterations per step to ensure stability and achieve monotonic convergence. However, further investigation should be conducted to determine the source of this instability.

Also observed in Figure 65 is that there is very little discernable difference in the convergence of the 9 group and 23 group SEAM runs when using 2 power iterations. Both require 44 fine group transport sweeps before the source is converged. This was somewhat surprising since the 9 group coarse group structure does not provide a detailed representation of the resonances in the 361 group total cross sections. In retrospect, this seems reasonable since any large sources initially present in the resonances at the beginning of the calculation will converge very quickly due to the very quick exponential decay of neutrons in the resonances. Since this will be picked up very quickly in the 361 group transport sweeps, then it may be important to take other spectral effects, such as upscatter, into consideration when choosing the coarse group structure. While

optimal choice of coarse groups is an important step in further optimizing the convergence of SEAM for fine group problems, this subject remains outside the scope of this thesis.

More important is to show that both the 9 and 23 group SEAM transport calculations remain equivalent to the full 361 group one. Therefore, the fission rate errors for each flat source region were calculated relative to the converged solution using the normal power iteration procedure in OpenMOC.

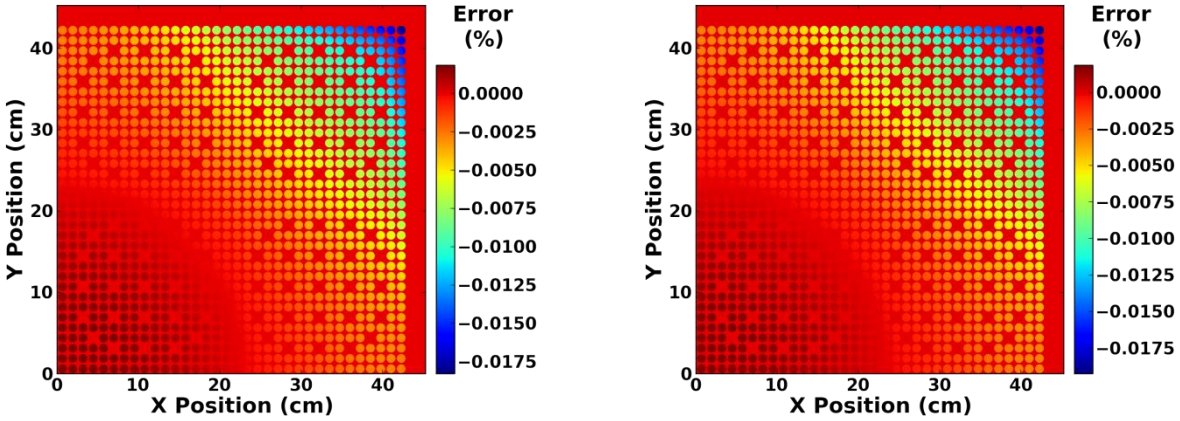

Figure 66 Comparison of fission errors in the 9 group (left) and 23 group (right) problems using SEAM

Figure 66 shows that all of these errors remain well below 0.1%. In fact, the majority the fission rate errors across over 90% of the core are below 0.01% for both the 9 group and 23 group SEAM runs using 2 power iterations per recondensation step. This proves that even for a 361 group problem, SEAM can still create an equivalent transport problem at each recondensation step even when using only 9 groups.

CMFD for both the 9 and 23 group calculations is able to converge quickly but not as much as SEAM. Both CMFD calculations using a relaxation factor of 0.45 were stable and required 50-51 fine group transport sweeps. So, in this case, SEAM was able to converge using 6-7 fewer fine group sweeps than CMFD. While using a larger relaxation factor was thought to further reduce fine group sweeps, the pin cell homogenized CMFD calculation became unstable when using the 0.6 relaxation factor.

Further inspection of Figure 65 show that initially the convergence of CMFD and SEAM match very closely to one another, but when their respective residuals drop below somewhere between 10^{-2}-10^{-3} then the convergence of CMFD begins to slow relative to SEAM. Since the coarse group structures, number of azimuthal angles and track spacings are all equivalent, the main difference must lie within the pin cell homogenization used in CMFD. This makes sense because of the spatial approximations used in CMFD.

While CMFD, even in its pin cell homogenized form, is able to produce a much better global solution after each iteration, eventually it has to rely on the fine group transport sweeps to sufficiently converge the spatial dependence of the scalar fluxes within each pin cell. This is likely what's occurring when CMFD deviates from SEAM. SEAM, on the other hand, doesn't apply any spatial homogenization and therefore is able to accelerate the scalar fluxes for each and every flat source region, providing a level of detail to the equivalent problem that CMFD can't achieve when applying spatial homogenization.

Table 20 Comparison of performance between PI, SEAM and CMFD for the 4 az, 0.14 cm C5G361 problem

Method	FG Sweeps	FG Time	CG Sweeps	CG Time	Total Time
Normal PI	2,101	86,983.0	N/A	N/A	86,983.0
9G SEAM	44	1,752.8	2,290	2,383.2	4,457.8
23G SEAM	44	1,790.8	2,428	4,729.7	6,883.9
9G CMFD	50	5,098.6 (3,059)**	50*	3,673.2	8,771.9
23G CMFD	50	4,926.3 (2,963)**	50*	14,788.7	19,715.0

* The number of diffusion eigenvalue solves used in CMFD

** Normalization accounting for overhead

In terms of performance, the 9 group SEAM calculation definitely provides the most efficient approach for creating an equivalent coarse group transport calculation. In this case, SEAM is able to speed up the normal fine group power iteration calculation by a factor of 19.5. On the other hand, CMFD is only able to speed it up by less than a factor of 10. This is primarily due to how expensive the solving the diffusion problem becomes when doing the full fine group CMFD calculation. The substantial increase in the diffusion solve time moving from 9 groups to 23 groups also highlights the need for optimization for many groups in OpenMOC CMFD solver.

Further angular refinement of this problem will also lead to additional overhead in the fine group sweeps for CMFD. All of this taken into consideration, it appears that SEAM is better suited to meet the challenges of fine group transport acceleration.

6.5 Summary

SEAM has now been successfully implemented in OpenMOC and tested on two different 2D problems, the C5G7 benchmark core and a newly created problem called the C5G361 core.

Results in the C5G7 testing show that SEAM can create a substantially cheaper transport problem that is equivalent to the full transport problem. One of the best examples of this was demonstrating that a 1 group calculation using 4 azimuthal angles, a 0.14 cm track spacing, 1 polar angle with step difference could be made equivalent to the full 7 group transport calculation using 128 azimuthal angles, a 0.01 cm track spacing, 3 polar angles with MOC. This test highlights SEAM's flexibility in forming a cheap but equivalent transport problem.

Initial comparison with CMFD showed that although 4 group SEAM calculations could match the number of full transport calculations used in CMFD, the relatively cheap transport calculations still required much more effort than conducting the incredibly cheap diffusion calculation used in CMFD. However, when moving to large numbers of azimuthal angles like the 128 azimuthal angle example, 4 group SEAM is able to outperform CMFD because of the additional overhead required to tally the surface partial currents for each of the homogenized pin cells to create the equivalent diffusion problem. While the equivalent transport problem SEAM forms is more expensive to solve, only a single fixed source sweep of the low order problem is required to form equivalence. Therefore, the same overhead is required even though high order problem becomes more expensive. Since the net current tally overhead is proportional to the angular order and track spacing used, CMFD eventually reaches the point where the overhead is much greater than the gains in the diffusion calculations. This point was reached in the 128 azimuthal angle, 0.01 cm calculation, at which point the 4 group SEAM calculation was able to converge in about 1500 seconds as opposed to CMFD's time of 1900 seconds.

SEAM was also able to accelerate the C5G7 problem without using group collapse. The result was a purely angular acceleration approach that further improved the 128 azimuthal angle results. The 7 group SEAM calculation was able to cut the number of full transport sweeps nearly in half relative to the CMFD calculation. In doing so, 7 group SEAM was able to converge in only 1100 seconds, resulting in a final speed up of 45 relative to the normal power iteration calculation.

Unfortunately, testing with the C5G7 benchmark core was unable to provide an adequate testing base for a true fine group transport problem, so a new 361 group problem was formed to truly test the performance of energy acceleration using SEAM. This "C5G361" problem is based almost exactly off of the geometry of the C5G7 core, but instead uses the 361 group microscopic cross section data used to test SEAM in the 1D BWR calculations of Section 5.2. This new problem was tested with 9 and 23 group CMFD acceleration pin cell homogenization and compared with the 9 and 23 group calculation using SEAM. Both 9 and 23 group SEAM calculations were able to further converge the solution using only 44 fine group transport sweeps as compared to the 50-51 required to converge using the most detailed CMFD calculation. This result suggests that convergence below a source residual of 10^{-2} is dictated more by changes in the flux inside the pin cell which can't be captured using a homogenized pin cell approximation. Since SEAM doesn't use spatial homogenization, the scalar fluxes for each flat source region can be accelerated and leads to better performance.

SEAM has now been shown to accelerate both fine group and high fidelity transport calculations, in many cases even improve upon CMFD acceleration. While further study is necessary, it is important to recognize that the recondensation process used in SEAM is very similar to that used in DGM. Even though DGM proved ultimately too memory intensive, it may still hold valuable intuition into how to further improve SEAM in the future. Therefore, recondensation using DGM will be looked at both as a nonlinear method and as a form of energy preconditioning.

7 Nonlinear DGM and Energy Preconditioning

Although DGM does not appear to be feasible in direct use for accelerating fine group transport calculations, it may still provide useful insights into how to apply nonlinear methods to recondensation. One only needs to look at the definitions of our cross section moments in DGM to understand this. Any changes in the flux moments will lead to a nonlinear change in the cross section because the reaction rate is divided by the coarse group flux, the leading order flux moment.

$$\Sigma_{i,g} + \delta\Sigma_{i,g} = \frac{\sum_j \Sigma_{g,ij}\left(\phi_{j,g} + \delta\phi_{j,g}\right)}{\phi_{0,g} + \delta\phi_{0,g}} \tag{7.1}$$

The recondensation process which has been applied thus far is essentially a fixed point iteration which solves this nonlinear problem. Throughout this process it is assumed that the cross section moments do not change as the coarse group eigenvalue is conducted and the fixed source DGM solved. It is only after the fine group fluxes have been reconstructed, which is assume to be closer to the true solution, that the cross sections are actually updated. Therefore, instead of using this fixed point iteration concept, there should be a way to apply a nonlinear solver to recondensation and solve for the flux moments, eigenvalue and cross section moments simultaneously. Newton's method will be applied for this purpose.

Newton's method is one of the most widely used methods for solving nonlinear problems. In order to apply this method, the definition of the problem is changed in such a way that a function or set of operations, *F*, performed on the solution *x* equals 0. The operator *F* is defined as a set of residual nonlinear equations acting on all of the unknowns in the vector *x*.

$$\vec{F}(\vec{x}) = 0 \tag{7.2}$$

In deriving Newton's method, a Taylor Expansion is performed on the residual equations centered about the initial guess or current iterate of the solution. All quadratic or higher terms are dropped thereby forming an approximation that linearizes the nonlinear problem. Index *k*

will denote the current iterate of *x* within the Newton method, while *i* and *j* will be used to distinguish between the residual equations in F and the unknowns in *x* respectively.

$$\vec{F}(\vec{x}^k + \delta\vec{x}) \approx \vec{F}(\vec{x}^k) + J(\vec{x}^k)\delta\vec{x} \tag{7.3}$$

J is used to describe the first order partial derivatives of the residual equations with respect to all the unknowns in *x*. This matrix is referred to as the Jacobian and its elements are defined for the i^{th} residual equation and the j^{th} unknown as follows.

$$J_{ij}(\vec{x}^k) = \frac{\partial F_i(\vec{x})}{\partial x_j}\Bigg]_{\vec{x}=\vec{x}^k} \tag{7.4}$$

Using this approximation, a change in our current solution can be found that will drive the linearized residual to zero. Assuming that this is a good approximation to our problem, this will drive the current iterate closer to the true solution. This step amounts to setting the linearized form of *F* equal to zero and solving a linear system of equations for the search direction δx.

$$\begin{aligned} \delta\vec{x}^k &= -J(\vec{x}^k)^{-1}\vec{F}(\vec{x}^k) \\ \vec{x}^{k+1} &= \vec{x}^k + \delta\vec{x}^k \end{aligned} \tag{7.5}$$

Then the current solution is updated with the evaluated search direction and the Jacobian is reevaluated at the new iterate for use in the next iteration. If this linear approximation holds, then the solution will converge. One of the very nice properties of Newton's method is that it achieves quadratic convergence. For example, if the current iterate is roughly accurate to one digit, then the first Newton step will provide accuracy to two digits, the second provides 4 digit accuracy and so on.

Unfortunately, this behavior is achieved only when the linear approximation used in the Taylor expansion is good. Since nonlinear problems are rarely well-behaving function, the initial guess could also provide a Jacobian that is singular and therefore not invertible. There are many potential fixes for this, such as using a continuation scheme. This approach adds a linear problem on top of the current nonlinear problem using a weighting that, at first, heavily favors the linear problem. After each newton step, the weight is reduced such that the modified problem better resembles the original nonlinear problem. While choosing how much to relax

this weight is tricky, this allows the method to avoid many of the convergence issues associated with Newton's method. For the infinite homogeneous problems being considered, however, it is assumed that they are well-behaved nonlinear functions and therefore it is unlikely using a continuation scheme will be necessary.

While Newton's method provides very good convergence for nonlinear problems, it does require the solution of a linear system of equations at each Newton step. For large problems, providing an exact inverse of the Jacobian is computationally prohibitive and, for the most part, unnecessary. If the solution to the linear system is found within machine precision, but only provides one or two orders of magnitude improvement in the residual norm of our Newton method, then a significant amount of time has been wasted in obtaining unnecessary accuracy for the search direction. On the other hand, if the search direction is too inaccurate, then the quadratic convergence of the Newton method can deteriorate to linear convergence. Therefore, inexact Newton can be used instead to reduce the burden of the linear solves at each Newton step.

For determining the approximate solution, any number of Krylov methods can be applied to solve the linear system of equations. For our purposes, since the nature of the Jacobian for the infinite homogeneous problem is largely material dependent, the Generalized Minimum RESidual (GMRES) method will be applied for the linear solve.

7.1 Nonlinear Derivation with Explicit Recondensation

In order to apply recondensation to Newton's method, the residual equations must be defined in order to build the Jacobian at each Newton step. The infinite homogeneous form of the Discrete Generalized Multigroup equations serves as the starting point for this derivation. All terms on the right hand side are moved to the left to form the residual equation for DGM. λ is used instead of $1/k_{eff}$ for simplicity in the definition of the nonlinear problem.

$$R_{i,g}^{DGM} = \Sigma_{T,g}\phi_{i,g} - \sum_{g'} \Sigma_{s,ig\leftarrow g'}\phi_{0,g'} - \lambda\chi_{i,g}\sum_{g'} \nu\Sigma_{f,0,g'}\phi_{0,g'} + \delta_{i,g}\phi_{0,g} \tag{7.6}$$

This provides a different equation for each fine group, but each has considerably more unknowns to solve for. Since the point is to incorporate the cross section moments into the Newton method, all of these must be included in the total list of unknowns. For this infinite homogeneous problem, the total number of unknowns associated with these cross sections is *FG*x*CG*+*FG*+2x*CG*.

To incorporate the definitions of the cross section moments and the del terms into the nonlinear solve, the original definitions of these material properties are put into residual form as was done for the DGM equations.

The definition for the coarse group total cross section is used to derive a residual equation that will be used to evaluate the Jacobian.

$$R_g^{\Sigma_T} = \Sigma_{T,g}\phi_{0,g} - \sum_{j\in g} \Sigma_{T,g,0j}\phi_{j,g} \tag{7.7}$$

The same process is applied to the del term, scattering cross section moments and the coarse group fission cross section. The equations for chi are not included in the nonlinear recondensation process because the group collapse of chi is independent of the fine group fluxes and only depends on the fine group values for chi.

$$R_{i,g}^{\delta} = \delta_{i,g}\phi_{0,g} + \Sigma_{T,g}\phi_{i,g} - \sum_{j\in g} \Sigma_{T,g,ij}\cdot\phi_{j,g}$$

$$R_{i,g,g'}^{\Sigma_s} = \Sigma_{s,ig\leftarrow g'}\phi_{0,g'} - \sum_{j\in g} \Sigma_{s,g\leftarrow g',ij}\cdot\phi_{j,g} \tag{7.8}$$

$$R_g^{\nu\Sigma_f} = \nu\Sigma_{f,0,g}\phi_{0,g} - \sum_{j\in g} \nu\Sigma_{f,g,0j}\cdot\phi_{j,g}$$

Although the correct number of residual equations for each of the cross section definitions has been added, there is still the eigenvalue to account for. An additional equation must be defined so that the problem becomes solvable. A normalization residual is typically used for this purpose. Due to the nature of the eigenvalue problem, changing the magnitude of a mode solving the eigenvalue problem will not affect the validity of the solution. Therefore a constraint can be placed on the flux limiting the magnitude of the flux moment vector. This constraint only needs to be placed on the flux moments since any changes in the overall magnitude of the flux moment vector will not change the cross section values.

To ensure that the method is equivalent to the nonlinear fine group equations, the normalization residual used in the fine group calculation is used.

$$R(\phi) = \frac{1}{2}\left(\sum_{g}\sum_{K\in g}\phi_g(K)^2 - 1\right) \tag{7.9}$$

The definition of the flux moments can be used to substitute the flux moments in place of the fine group fluxes. Since the flux moments are defined through the application of the discrete basis function and these basis functions satisfy an orthonormality condition, the fine group fluxes can be expanded in terms of the flux moments, $\check{\phi}$, which are the fine group fluxes transformed with the discrete basis function, *P*, and normalization, *A*.

$$R(\check{\phi}) = \frac{1}{2}\left(\check{\phi}^T A^T P^T P A \check{\phi} - 1\right) \tag{7.10}$$

This substitution defines a new normalization residual for the flux moments. If the orthonormality condition still holds true for the discrete basis functions, then the following relation must hold as well.

$$AP^TP = I \tag{7.11}$$

Since *A* is a diagonal matrix, there is no difference between *A* and A^T. In the end, this means that the product of the four matrices in Equation (7.10) reduces down to A. Thus the normalization equation reduces to Equation (7.12).

$$R(\check{\phi}) = \frac{1}{2}\left(\sum_{g}\sum_{i\in g} A_{i,g}\phi_{i,g}^{2} - 1\right) \tag{7.12}$$

This ensures not only that the converged solution will match the true fine group solution, but this also ensures that their norms will match perfectly.

In order to implement the Newton method, the Jacobian must be formed using the residual equations. While employing a Jacobian-Free Newton-Krylov method would be the most clear choice, since this is a first of its kind treatment of the infinite homogeneous problem, defining and analyzing the structure of the Jacobian will be important in determining the practicality of this method moving forward.

The first part of the Jacobian is formed by taking the partial derivatives of the DGM Residual Equation with respect to the flux moments, eigenvalue, coarse group total cross section, del term, scattering moment cross sections and fission cross sections in that order. For clarity, upper case deltas in the following definitions represent Kroenecker deltas if no comma is present. All lower case deltas represent the del term used in DGM.

$$\frac{\partial R_{i,g}^{DGM}}{\partial \phi_{k,h}} = \begin{cases} \left(\Sigma_{T,g} + \delta_{i,g}\right)\Delta_{hg} - \sum_{g'} \Sigma_{s,ig\leftarrow g'}\Delta_{hg'} - \lambda\chi_{i,g}\sum_{g'} \nu\Sigma_{f,0,g'}\Delta_{hg'} & k = 0 \\ \Sigma_{T,g}\Delta_{hg} & k > 0 \end{cases} \tag{7.13}$$

It is noted that as long as the number of coarse group is much less than the fine group problem it represents, this region of the Jacobian will be very sparse. This is because only the columns corresponding to the coarse group fluxes will be completely filled. All other columns have a single value along the diagonal.

The derivative with respect to the eigenvalue produces only a single column vector relating directly to the fission term.

$$\frac{\partial R_{i,g}^{DGM}}{\partial \lambda} = -\chi_{i,g}\sum_{g'} \nu\Sigma_{f,0,g'}\phi_{0,g'} \tag{7.14}$$

The partial derivative with respect to the coarse group total cross section produces a diagonal matrix whose diagonal is filled with the flux moments.

$$\frac{\partial R_{i,g}^{DGM}}{\partial \Sigma_{T,h}} = \phi_{i,g}\Delta_{hg} \tag{7.15}$$

The partial derivative with respect to the del term results in a diagonal matrix filled only with the coarse group fluxes.

$$\frac{\partial R_{i,g}^{DGM}}{\partial \delta_{k,h}} = \phi_{0,g}\Delta_{hg}\Delta_{ki} \tag{7.16}$$

$$\frac{\partial R_{i,g}^{DGM}}{\partial \Sigma_{s,kh\leftarrow h'}} = -\phi_{0,h'}\Delta_{hg}\Delta_{ki} \tag{7.17}$$

$$\frac{\partial R_{i,g}^{DGM}}{\partial \nu\Sigma_{f,0,h}} = -\lambda\, \chi_{i,g}\phi_{0,h} \tag{7.18}$$

Next, the coarse group total cross section residual is incorporated into the Jacobian by taking its partial derivative with respect to the flux moments and the coarse group cross sections. The first set of equations is defined by a row of diagonal matrices and the second is a single diagonal matrix.

$$\frac{R_g^{\Sigma_T}}{\partial \phi_{k,h}} = \begin{cases} \left(\Sigma_{T,g} - \vec{\Sigma}_{T,g,00}\right)\Delta_{hg} & k = 0 \\ -\Sigma_{T,g,0k}\Delta_{hg} & k > 0 \end{cases} \tag{7.19}$$

$$\frac{R_g^{\Sigma_T}}{\partial \Sigma_{T,h}} = \phi_{0,g}\Delta_{hg} \tag{7.20}$$

The partial derivative of the del tem with respect to the flux moments forms a block diagonal matrix which is blocked by coarse group. With respect to the coarse group total cross section, the derivative forms a column of diagonal matrices. Taking the derivative with respect to the del term produces an additional diagonal matrix.

$$\frac{R_{i,g}^{\delta}}{\partial \phi_{k,h}} = \begin{cases} \left(\delta_{i,g} + \Sigma_{T,g} - \vec{\Sigma}_{T,g,i0}\right)\Delta_{hg} & k = 0 \\ \left(\Sigma_{T,g} - \Sigma_{T,g,ik}\right)\Delta_{hg} & k > 0 \end{cases} \tag{7.21}$$

$$\frac{R_{i,g}^{\delta}}{\partial \Sigma_{T,h}} = \phi_{0,g}\Delta_{hg} \tag{7.22}$$

$$\frac{R_{i,g}^{\delta}}{\partial \delta_{k,h}} = \phi_{0,g}\Delta_{hg}\Delta_{ki} \tag{7.23}$$

The scattering term residual is now factored into the Jacobian.

$$R_{i,g,g\prime}(\Sigma_s) = \Sigma_{s,ig\leftarrow g\prime}\phi_{0,g\prime} - \sum_{j\in g}\Sigma_{s,g\leftarrow g\prime,ij}\cdot\phi_{j,g\prime} \tag{7.24}$$

The partial derivative of this residual is taken with respect to the flux moments and the cross section moments to determine the linear response due to changes in these parameters.

$$\frac{R_{i,g,g\prime}^{\Sigma_s}}{\partial \phi_{k,h}} = \begin{cases} \left(\Sigma_{s,ig\leftarrow g\prime} - \Sigma_{s,g\leftarrow g\prime,i0}\right)\Delta_{hg\prime} & k = 0 \\ -\Sigma_{s,g\leftarrow g\prime,ik}\delta_{hg\prime} & k > 0 \end{cases} \tag{7.25}$$

$$\frac{R_{i,g,g\prime}^{\Sigma_s}}{\partial \Sigma_{s,kh\leftarrow h\prime}} = \phi_{0,g\prime}\Delta_{h\prime g\prime} \tag{7.26}$$

The contribution of the fission cross section term to the Jacobian is also derived. The partial derivative with respect to the flux moments and fission cross sections are taken.

$$\frac{R_g^{\nu\Sigma_f}}{\partial\phi_{k,h}} = \begin{cases} \left(\nu\Sigma_{f,0,g} - \nu\Sigma_{f,g,00}\right)\Delta_{hg'} & k = 0 \\ -\nu\Sigma_{f,g,0k}\Delta_{hg'} & k > 0 \end{cases} \tag{7.27}$$

$$\frac{R_g^{\nu\Sigma_f}}{\partial\nu\Sigma_{f,0,h}} = \phi_{0,g}\Delta_{hg} \tag{7.28}$$

Lastly, the contribution of the normalization term to the Jacobian is found by taking the derivative with respect to the flux moments.

$$\frac{R(\check{\phi})}{\partial\phi_{k,h}} = A_{i,g}\phi_{i,g}\Delta_{hg}\Delta_{ki} \tag{7.29}$$

This leads to a Jacobian that is much larger in total size than that for the original fine group problem, but is sparser and actually has fewer nonzero terms. The only dense region in the Jacobian is due to the partial derivative of the scattering cross section with respect to the flux moments. Most of the other partial derivatives return a diagonal or mostly diagonal contribution to the Jacobian. The sparsity pattern of the Jacobian mapping a 295 fine group problem onto an 8 coarse group problem can be seen in Figure 67.

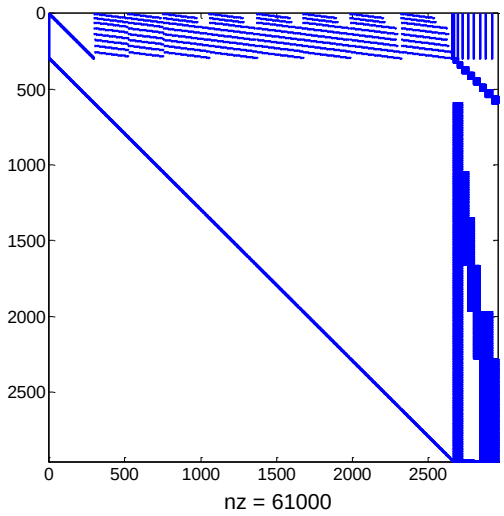

Figure 67 Sparsity pattern for the Jacobian built for a 295 fine group problem using explicit nonlinear recondensation

This Jacobian is very sparse due to the additional residual equations introduced to relate the flux moments to the cross section moments. A one coarse group calculation roughly doubles the number of nonzeros relative to the 8 group case and leads to a much more dense Jacobian. However, in total, the number of nonzero Jacobian elements for the original fine group problem is roughly 87,000, therefore introducing this approach for a one coarse group representation actually increases the number of nonzeros. There seems to be a minimum in the number of total nonzeros as one moves to splitting the fine group problem into more and more coarse groups. This certainly is interesting and could translate into improved performance if a Jacobian-Free method were applied to this problem. For now, though, application of this approach will continue to be focused on using Newton-Krylov methods.

7.2 Results for Infinite Homogeneous Problems

For our infinite homogeneous problem, the inexact Newton solve uses a convergence criterion that solves the linear system of equation to near machine precision and therefore is an exact Newton method. Therefore, if the linear approximation holds, quadratic convergence should be observed. The tolerance for the residual of the Newton method is set to 10^{-12}. The problem analyzed here is for a 295 group infinite homogeneous material representative of fuel in an HTR.

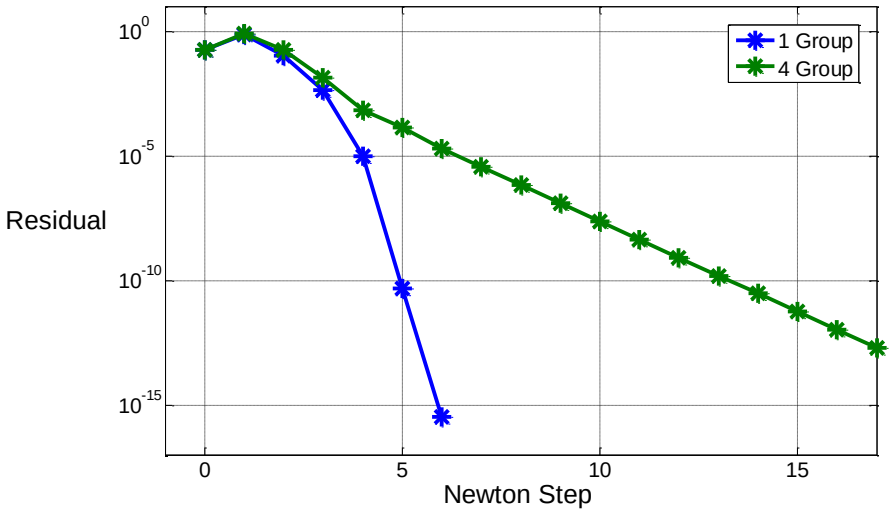

Figure 68 Convergence of the residual using explicit nonlinear recondensation and the del term

In Figure 68, the convergence rate is shown to depend strongly on the number of coarse groups representing the fine group problem. While quadratic convergence is observed in the one group problem, the 4 group problem only attains linear convergence. It seems that somewhere in the explicit nonlinear recondensation equations, convergence is forced to be linear instead of the optimal quadratic convergence Newton's method is known for. Not shown in this figure is the convergence plot for the 8 group problem, which is unstable and does not converge after 20 Newton iterations. Therefore, this approach to nonlinear recondensation does not appear to be viable.

One thought was that the del term, which was a necessity in spatial calculations, is not really required in an infinite homogeneous problem and could be causing the Newton method to lag behind its optimal convergence properties. Therefore, the del term definition was combined with the definition of the total coarse group cross section to produce a coarse group total cross section moment in hopes of improving the convergence properties.

$$\Sigma_{T,g}\phi_{i,g} + \delta_{i,g}\phi_{0,g} \rightarrow \Sigma_{T,i,g}\phi_{0,g} \tag{7.30}$$

As a consequence of combining these two terms, a new residual equation is defined to incorporate this new definition into the Jacobian.

$$R_{i,g}(\Sigma_T) = \Sigma_{T,i,g}\phi_{0,g} - \sum_{j\in g} \Sigma_{T,g,ij} \cdot \phi_{j,g} \tag{7.31}$$

Unfortunately, this doesn't seem to be the true source of the problem. In Figure 69 below, the exact same convergence properties are observed as in the previous approach. Here, the convergence behavior for the 8 group case is also shown to highlight its divergence.

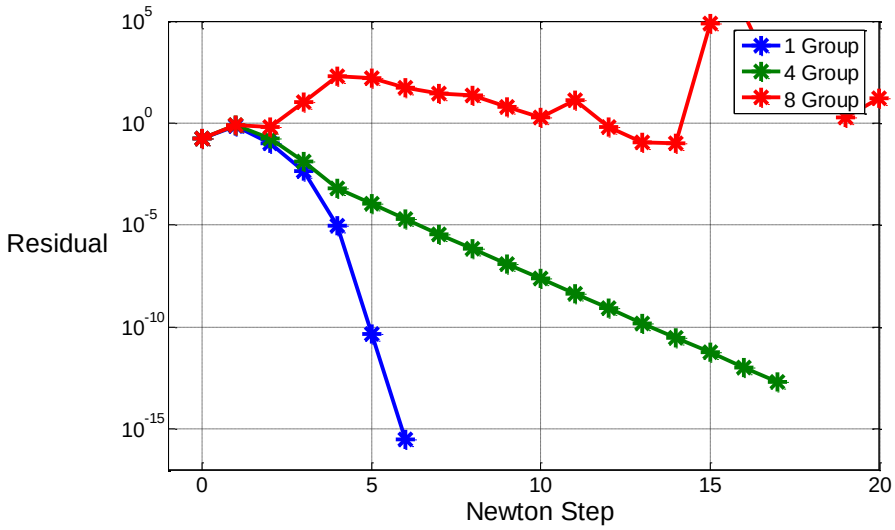

Figure 69 Convergence of residual for nonlinear explicit recondensation without using a del term

7.3 Nonlinear Implicit Recondensation

For explicit nonlinear recondensation, all flux moments and cross sections were required to be searched during the Krylov solve for each Newton step. This means that using a Jacobian Free Newton Krylov method would require evaluating all 2x*FG*+2x*CG*+*FG*x*CG*+1 number of residual equations for each Krylov iterations within a newton step. If applied to a spatially dependent problem, then this number would need to be further multiplied by the total number of spatial unknowns. This is considerably more number crunching and storage than would be required if the fine group equations were used to begin with.

The primary issue with this approach is that it seems that explicitly defining the recondensation process in its true nonlinear form does not lead to any gains in performance and actually can lead to instability and divergence. Therefore, in order to make the method practical regarding performance and memory requirements, the moment cross sections definitions are folded back into the DGM moment residual equations to reduce the number of searched variables to the number of flux moments.

$$
\begin{aligned}
R_{i,g}(DGM) = \Sigma_{T,g}\phi_{i,g} &- \sum_{g'} \frac{\sum_{j\in g} \Sigma_{s,g\leftarrow g',ij} \cdot \phi_{j,g'}}{\phi_{0,g'}} \phi_{0,g'} \\
&- \lambda\chi_{i,g} \sum_{g'} \frac{\sum_{j\in g'} \nu\Sigma_{f,g',0j} \cdot \phi_{j,g'}}{\phi_{0,g'}} \phi_{0,g'} \\
&+ \left(\frac{\sum_{j\in g} \Sigma_{T,g,ij} \cdot \phi_{j,g} - \Sigma_{T,g}\phi_{i,g}}{\phi_{0,g}} \right) \phi_{0,g}
\end{aligned}
\tag{7.32}
$$

Once these residual definitions are substituted, many cancellations occur which lead to the definition of Equation (7.33).

$$
\begin{aligned}
R_{i,g}(DGM) = &\sum_{j\in g} \Sigma_{T,g,ij} \cdot \phi_{j,g} - \sum_{g'} \sum_{j\in g} \Sigma_{s,g\leftarrow g',ij} \cdot \phi_{j,g'} \\
&- \lambda\chi_{i,g} \sum_{g'} \sum_{j\in g} \nu\Sigma_{f,g,0j} \cdot \phi_{j,g}
\end{aligned}
\tag{7.33}
$$

This new set of equations is essentially the original fine group residual equations transformed using the discrete basis functions. If this is placed in matrix-vector form, the similarity is more striking. First the vector of flux moments is organized by stacking the moments from each group in sequential order. This organization dictates the order in which the rest of the mass matrices must be ordered as well.

For the DGM scattering matrix, it is defined in block form by coarse group.

$$
\overline{\overline{\Sigma_s}} = \begin{pmatrix} \overline{\overline{\Sigma}}_{s,1\leftarrow 1} & \cdots & \overline{\overline{\Sigma}}_{s,1\leftarrow G} \\ \vdots & \ddots & \vdots \\ \overline{\overline{\Sigma}}_{s,G\leftarrow 1} & \cdots & \overline{\overline{\Sigma}}_{s,G\leftarrow G} \end{pmatrix}
\tag{7.34}
$$

Each block is defined by the action of discrete basis functions on all the in-going and outgoing fine group scattering information when scattering from one coarse group to another.

$$
\Sigma_{s,g\leftarrow g',i\leftarrow j} = \sum_{K\in g} \sum_{L\in g'} P_i(K)\Sigma_s(K\leftarrow L)P_j(L)a_j
\tag{7.35}
$$

The total cross section information, originally a diagonal matrix, becomes block diagonal.

$$\overline{\overline{\Sigma_T}} = \begin{pmatrix} \overline{\overline{\Sigma}}_{T,1} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \overline{\overline{\Sigma}}_{T,G} \end{pmatrix} \tag{7.36}$$

Each block is a dense matrix which looks like the discrete equivalent of mass matrices used in finite element methods.

$$\Sigma_{T,g,ij} = \sum_{K \in g} \sum_{L \in g'} P_i(K) \Sigma_T(L) P_j(L) a_j \tag{7.37}$$

The fine group fission cross sections and fission spectrum are transformed by the discrete basis functions to produce a new set of vectors which are grouped by coarse group.

$$\overrightarrow{\nu\Sigma_f} = \begin{pmatrix} \overrightarrow{\nu\Sigma_{f,1}} \\ \vdots \\ \overrightarrow{\nu\Sigma_{f,G}} \end{pmatrix}, \ \vec{\chi} = \begin{pmatrix} \overrightarrow{\chi_1} \\ \vdots \\ \overrightarrow{\chi_G} \end{pmatrix} \tag{7.38}$$

Each vector group is comprised of the fission and spectrum moments, respectively, each of which are produced by applying the discrete basis functions to the fine group data.

$$\begin{aligned} \nu\Sigma_{f,g,i} &= \sum_{K \in g} a_i P_i(K) \nu\Sigma_{f,g}(K) \\ \chi_{g,i} &= \sum_{K \in g} P_i(K) \chi_g(K) \end{aligned} \tag{7.39}$$

The complete flux moment vector is similarly defined for this set of DGM residual equations and is organized by coarse group.

$$\vec{\phi} = \begin{pmatrix} \overrightarrow{\phi_1} \\ \vdots \\ \overrightarrow{\phi_G} \end{pmatrix} \tag{7.40}$$

Each group of column vectors contains all of the flux moments for that coarse group. These are also defined through application of the discrete basis functions.

$$\phi_{g,i} = \sum_{K \in g} P_i(K)\phi_g(K) \tag{7.41}$$

When all the components of the DGM residual equations are defined in this manner, the equations simplify down to the following form.

$$\vec{R}(DGM) = \left(\overline{\overline{\Sigma_T}} - \overline{\overline{\Sigma_s}} - \lambda\left(\vec{\chi} \otimes \overrightarrow{\nu\Sigma_f}\right)\right)\vec{\phi} \tag{7.42}$$

The only unknowns in this implicit formulation of the recondensation process are the flux moments and eigenvalue. All the cross section moments have been implicitly included into the mass matrix formalism. If a matrix were created where the discrete basis functions are placed in block diagonal fashion for each coarse group, then the summations are replaced by a set of matrix-matrix and matrix-vector multiplications.

$$\overline{\overline{P}} = \begin{pmatrix} \overline{\overline{P}}_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \overline{\overline{P}}_G \end{pmatrix} \tag{7.43}$$

Each block's columnspace is defined by the discrete basis functions that are being applied to each distinct coarse group. These blocks are dense, square and are of the same size as the number of fine groups within that particular coarse group.

$$\overline{\overline{P}}_g = \begin{pmatrix} P_0(0) & P_1(0) & \cdots & P_K(0) \\ P_0(1) & P_1(1) & & P_K(1) \\ & \vdots & \ddots & \vdots \\ P_0(K) & P_1(K) & \cdots & P_K(K) \end{pmatrix} \tag{7.44}$$

The normalization constants must also be organized by group and then moment to match the ordering of the discrete basis functions in *P*. Matrix *A* is block diagonal by coarse group and each block is a diagonal matrix comprised of the normalization constants for each moment within that coarse group.

$$\bar{\bar{A}} = \begin{pmatrix} \bar{\bar{A}}_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \bar{\bar{A}}_G \end{pmatrix} \tag{7.45}$$

With the discrete basis functions and residual equations organized in this manner, the summation definitions previously used can be defined instead by the following matrix notation.

$$\begin{aligned} \overline{\overline{\Sigma_T}} &= \bar{\bar{P}}^T \, diag(\vec{\Sigma}_{T,FG}) \bar{\bar{P}} \bar{\bar{A}} \\ \overline{\overline{\Sigma_s}} &= \bar{\bar{P}}^T \, \bar{\bar{\Sigma}}_{s,FG} \bar{\bar{P}} \bar{\bar{A}} \\ \overrightarrow{\nu\Sigma_f} &= \overrightarrow{\nu\Sigma_{f,FG}} \bar{\bar{P}} \bar{\bar{A}} \\ \vec{\chi} &= \overrightarrow{\chi_{FG}} \bar{\bar{P}} \end{aligned} \tag{7.46}$$

Substituting these definitions into our residual equation to produce Equation (7.47).

$$\vec{R}(DGM) = \bar{\bar{P}}^T \left(diag(\vec{\Sigma}_{T,FG}) - \bar{\bar{\Sigma}}_{s,FG} - \lambda \left(\overrightarrow{\chi_{FG}}^T \overrightarrow{\nu\Sigma_{f,FG}} \right) \right) \bar{\bar{P}} \bar{\bar{A}} \bar{\bar{P}}^T \overrightarrow{\phi_{FG}} \tag{7.47}$$

Now the nonlinear implicit recondensation method is nothing more than the original nonlinear fine group problem projected using the discrete basis functions. Since this new mapping is linear, the search direction at each Newton step is transformed in the same linear fashion and will result in the same convergence properties as the original fine group problem. Therefore, there is no reason to expect that this new formulation would cause the Newton method to converge faster. While this may be the case, one can see that this DGM formulation is essentially a left and right preconditioning of the original fine group problem using the discrete basis functions. This means that there is the possibility of acceleration in determining the new search direction if a method such as GMRES were used. This idea will be later explored in Section 7.4.2, but for now the focus will be solely on the Newton method without worrying about the solver used in the linear problem.

7.4 Results for Infinite Homogeneous Problems

For initial tests of this implicit nonlinear recondensation method, the same 295 group data set applied to the explicit case was used. The initial guess for the fine group problem is a flat spectrum and the discrete basis function is used to produce the corresponding initial guess for the DGM problem to maintain consistency.

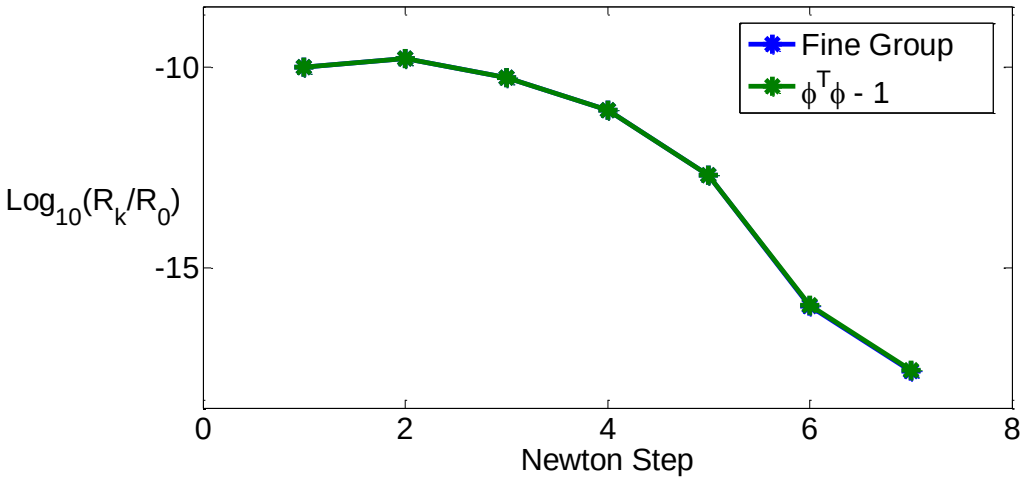

Figure 70 Convergence of the Newton method when applied to the fine group problem and implicit recondensation. The two convergence plots match exactly, so the flue line is directly beneath the green.

As can be seen, the two methods produce exactly the same convergence profiles. As was mentioned before, this is to be expected since the Newton method uses a linearization of the nonlinear problem. Therefore, projecting the original problem using a linear mapping will also lead to mapping the search direction into the new space as well.

This seems fairly counterproductive since the whole reason behind developing a nonlinear DGM method is to achieve near fine group accuracy with much less work than a fine group solve. It is concluded, therefore, that this linear mapping using the discrete basis functions doesn't provide any acceleration to the Newton method.

7.4.1 Relaxing the Residual Equations

While this projection of the fine group equations does not improve convergence, it does provide different ways to relax some of the residual equations placed on the original fine group problem if implicit recondensation is interpreted in terms of the original recondensation process. By doing so, some of the constraints can be relaxed knowing that coarse group values and reaction rates must match those of the fine group exactly.

For the original recondensation procedure, the coarse group problem was separated from the moment solve because, while the higher moments provided the spectral shape functions to reconstruct fine group fluxes, the coarse group flux is still what preserves the reaction rates. Now normalization residual will be considered in this context. If the coarse group flux is what actually maintains neutron balance, then the normalization residual need only to be placed on the coarse group flux. This should allow the problem to converge to the correct solution since the moments will automatically adjust with any changes in the coarse group flux.

In fact, considerable acceleration in our Newton method is observed just by using the fact that the higher moments in the original recondensation process depended only on the coarse group solution. There is one slight issue that is worth mentioning, though, when changing the normalization residual to depend solely on the coarse group fluxes and that is consistency. A direct sum of the fine groups to produce the coarse group solution represents an L_1 norm for each coarse group, but when these coarse groups are applied to the normalization, which represents an L_2 norm, the result is a slight difference in the initial guess. Therefore, the summed coarse groups are renormalized such that they satisfy the normalization residual at the initial guess.

Next, fission contributions to the residual can be trimmed using the same logic. From the coarse group equations in the original recondensation process, the outgoing fission spectrum is assumed to be completely independent of the neutron spectrum coming into the fission reaction. Under this assumption, only the magnitude, not the shape, of the fission spectrum can be affected by the scalar fluxes through the total fission rate. In this coarse group formulation, though, the total fission rate is only tied to the coarse group flux and not the higher flux moments within the

coarse group. Therefore, the higher order flux moments' contributions to the fission part of the residual can be removed, leaving only the contributions from the coarse group fluxes.

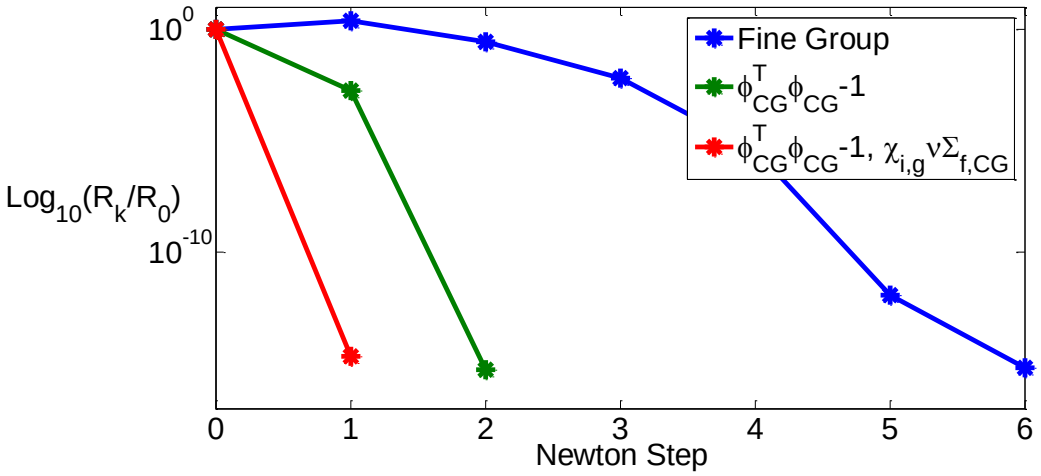

Figure 71 Comparison of Newton method convergence when normalization and fission residuals are simplified.

For this problem, trimming the fission contribution to the residual leads to even more improvement in the Newton method. Since the method converges in one iteration, it can be said that these modifications to the residual have essentially changed this nonlinear eigenvalue problem in such a way that it becomes linear or at least very close to linear. This is the only logical way to explain how it is that the Newton method converged in only one step

7.4.2 Discrete Cosine Transforms as a Preconditioner

In addition to accelerating the Newton method through changes in the residual, the method used to solve for the correct search direction at each Newton step is also analyzed. Since application of the discrete basis functions are, in a sense, preconditioning the fine group problem, their impact on the convergence of GMRES in solving for the search direction will be determined. All three types of basis functions were tested to determine if any performed better than the others. In order to apply Discrete Legendre Polynomials to the 295 group problem, an 8 coarse group structure was used to minimize accrual of round off error in forming this particular basis.

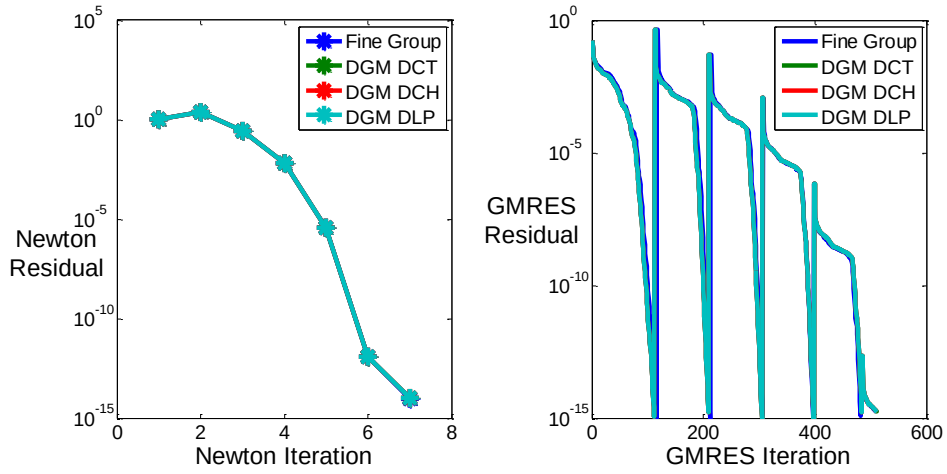

Figure 72 Comparison of convergence for the Newton method and GMRES when using various discrete basis functions.

As seen in Figure 72, preconditioning using any of the discrete basis functions is largely ineffective. This is to be expected because the transformations applied contain no information pertaining to the nature of the problem they are trying to precondition.

The eigenvalue distribution of the problem also provides a good indicator as to the effectiveness of the preconditioner. This is extremely important when preconditioning the problem for Krylov methods like GMRES. These methods work basically by producing an approximate function that interpolates the eigenvalues of the linear system. If many of eigenvalues are clumped close together, then the Krylov method will produce a good approximation for the solution very quickly. If they are spaced out, then more GMRES iterations will be required to converge to an approximate solution. Therefore, comparing the eigenvalue distribution of the original fine group problem to the preconditioned problem using the discrete basis functions should provide a good indicator at how effective they are at preconditioning the system.

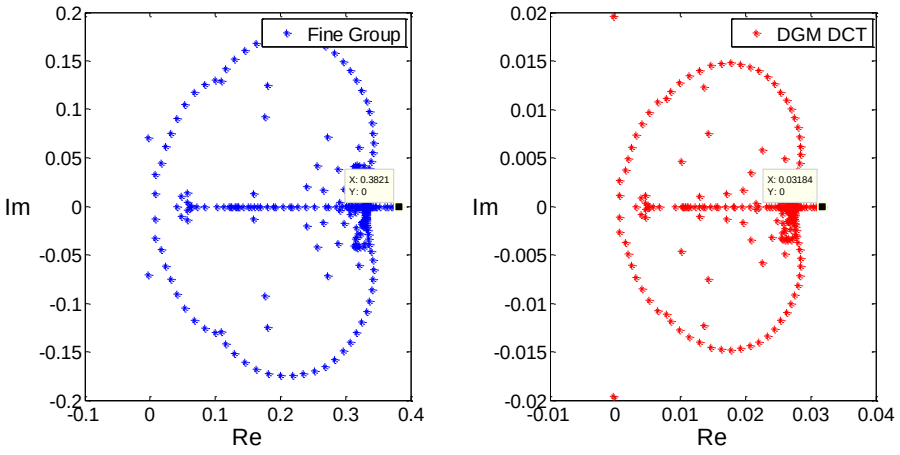

Figure 73 Plot of the eigenvalues for the infinite homogeneous problem before and after applying discrete basis functions

The biggest discernible difference in Figure 73 is the scaling of the eigenvalues that takes place when the discrete basis functions are applied to the original problem. The overall shape and distribution of the eigenvalues distribution has not changed significantly, therefore, there is no observed improvement in the convergence of GMRES at each newton step.

Looking forward in applying this methodology to spatial problems, accelerating the Newton method will only go so far in reducing the total number of transport sweeps. In fact, for infinite homogeneous and spatial problems, it is possible that relaxing the constraints previously placed on the basis functions could reduce the total number of sweeps more than the Newton acceleration could. Therefore, instead of applying discrete basis functions, generic left and right preconditioners will be applied to the fine group problem in the following fashion. Note the fine group residual matrix from the fine group problem has been replaced with the short hand notation M(λ).

$$\vec{R}(\phi,\lambda) = P_L M(\lambda) P_R P_R^{-1} \phi \tag{7.48}$$

Preconditioning is focused on this part of the Residual and Jacobian because it constitutes the bulk of the physics. Preconditioning this quantity will incorporate energy information, as well as spatial, if solving a diffusion problem. Applying the preconditioners only to this subset of the

problem, though, will lead to the formation of a different nonlinear problem and will become inconsistent with the original fine group problem. Therefore, these matrices must be applied with the full Jacobian in a linear fashion to maintain consistency while still achieving the desired effect. First the original linear problem solved at each newton step.

$$\begin{bmatrix} M(\lambda) & -F\phi \\ \phi^T & 0 \end{bmatrix} \begin{bmatrix} \delta\phi \\ \delta\lambda \end{bmatrix} = \begin{bmatrix} M(\lambda)\phi \\ (1/2)(\phi^T\phi - 1) \end{bmatrix} \tag{7.49}$$

In this notation, *F* is used to represent the "scattering" due to fission term, which is the outer product of chi and the fission cross sections. Then the preconditioner can be drawn out of the identity matrix if the orthonormality condition placed on the preconditioners is satisfied. Equation (7.50) shows the end result of this transformation.

$$\begin{bmatrix} P_L M(\lambda) P_R & -P_L F\phi \\ \phi^T P_R & 0 \end{bmatrix} \begin{bmatrix} P_R^{-1}\delta\phi \\ \delta\lambda \end{bmatrix} = \begin{bmatrix} P_L M(\lambda)\phi \\ (1/2)(\phi^T\phi - 1) \end{bmatrix} \tag{7.50}$$

Another set of transformations must be applied such that all the necessary unknowns are right preconditioned to produce a set of flux moments. To do this, the orthonormality condition can be substituted in place of the identity matrices, allowing the preconditioners to be applied directly to the fine group fluxes. Upon further simplification, Equation (7.51) is obtained.

$$\begin{bmatrix} \widetilde{M}(\lambda) & -\tilde{F}\tilde{\phi} \\ \tilde{\phi}^T P_R^T P_R & 0 \end{bmatrix} \begin{bmatrix} \delta\tilde{\phi} \\ \delta\lambda \end{bmatrix} = \begin{bmatrix} \widetilde{M}(\lambda)\tilde{\phi} \\ (1/2)\left(\tilde{\phi}^T P_R^T P_R \tilde{\phi} - 1\right) \end{bmatrix} \tag{7.51}$$

This new form of the linearized problem is now consistent with the original fine group problem as long as the left preconditioner is unitary. This assures that no scaling of the residual occurs. This is important because, while a unitary transformation will maintain the correct search direction, any additional scaling will cause the magnitude of the search vector to expand or contract and produce a different starting point for the next newton step. Therefore a unitary left preconditioner should minimize any scaling of the residuals and therefore reduce the number of GMRES steps for each linear solve without affecting the newton method.

7.4.3 Forming a Set of Natural Basis Functions

The best preconditioners that can be used are those which, in addition to satisfying the unitary behavior noted above, diagonalize or approximately diagonalize $M(\lambda)$. If this is done exactly, then the eigenvalues of the Jacobian will be grouped together very closely and lead to fewer GMRES iterations. Obviously the ideal preconditioner would be the exact inverse of $\mathrm{M}(\lambda)$, but if enough work has been put in to produce the inverse at every newton step then problem is nearly solved already. What can be done instead is calculate the inverse of M at a fixed value for lambda, say 1, knowing beforehand that realistic reactors will have an eigenvalue near 1. Therefore, inverting *M* at this value should give us a very good approximation of the inverse at every Newton step.

Another quality of a good preconditioner is that it is easily invertible. There's little point to preconditioning the problem if calculating the preconditioner's inverse is just as time consuming and computationally intensive as inverting the original problem. Therefore, it would be ideal for the preconditioner to be sparse or, as discussed before, unitary since the inverse of a unitary matrix is its transpose. For the purposes of preconditioning an infinite homogeneous problem, since the problem from the beginning is not sparse, the focus will be placed on preconditioners that have this unitary property.

The first possible method which satisfies all these properties is eigendecomposition. This method decomposes the original matrix into two matrices, one which is a dense square matrix and the other diagonal. The diagonal matrix, Λ, contains all of the eigenvalues of the original matrix while the full matrix, *Q*, contains their respective eigenvectors. The original matrix, *M*, is decomposed using Λ and *Q* in the following manner.

$$M = Q\Lambda Q^{-1} \tag{7.52}$$

The matrix of eigenvectors can be unitary if the original matrix is real and symmetric. While the matrices considered are always real, they are never symmetric. This asymmetry can't be avoided since it is a direct result of the cross section data. Therefore, eigendecomposition will likely lead

to complex numbers defining the columnspace of the preconditioner and they won't be unitary. Therefore, inverting Q would also be just as costly as inverting the initial problem. Therefore this method was not considered further for producing preconditioners.

Singular Value Decomposition (SVD), on the other hand, takes in a matrix M and produces a single diagonal matrix, S, and two different unitary matrices, U and V. The values in the diagonal matrix are called singular values. SVD can be applied to non-symmetric matrices and still yield U and V matrices which are orthonormal to their respective transposes, providing two different preconditioners that are easily invertible. The matrix M can then be decomposed using the different basis functions produced via SVD.

$$M = USV^T \tag{7.53}$$

If the matrix M is real, then SVD is closely related to eigendecomposition according to Equation (7.54).

$$\begin{aligned} M^T M &= VSU^T USV^T = VSSV^T \\ MM^T &= USV^T VSU^T = USSU^T \end{aligned} \tag{7.54}$$

In this case, SVD produces eigenvectors for M^TM and MM^T and the common eigenvalues the two products share. Therefore, V is a matrix composed of the eigenvectors of M^TM, U is a matrix composed of the eigenvectors of MM^T and the singular values are the square roots of the eigenvalues.

One of the truly interesting aspects of using SVD on this type of problem rests in the columnspace of U and V. Since these matrices are related to the eigenvectors of M^TM and MM^T, then these eigenvectors should have some relation to the solution of these modified problems. So if φ is the solution to the fine group problem, $M\varphi=0$, then it is also the solution to $M^TM\varphi=0$ as well. This means that the columnspace of V is composed of eigenvectors which have some relation to the approximate fine group solution. This can be observed by comparing the columnspace of V to fine group solution, as is done in Figure 74.

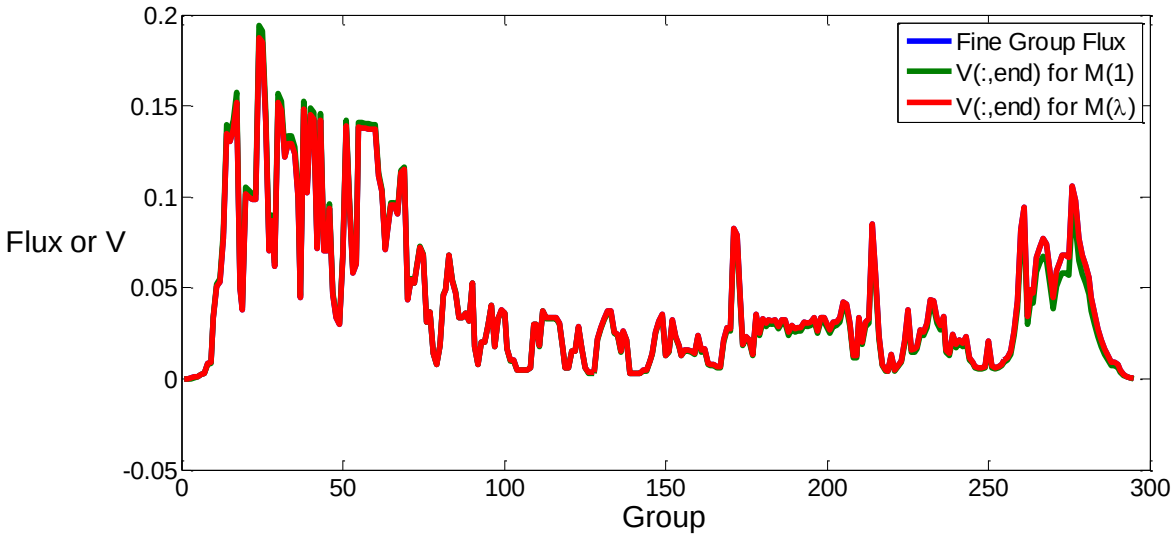

Figure 74 Comparison of the columnspace of V associated with the smallest singular value for different assumed eigenvalues

In Figure 74, the fine group flux is compared to the column of matrix *V* corresponding to the lowest singular value in S. When lambda is equal to the exact value converged upon in the Newton iteration, this column of *V* is exactly the same as the true fine group solution. This means that there are *N*-1 additional columns that contain a natural set of basis functions which are all orthonormal to the exact fine group solution. This, of course, the eigenvalue is known beforehand, which means that the solution is already full known at this point. In order to apply this technique to the infinite homogeneous problem without conducting an eigensolve first, lambda is assumed to be 1 when SVD is applied. In this case the SVD only needs to be computed for the approximate problem once and then applied accordingly at each Newton step. This approximation provides us with an approximate fine group spectrum for the last column of *V*. The remaining column space of *V* then provides basis functions which can be used to correct the approximate spectrum.

Since *V* will be used to perform the right preconditioning of our problem, seeing how V^T acts on φ provides some insight into what *V* is really doing to our infinite homogeneous problem.

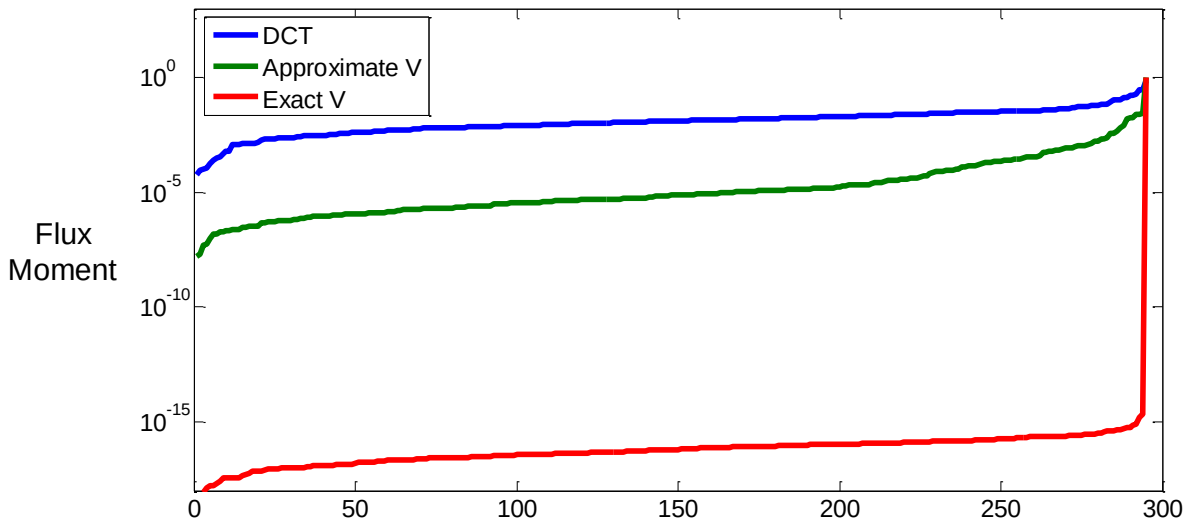

Figure 75 Flux moment values sorted from highest to lowest values after applying DCT, approximate V and exact V as the discrete basis functions used in DGM. A faster drop in the magnitude of the flux moment moving right to left means fewer moments are required to represent an approximate solution.

The compactness of the flux moment representation is compared when using DCTs, the approximate *V* matrix and the exact *V* matrix as discrete basis functions. This was done in Figure 75 by sorting the flux moments from their largest value to lowest. As expected the best representation occurs when the exact *V* matrix is used as the discrete basis function. In this case, the last column of *V* is the exact fine group spectrum and therefore provides the sole source of information necessary for the problem. All other moments, since they are orthonormal to the fine group solution, are within machine precision of 0. The application of DCT's performs poorly because little to no physics has been incorporated in generating these basis functions. The first moment in DCTs assumes a flat spectrum and the rest of the higher moments are used after to correct for this poor approximation. Application of *V* from the approximate problem, however, provides a significant improvement upon DCTs without needing to fully solve the eigenvalue problem beforehand. Because the last column contains an approximate fine group spectrum, this moment within *V* is able to incorporate the bulk of the spectral information while the rest of the higher moments provide small corrections to the approximation.

The same logic can be applied to the use of U as our left preconditioner. Since it is known that the adjoint fine group flux solves the following problem, $M^T\varphi^\dagger=0$, then this should also lead to $MM^T\ \varphi^\dagger=0$. This means that the eigenvectors in U must relate somehow to the solution of the adjoint problem. A comparison of the calculated adjoint flux to the columnspace of U is provided in Figure 76.

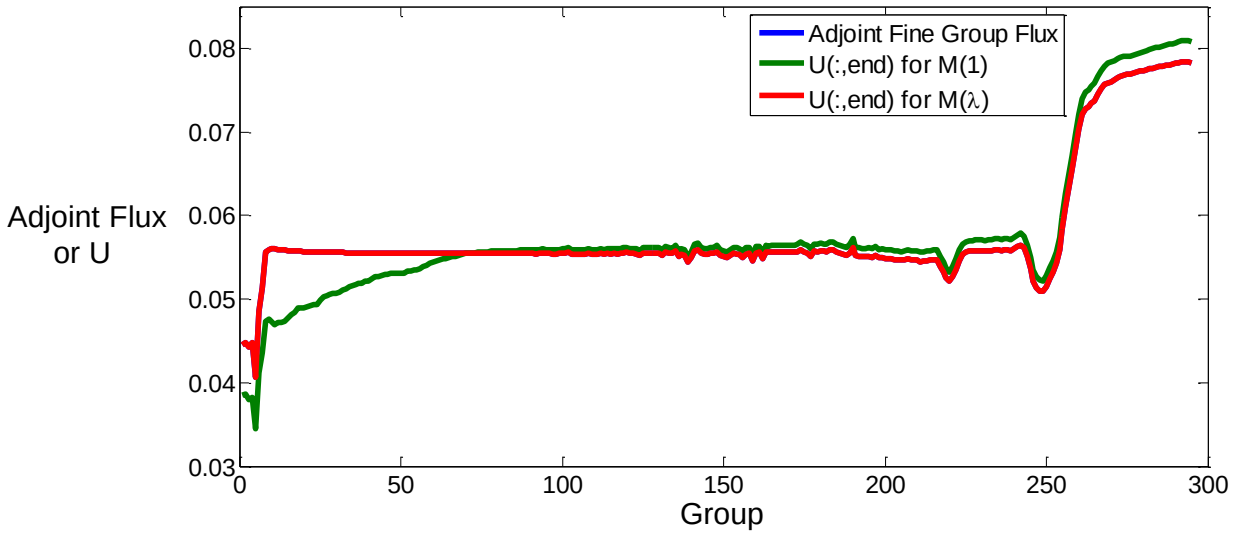

Figure 76 Adjoint flux moment values sorted from highest to lowest values after applying DCT, approximate U and exact U as the discrete basis functions used in DGM. The blue line lies beneath the red line since use of the exact eigenvalue produces the exact fine group adjoint flux in *U*.

When U is produced from the SVD of the residual problem assuming the exact eigenvalue the last column of U matches, the exact fine group adjoint flux is calculated to within machine precision. What this tells us is that application of U^T as a left preconditioner provides a transformation of the residual using the adjoint flux solution. If the adjoint flux is interpreted as an indicator of neutron importance, then this application of U^T as a left preconditioner weights the residual for each fine group according to their relative importance to the problem. However, since the correct eigenvalue is not known beforehand, it is assumed to be close to 1. This provides us with a different U which carries an approximate adjoint flux solution in its last

column. The rest of the columnspace then provides a means to correct for the approximation used.

One of the other interesting aspects of SVD lies in the behavior of the singular values. In the fine group problem, when lambda approaches the correct eigenvalue, the minimum singular value approaches 0. Since this is the singular value directly related to the approximate flux and adjoint flux solutions, this appears to be a major problem. This seems to contradict the Eckartt-Young theorem. This theorem states that if the SVD is applied in such a way that the singular values in S are ordered from largest to smallest, then choosing the highest N singular values produces a rank N matrix which optimally approximates the original matrix in a least-squares sense. In our problem, though, the minimum singular value corresponds to the fine group flux vector within V. Therefore, a singular value of 0 is counterintuitive since it actually coincides with the most important column of matrix V.

Upon further reflection, a singular value of 0 is the only way there can be solution of the form $M\varphi=0$. This becomes apparent if SVD of M is decomposed and the actions of U, S and V^T on φ are separated. First, V^T acts on φ, the exact fine group solution, and produces a vector with the only nonzero occurring at the bottom. If the minimum singular value were nonzero, then this value would have scaled according to the singular value and multiply U as a scalar. This would then produce a nonzero vector and not satisfy the eigenvalue problem. The only way $M\varphi=0$ can be satisfied in this case is if the minimum singular value of M is zero. As long as the initial choice of lambda is far enough away from the correct value, this will have little effect on our application.

From this analysis, it appears that SVD provides a set of natural basis functions in U and V that can include some of the important physics of the infinite homogeneous solution into their construction. Therefore, SVD is chosen as the primary method of preconditioning based on its ability to approximately diagonalize the problem and on the physical interpretation of the left and right preconditioners that SVD produces. These preconditioners can be defined according to Equation (7.55).

$$P_L = U^T$$
$$P_R = VS^{-1} \tag{7.55}$$

Assuming that the approximations applied in our SVD are good, application of these natural basis functions should shift the eigenvalues of the Jacobian very close to 1 and allow GMRES to converge much faster at each Newton step.

7.4.4 Results for Natural Basis Preconditioning

The DGM-SVD preconditioning method was first applied to the 295 group infinite homogeneous HTR problem. The cross sections chosen for this problem are those for the type 3 fuel specified in Section 5.2. Inexact Newton was used to prevent GMRES from wasting too much effort on unnecessary accuracy in the linear solve when the current iterate is still far away from the correct nonlinear solution. Since there are only 296 unknowns in this infinite homogeneous problem, memory requirements are low and full GMRES was chosen over restarted GMRES. The initial guess was assumed to be a flat spectrum. Nothing was used to initially seed the calculation to provide a better guess.

The natural basis functions were produced when SVD was applied to the original fine group problem, *M*, assuming an eigenvalue of 1. These energy preconditioners were then included into GMRES for the linear solve at each Newton step.

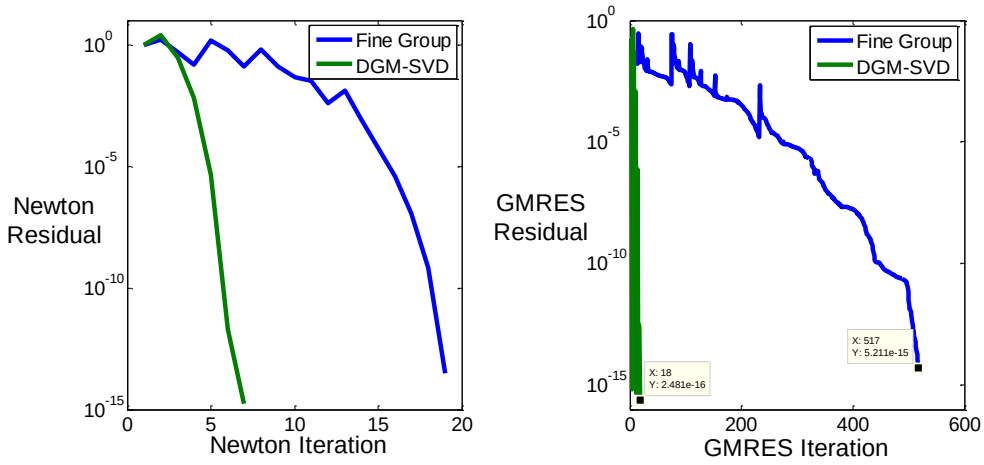

Figure 77 Accelerating convergence using the natural basis functions as a preconditioner. Newton iteration convergence located to the left and cumulative GMRES iterations to the right.

These natural basis functions are shown in Figure 77 to accelerate convergence of GMRES at each Newton step relative to the original fine group problem. Without preconditioning, the fine group problem took 19 total Newton steps and 517 total GMRES iterations to converge. No cutoff is applied for each Newton step and no restart is used either. When the natural basis functions were used as preconditioners for the energy problem, convergence required only 7 Newton steps and 18 total GMRES iterations. The original problem takes more Newton steps to converge due to the inexact nature of the linear solve at each step. When DGM-SVD is applied, convergence of the Newton residual matches exactly what the results when converging GMRES to very high tolerances. The natural basis function precondition the problem so well that the GMRES residual drops very faster, essentially producing an exact inversion at each Newton step. Since the residual drops far below the given tolerance for each linear solve, the inexact nature of these calculations becomes more like the exact Newton method. On average, this DGM-SVD method solves for the next search direction using only 3 GMRES iterations. The reason for this acceleration can be seen in the structure of our old and new Jacobians at the final Newton step in Figure 78.

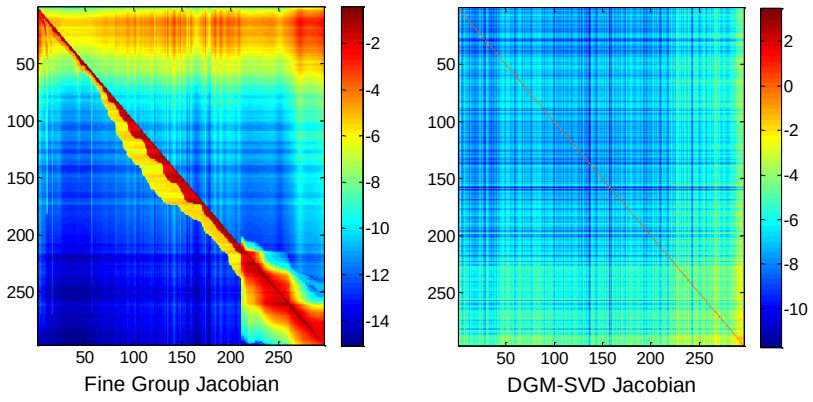

Figure 78 Comparing the magnitude of the elements in the Jacobian before and after preconditioning using the approximate SVD

Our old fine group problem contains a significant amount of off-diagonal information most of which are within 3 orders of magnitude of the diagonal values. Application of the natural basis functions, even though derived using an incorrect eigenvalue, is able to approximately diagonalize the original problem very well. Most of the off-diagonal elements are now 5 to 8 orders to magnitude lower than the diagonal. On top of this, it is noticed that much of the larger couplings in the preconditioned Jacobian are moved towards the last few columns and rows. This is expected since the columns and rows associated with the smallest singular values, at least in the case of the eigenvalue problems, contribute the most in terms of the physics of this problem.

While the diagonalization is easily seen by comparing these two Jacobians, the eigenvalue distributions are the true indicators of how much faster GMRES will converge. Therefore, the eigenvalues for each Jacobian were also calculated and plotted on the complex plane in Figure 79.

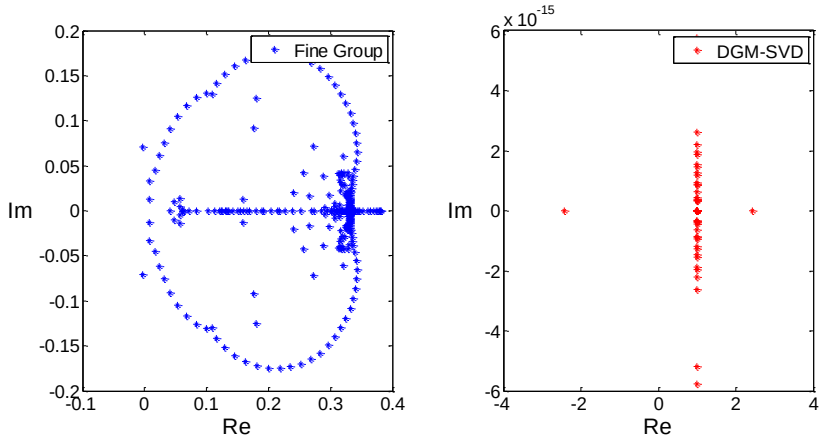

Figure 79 Comparing the eigenvalue distributions of the fine group problem before (left) and after (right) preconditioning with the natural basis functions

This is where the difference is truly interesting. While the eigenvalue distribution for the fine group problem is spread out across the entire complex plane, the DGM-SVD method has effectively clustered all but two of the eigenvalues near 1 and eliminated nearly all of the imaginary components of the eigenvalues. This accounts for the average convergence in roughly 3 GMRES iterations at each Newton step, since there are effectively only 3 distinct eigenvalues in the preconditioned Jacobian.

The DGM-SVD method comes at a cost, though, requiring $O(n^3)$ operations for dense matrices. This is a considerable cost for just solving an infinite homogeneous problem. To fairly compare this method against the original fine group problem, the cost of the SVD must be compared to the cost of the GMRES iterations, each of which is $O(n^2)$ due to the matrix-vector multiplication. For this 295 group problem, the DGM-SVD method required roughly 295 $O(n^2)$ calculations for the SVD while only 18 were necessary for the actual solver. The original problem required a total of 517 $O(n^2)$ calculations. Therefore, even with the cost of the SVD factored in, the DGM-SVD method was still able to outperform the Newton solver on the original fine group problem, reducing the total cost by roughly 40%. Still, this is considerably more expensive than conducting a single power iteration to converge the problem, since this only requires $O(n^2)$ operations.

One way to potentially reduce the cost of the SVD would be to reduce the total number of nonzeros in the original fine group problem. It is possible that something as simple as a cutoff could be applied to the problem would be effective, but remains outside the scope of this work.

7.5 Results for the 1D HTR Model Using Diffusion

While the nonlinear and preconditioning methods highlighted thus far show limited improvements in performance for infinite homogeneous problems, they could prove to be beneficial in spatially dependent problems. However, when applying Newton's method to a spatially dependent problem, it is not guaranteed that the method will converge to the fundamental mode of the eigenvalue problem. A good practice to avoid this complication is to apply a small number of power iterations to the initial guess first before using it in Newton's method. This will remove many of the higher order modes for the eigenvalue problem, allowing Newton's method to focus in on the fundamental mode. For these 1D diffusion examples, the 3 power iterations should remove the higher modes and allow Newton's method to converge to the fundamental mode.

7.5.1 Preconditioning Using Natural Basis Functions

The true potential of DGM-SVD really lies in its application to spatial methods, where the number of spatial unknowns greatly exceeds the number of different material types in the problem. This is because the SVD used to calculate the preconditioners needs to only be applied to an infinite homogeneous problem for each material. Using this approach, the $O(n^3)$ cost of the SVD is applied to a much smaller subset of the total problem and the preconditioners only need to be stored for each material type. The primary issue then becomes how to organize the problem and apply the preconditioners.

The typical method of arranging the fluxes in a given multigroup diffusion problem is by grouping all the fluxes associated with a single energy group into a block. This approach will be called energy blocking. Within each energy block, the fluxes are organized by some indexing with respect to their position in the geometry. Application of our preconditioners to the problem

in this form is not straight forward since the matrices produced by SVD need to be applied to all groups within a cell, not the other way around. Therefore, cell blocking will be applied instead. This new organizational scheme groups all fine group fluxes associated with a single cell together. Within each cell block, the fluxes are arranged from the lowest (fastest) group to the highest (most thermal) group. For the 13 cell diffusion fine group problem that will be considered later, the sparsity pattern seen on the right side of Figure 80 is much more compact.

The matrix is now block tridiagonal with the diagonal blocks, *D*, being dense and the off-diagonal blocks, *C* and *E*, being diagonal themselves. When the problem is formed in this manner, the left and right preconditioners can be applied in block diagonal form. The following serves as an example for a 6 cell problem.

$$\begin{bmatrix} P_{L,1} & 0 & \cdots & & 0 \\ 0 & P_{L,2} & & & \\ \vdots & & \ddots & \vdots & \\ & & P_{L,5} & 0 & \\ 0 & & \cdots & 0 & P_{L,6} \end{bmatrix} \begin{bmatrix} D_1 & C_1 & \cdots & & 0 \\ E_2 & D_2 & & & \\ \vdots & & \ddots & \vdots & \\ & & D_5 & C_5 & \\ 0 & & \cdots & E_6 & D_6 \end{bmatrix} \begin{bmatrix} P_{R,1} & 0 & \cdots & & 0 \\ 0 & P_{R,2} & & & \\ \vdots & & \ddots & \vdots & \\ & & P_{R,5} & 0 & \\ 0 & & \cdots & 0 & P_{R,6} \end{bmatrix} \quad (7.56)$$

Since DGM-SVD is only preconditioning the energy dependence of the problem, this method doesn't incorporate any spatial information from the problem. Therefore, the spatial problem also needs to be preconditioned. To do this without having to recalculate the spatial preconditioner at every Newton step, lambda is assumed to be 1. The first spatial preconditioner tested is just the inverse of the main diagonal from the SVD preconditioned problem. The next step is to conduct ILU(0) on the preconditioned problem. Unfortunately, applying any ILU decomposition, even if it is ILU(0), would prove unfruitful since the off diagonal blocks have been made dense in the process. Conducting ILU(0), since all possible fill-in spots from the original fine group problem are already nonzero, would lead to the exact LU decomposition.

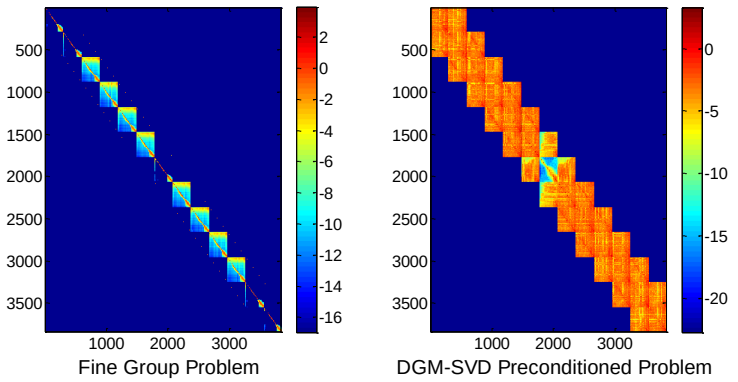

Figure 80 Comparing the magnitude of the elements in the 1D Jacobian before (left) and after (right) preconditioning with the approximate SVD

Therefore, the elements of the preconditioned matrix are truncated through application of a fixed ratio respect to the lowest valued term on the main diagonal. While naïve truncation in a matrix would be less than ideal, preconditioning with SVD has increased the magnitude of the most important moments through application of S^{-1}. It is this action that allows truncation to be applied in this simplistic fashion. Decreasing this ratio will allow more spatial information from the preconditioned problem to be used in the ILU(0) decomposition. Once this is completed, then ILU(0) is applied to the truncated problem. The resulting ILU decomposition should incorporate the necessary spatial information into the problem that is missing in the block cell SVD preconditioners. The right and left preconditioners must be modified in order to incorporate the incompletely LU decompositions.

$$P_L = \tilde{L}^{-1}U^T$$
$$P_R = VS^{-1}\tilde{U}^{-1} \tag{7.57}$$

Comparing these results to the unpreconditioned fine group problem is unfair since most whole core methods incorporate some form of ILU decomposition to precondition the problem. Therefore, normal ILU decomposition will also be used to precondition the fine group problem. First only the main diagonal will be used to precondition, then ILU(0) will be applied and lastly limited fill-in through application of a pivot threshold, τ, will be allowed. To incorporate more

information into the ILU(τ), the threshold will be reduced by an order of magnitude and the results compared to the DGM-SVD preconditioner. An exact LU decomposition of the fine group problem will be used as a reference to determine the relative performances of all these preconditioners.

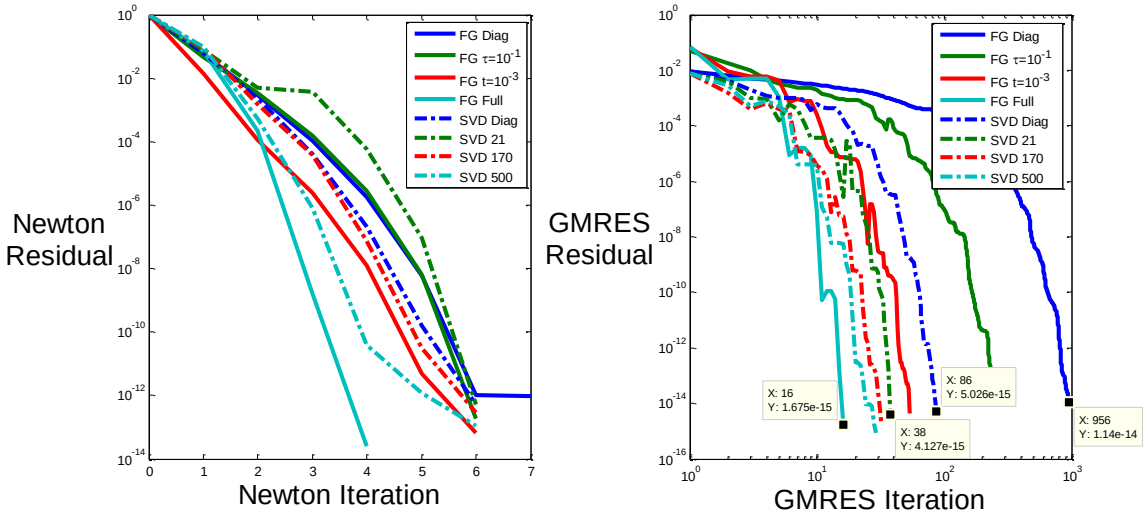

Figure 81 Comparison of precondition achieved using ILU(t) on the fine group problem and truncation of the SVD preconditioned problem.

In Figure 81 above, the diagonal preconditioner used in conjunction with the DGM-SVD provides significant improvements in convergence even relative to applying a pivot threshold of .1 to the original fine group problem. Even when applying a very large truncation ratio, 1/21, to the DGM-SVD problem, enough spatial information is incorporated into the preconditioner to allow further acceleration. At this point, DGM-SVD is converging using 38 iterations, only 20 additional iterations away from the fine group problem preconditioned using the exact LU decomposition. This is very interesting since preconditioning with the exact LU decomposition is equivalent to computing the exact inverse at each Newton step. However, further reducing the truncation ratio only produces slight improvements in convergence. This was thought to be a result of applying ILU(0) to building the spatial preconditioner instead of allowing some fill-in according to some tolerance. Therefore, ILU(τ) was also applied to each level of truncation

using a threshold of .001. The performance is compared to the original ILU(0) results in Figure 82.

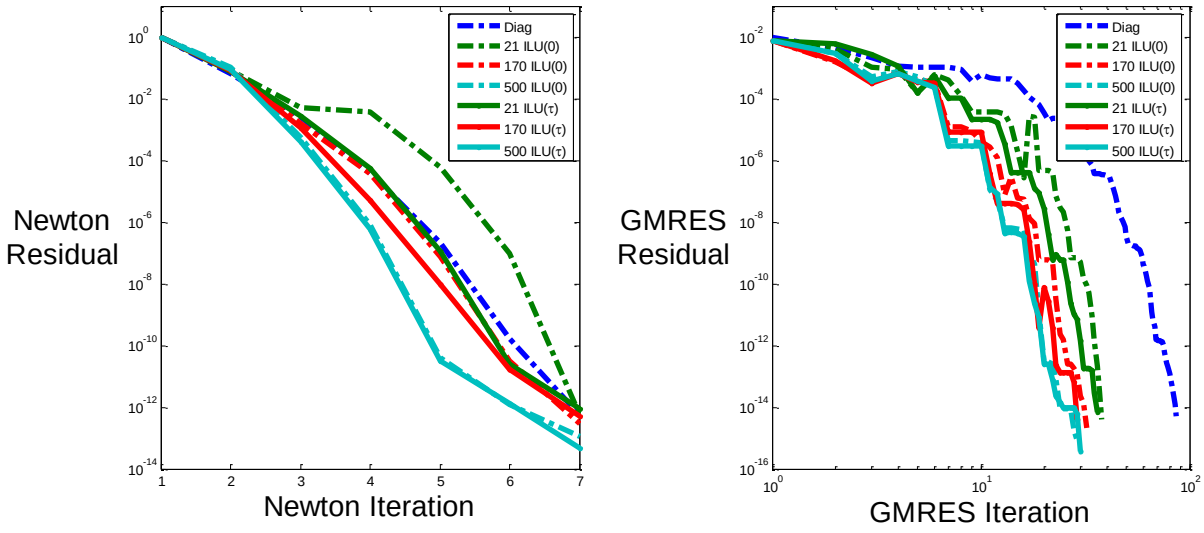

Figure 82 Comparison of SVD preconditioning when using ILU(t) in conjunction with naïve truncation

Even when a limited amount of fill-in is included into the ILU, this still doesn't provide any significant improvement on the convergence of the problem. The only way this makes sense is if the no-fill in approximation being used is of second order importance when compared to the truncation of elements in the first preconditioned matrix. In order to see what is really taking place, truncated ILU calculated using the 1/21 cutoff ratio was compared to the fine group ILU(τ) that produces the same total number of GMRES iterations, 38. The value of τ producing this level of preconditioning is 5.5×10^{-5}.

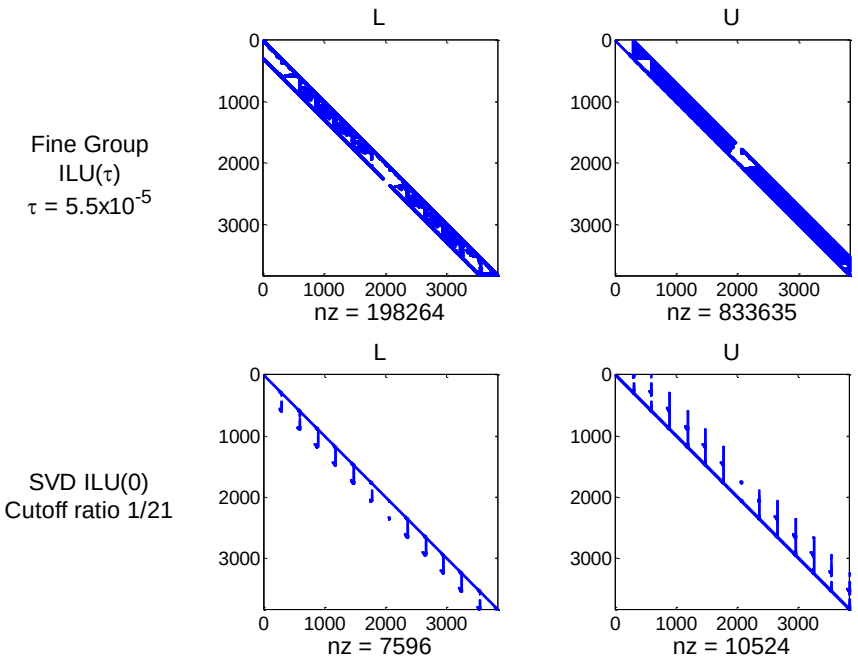

Figure 83 Comparison of the sparsity patterns of ILU(t) applied to the normal fine group problem (top) and ILU(0) applied the naïve truncation of the SVD preconditioned problem (bottom)

Indeed, DGM-SVD preconditioning is shown in Figure 83 to have placed the most important information in the column within each cell corresponding to the last column of *V*, which is an approximate fine group spectrum. This result is a direct consequence of the application of the inverse singular values as part of the DGM-SVD right preconditioner. Since an eigenvalue of 1 is assumed to be close to the correct value for the infinite homogeneous problem, this should result in a small singular value corresponding to the approximate fine group flux in the last column of *V*. Multiplying by the inverse of *V* therefore provides a large weight to the flux moment corresponding to this singular value. This can be seen when the DGM-SVD preconditioner is applied to the off-diagonal blocks.

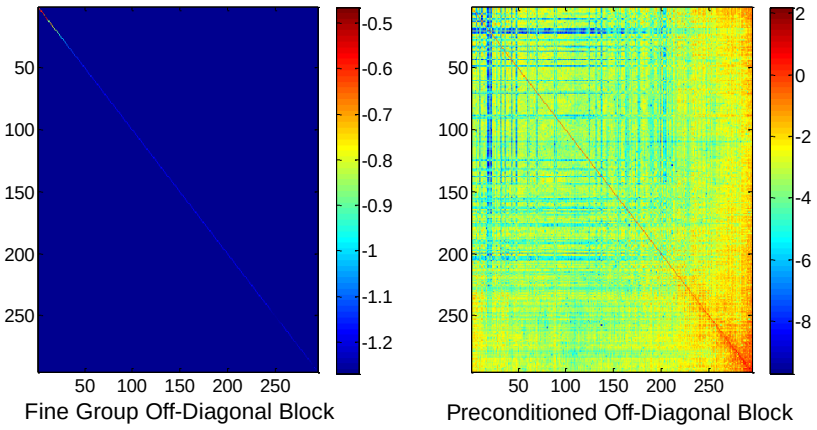

Figure 84 Comparing the off-diagonal components of the Jacobian for a 1D problem before (left) and after (right) preconditioning using the approximate SVD

The original off-diagonal block, itself, is diagonal. Once the right and left DGM-SVD preconditioners are applied, though, this becomes a dense matrix. Before we apply the inverse of the singular values, the matrix is filled with elements of roughly the same order of magnitude. In order to invert the approximate infinite homogeneous problem, it must also be multiplied by the inverse of the singular values. This produces a drastic shift in the importance placed in the right most columns of *V* corresponding to the smallest singular values. Even though there is still a diagonal clearly visible after this transformation, most of the information has been shifted to the right most column corresponding to the column in *V* containing the approximate fine group spectrum.

This shift also explains why including fill-in into this method did not produce significant gains in performance. Fill-ins due to LU decomposition occur when there are blank spaces located between the main diagonal and any off-diagonal elements within that matrix. Since the rightmost column of each cell block is the most important and is completely filled, allowing fill-in to take place will have no effect. The only fill-ins that occur are in the other columns that are mostly empty. If the assumption made in the SVD is good, then these fill-ins contribute very little information to the problem and result in almost no noticeable improvement in preconditioner performance.

In a way, the DGM-SVD preconditioning is able to highlight the most important information necessary for the spatial preconditioner and explains why such a sparse ILU with no fill-ins can outperform the ILU with fill-in applied to the original fine group problem. The true ability of DGM-SVD preconditioning is its ability to drastically reduce the storage required to effectively precondition a spatially dependent fine group problem. We can show this by plotting the total number of GMRES iterations required to converge the Newton method against the total number of nonzeros stored in the ILU decomposition.

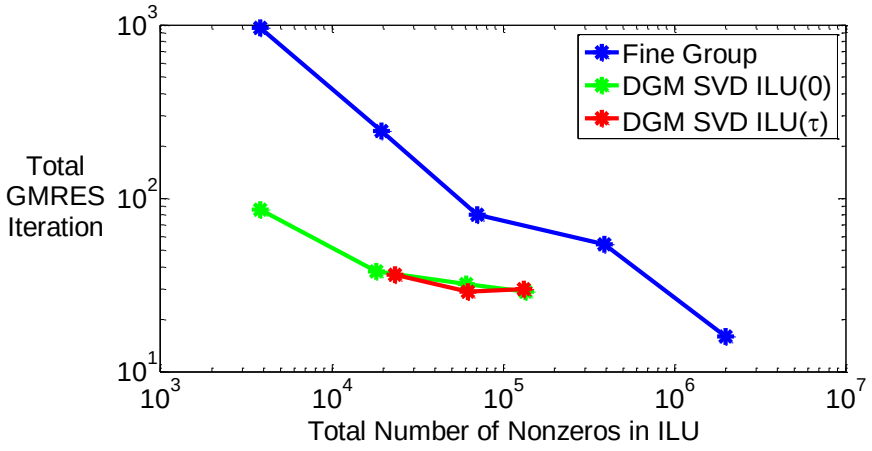

Figure 85 Total number of nonzeros in preconditioners using normal ILU and DGM-SVD

In Figure 85, it is observed that the true power of the DGM-SVD preconditioning method is in its ability to reduce the total storage required for the ILU decomposition. For the 38 GMRES iteration case in Figure 81, the DGM-SVD ILU decomposition required the storage of only 18,000 nonzeros while the equivalent fine group ILU decomposition required over 10^6. This is an improvement of almost two orders of magnitude. Of course, the reason so much information could be compressed into so little storage is because spatial preconditioning and energy preconditioning have become approximately decoupled.

If the total storage required for the SVD preconditioners were added into these results, Figure 85 above would look much less impressive. This is primarily because the 1D problem being tested still has roughly as many spatial unknowns as materials (13 cells and 4 materials). As the number of cells grows larger than the number of materials, it can be expected that the total

storage required for the SVD preconditioners will be significantly less relative to ILU with threshold pivoting for the full fine group problem.

7.6 Summary

The nonlinear process of recondensation has been successfully linearized and solved through use of an inexact Newton method, showing that further accelerations in the SEAM calculations of Sections 5 and 6 are possible. For the infinite homogeneous problem initially tested, improved acceleration of the Newton method curiously occurred when certain parts of the Jacobian were removed. When the normalization residual and the fission rate residuals were defined solely using the coarse group flux, the Newton method converged in only one step. This means that this transformation of the eigenvalue problem effectively linearized what would have been a nonlinear problem. Unfortunately this same approach did not work as well when applied to a 1D diffusion problem, likely because the lack of spectral swapping when the residuals are defined in this way produces inconsistencies in the model.

While the acceleration of the Newton method was ineffective for spatial problems, preconditioning the energy aspect of a coarse spatial problem resulted in very good reductions in the number of GMRES iterations required to converge the Newton method. In this application, taking the SVD of the infinite homogeneous problem produced a set of natural basis functions which have the very nice feature of containing an approximate fine group forward and adjoint flux in its columnspaces. This allows for a significant portion of the information to be moved towards the right most column of each energy block. In doing so, this allows for an extremely efficient preconditioner to be built with ILU(0) for the problem tested in this section. This definitely merits further work in determining if this approach with SVD can provide the same level of acceleration to more realistic problems.

8 Conclusions

Achieving high fidelity analysis for reactor neutronics is of utmost importance in analyzing the performance and safety of reactor designs. Since quantities such as power density are directly proportional to fission rate density, which in turn is dependent on the neutron distribution, accurate descriptions of neutrons in core are key to performing adequate steady-state and transient safety analyses. These analyses are conducted by solving the neutron transport equation.

While fine group whole core reactor analysis remains one of the long sought after goals of the reactor physics community, such a detailed analysis is typically too computationally expensive to be realized on anything outside of supercomputers. However, recondensation using the Discrete Generalized Multigroup (DGM) method offers a relatively cheap alternative to solving the full fine group transport problem. DGM forms a set of moment equations, of which the leading order is equivalent to a coarse group eigenvalue problem. All other moment equations can be described as fixed source sweeps relying solely on the solution of the coarse group flux to provide the source. Therefore, once the coarse group solution is found, solving the higher moment equations provides a cheap way to reconstruct an approximate fine group flux solution.

Since the full fine group solution isn't known *a priori*, though, a nonlinear iteration is used to update the coarse group constants and moment cross sections using the updated fine group flux upon solving the DGM equations. This process is called recondensation.

While this process was shown to converge to the exact fine group solution when applied to step difference Sn, inconsistencies arose when applied to high order spatial methods such as step characteristics.

Therefore, the goals of this research were two-fold. First, this spatial inconsistency needed to be fully corrected to ensure that the convergence to the fine group solution is guaranteed across all spatial methods. Second, recondensation was applied to significantly reduce the number of fine group transport sweeps required to converge to the fine group solution.

8.1 Spatial Recondensation

Spatial recondensation was the first step taken towards addressing spatial inconsistencies between DGM and the original fine group problem. The theory behind this work was that while the cross sections were collapsed assuming spatially flat angular and scalar fluxes, in step characteristics these fluxes actually do assume a distinct shape across a given cell. Including this spatial variation explicitly into the coarse group collapse process would then be able to reintroduce some of the spatial information usually lost in the normal group collapse procedure.

This approach was tested using a 1-D High Order MOC (HOMOC) method able developed for the purposes of spatial recondensation. 1D HOMOC provided an accurate description of the angular and scalar fluxes to an arbitrarily high order. The result of this was a first of its kind attempt at producing and incorporating high order spatial variation into cross sections within the context of multigroup collapse.

When these high order spatial cross sections were included into the DGM equations and the recondensation process, the result was an asymptotically improving agreement between the fine group and recondensation solution with every spatial order added to represent the cross sections. A similar result was also achieved for a 0^{th} order 1D MOC calculation by incorporating spatial information from a high order spatial DGM calculation back into the 0^{th} order problem.

While this first of its kind approach worked, albeit asymptotically, it also suffered from exceeding large memory requirements when considering realistic transport problems. The del term would be stored for each discrete angle, DGM moment and spatial quadrature point, increasing memory requirements relative to the original fine group problem by a factor of 100. While this approach was abandoned, maintaining consistency between fine and coarse group transport solutions using higher order spatial methods was still of utmost importance to realizing the objectives of this thesis. Therefore a different approach was taken that could exactly resolve the inconsistencies between DGM and the original fine group problem given any spatial method.

8.2 Exact Recondensation Using DGM

One of the initial assumptions made in the original derivation of DGM was that no spatial method was yet applied. While this created a general del term for use in correct transport calculations, it could not fully capture the effects of higher order methods such as step characteristics. Therefore, step characteristics first applied to the multigroup transport problem first and then DGM was derived through the application of discrete basis functions. The result of this was the production of two del terms, one that corrects the incoming-outgoing angular flux relationship and another which corrects the cell-averaged angular flux.

Since these two terms theoretically correct for all the errors accrued when moving to the DGM calculation, the fission and scattering cross sections could be treated through the normal coarse group collapse procedure. The results offered in this thesis revealed that the errors using the original del term slowly decrease as mesh size decreases while the exact del term could entirely remove all errors due to spatial inconsistencies.

Alternatively, all of these tools could be used instead to produce a considerably cheap improvement assuming a close initial guess of the coarse group flux from an assembly level calculation. Unfortunately, in order to accurately account for the missing spatial information in the DGM moment equations the exact del term definitions established in this section required an inordinate amount of memory for realistic transport problems. Once again, the del term needed to be stored for every angle, cell and DGM moment. If exact convergence to the fine group solution were not required, future work could investigate ways of approximating this exact del term in an efficient manner through representation with continuous basis functions in angle to provide cheaper corrections to a coarse group solution. An approximate del term could also be produced for a given pin cell or assembly so as to provide corrections informed by the errors of the high spatial order method used.

In its current state, though, exact recondensation remains limited by the memory requirements of the del term defined using DGM. Therefore, the next section will continue addressing this issue

by relaxing some of the assumption made in the original derivation of DGM to determine if exact recondensation can be made feasible for whole core transport calculations.

8.3 The Source Equivalence Acceleration Method (SEAM)

A new method called SEAM was derived through further simplifications of the del term using intuition gained from previous work with DGM and also from the SubGroup Decomposition method (SGD) developed at Georgia Tech. These simplifications drastically reduced the storage required to produce a fully consistent transport calculation despite the inherent differences in the transport operator when moving from the fine group to coarse group problem. Instead of storing a corrective term for each cell, fine group, and discrete angle, these new factors only need to be stored for each cell and coarse group.

Not only did this new method reduce memory requirements, but it also came at little extra computational cost on top of the normal problem. This method only requires a fixed source coarse group sweep using the collapsed fine group source from the previous fine group calculation. Initially, the collapsed fine group flux from the fine group sweeps was divided by this intermediate coarse group flux to produce a Source Equivalence Factor (SEF). This factor could be multiplied with the coarse group scalar flux after each coarse group sweep to produce a corrected coarse group flux. Further inspection revealed that this procedure was identical to using the fixed source coarse group flux to normalize the fine group reaction rates in the group collapse procedure instead of the collapsed fine group flux.

In both cases, the main effect was a recreation of the fine group source in the coarse group sweep regardless of how inconsistent the coarse group transport operator is relative to the fine group. Therefore this method was called the Source Equivalence Acceleration Method (SEAM). This method was found to be much more general in its application, though, than just correcting for errors accrued in the group collapse process. By definition, the formulation of the Source Equivalence Factors enables the use of any inconsistent method, whether it's a lower order angle approximation or an entirely different transport method altogether from the one used in the fine group calculation. SEAM's simplicity and flexibility are hard to match.

After testing on a number of 1D fine group cores, SEAM was successfully implemented in OpenMOC and tested on two different 2D problems, the C5G7 benchmark core and a newly created problem called the C5G361 core.

Results from C5G7 testing showed that SEAM can create a substantially cheaper transport problem that is still equivalent to the full transport problem although MOC is a higher order spatial method. As an example, a 4 group calculation using 4 azimuthal angles, a 0.14 cm track spacing, 1 polar angle with step difference was made equivalent to the full 7 group transport calculation using 32 azimuthal angles, a 0.02 cm track spacing, 3 polar angles with MOC using SEAM.

Initial comparison with CMFD showed that although 4 group SEAM calculations could match the number of full transport calculations used in CMFD, the relatively cheap transport calculations still required much more effort than conducting the much cheaper diffusion calculations in CMFD. However, when moving to large numbers of azimuthal angles like the 128 azimuthal angle example, 4 group SEAM was able to outperform CMFD because of the additional overhead required to tally the surface partial currents for each of the homogenized pin cells when creating the equivalent diffusion problem. While SEAM's equivalent transport problem is more expensive to solve, it's formation only requires a single fixed source sweep across the low order problem. In this way the overhead only scales with the low order problem, not the high order on. Since the partial current tally overhead is proportional to the angular order and track spacing used, CMFD eventually reaches the point where the overhead is much greater than the gains in the diffusion calculations. This point was reached in the 128 azimuthal angle, 0.01 cm calculation. In this case, the 4 group SEAM calculation converged in about 1500 seconds as opposed to CMFD's time of 1900 seconds.

SEAM was also able to accelerate the C5G7 problem without using group collapse. The result was a purely angular acceleration approach that further improved the 128 azimuthal angle results. The 7 group SEAM calculation was able to cut the number of full transport sweeps nearly in half relative to the CMFD calculation. In doing so, 7 group SEAM converged in only

1100 seconds, resulting in a final speed up of 45 relative to the normal power iteration calculation.

Unfortunately, testing with the C5G7 benchmark core was unable to provide an adequate testing base for a true fine group transport problem, so a new 361 group problem was formed to truly test the performance of energy acceleration using SEAM. This "C5G361" problem was based almost exactly off of the geometry of the C5G7 core, but instead used a 361 group microscopic cross section data set representative of a MOX-load LWR core. This new problem was tested with CMFD using pin cell homogenization to compare with the 9 and 23 group calculation using SEAM. The CMFD calculation was run using the same coarse group structures as the SEAM calculations

While the 9 group calculation using 1 power iteration converged using only 40 fine group transport sweeps, this approach would limit the flexibility of SEAM by requiring use of the same angular discretization as the fine group problem. When the 9 and 23 group SEAM calculations used 2 power iterations convergence still occurred after 44 fine group transport sweeps as compared to the 50 used in the CMFD calculation. While switching to 2 power iterations did slightly degrade performance, this also provided additional stability and would allow for further coupling between low and high order angular discretization if desired.

The improvement over CMFD in this case suggests that, while the diffusion calculation contains a good approximation of the global spatial problem, convergence is limited by changes in the solution at the pin cell level which aren't captured with a homogenized problem. Since SEAM uses no spatial homogenization, the scalar fluxes for each flat source region can be accelerated, leading to the better performance observed for this 361 group problem.

SEAM was shown to accelerate both fine group and high fidelity transport calculations, and in some cases even improving upon the acceleration of CMFD. While further study is merited, it is important to recognize that the recondensation process used in SEAM is very similar to that used in DGM. Even though DGM proved ultimately too memory intensive, it may still hold valuable intuition into how to further improve SEAM in the future.

8.4 Nonlinear Recondensation and Energy Preconditioning

The nonlinear process of recondensation was successfully solved through use of an inexact Newton method, showing that further accelerations in the SEAM calculations could be possible. For the infinite homogeneous problem initially tested, improved acceleration of the Newton method curiously occurred when certain parts of the Jacobian were removed.

When the normalization residual and the fission rate residuals were defined solely using the coarse group flux, the Newton method converged in only one step. This meant that this transformation of the eigenvalue problem effectively linearized what would have been an otherwise nonlinear problem. Unfortunately this same approach did not work as well when applied to a 1D diffusion problem, likely because the lack of spectral swapping when the residuals are defined in this way produces inconsistencies in the model.

While the acceleration of the Newton method was ineffective for spatial problems, preconditioning the energy aspect of a 1D coarse mesh spatial problem resulted in very good reductions in the cumulative number of Generalized Minimum RESidual (GMRES) iterations required to converge the Newton method. In this application, taking the Singular Value Decomposition (SVD) of the infinite homogeneous problem produced a set of natural basis functions which have the very nice feature of containing an approximate fine group forward and adjoint flux in the columnspaces of U and V.

This allows for a significant portion of the information to be moved towards the right most column of each energy block, resulting in a substantial reduction in storage for the preconditioner. Therefore, a very efficient preconditioner was built with ILU(0) for the problem tested in this section. This interesting approach merits further work in determining if ILU(0) and SVD can provide this same level of acceleration to more realistic problems or perhaps be used in conjunction with SEAM.

8.5 Summary and Future Work

The discoveries, tools and methods developed in this thesis provide a strong foundation for further accelerating fine group whole core neutron transport calculations. While this work began with the goal of maintaining consistency in the recondensation process using DGM, it led to the development of many first of their kind methods.

- 1D HOMOC was developed in order to provide arbitrarily high order spatial angular and scalar fluxes for use in spatial recondensation. In-depth analysis showed that this method provided $2(P+1)$ spatial convergence

- High order spatial cross section definitions were developed and employed in DGM and recondensation to enforce consistency between fine and coarse group calculations. First of its kind method, called spatial recondensation, produced asymptotic consistency but unfortunately at the expense of memory requirements.

- Exact recondensation was derived specifically for 1D step characteristics and allowed recondensation to converge the coarse group problem to the exact fine group solution for the first time. However, memory requirements made this approach largely infeasible when considering realistic 2D problems.

- SEAM, developed using various aspects of these previous attempts, enabled the first ever generalized, yet fully consistent, coupling between transport-transport and transport-diffusion calculations in the context of the multigroup collapse procedure. Implementation and testing in OpenMOC revealed its additional ability to accelerate the angular problem by forming equivalence between completely different numbers of azimuthal angles and track spacings. For fine group problems or problems using very fine angular discretizations, SEAM was able to match the performance of or outperform equivalent CMFD calculations.

- Natural basis functions for DGM were developed through the application of SVD on the infinite homogeneous problem and successfully preconditioned GMRES used in an

inexact Newton method to solve a 1D diffusion eigenvalue problem. The decoupling of energy and spatial dependences provided a very efficient preconditioner, both in terms of its storage, ease of inversion and reduction of GMRES iterations.

All of these new methods open up many avenues for future work.

8.5.1 Cheap Improvement of Whole Core Coarse Group Solution

One of the possibilities is to take the fine group fluxes from an assembly level calculation and conduct one or two power iterations using them at the whole core level. Using the collapsed fine group source, a set of core wide coarse group cross sections could be produced for each cell using a very cheap transport method with a low order angular approximation. Once the coarse group problem converged, theoretically, the reaction rates and eigenvalue should be closer to the true fine group solution than had the ordinary assembly collapsed cross sections been used to conduct the whole core analysis. One primary difference with this approach is the cell by cell cross sections produced using SEAM. This would give it an advantage over calculations using assembly or pin cell homogenized cross sections in terms of driving the solution closer on a single whole core sweep.

8.5.2 *A Priori* Generation of Assembly Level Method Corrected Cross Sections

Another possibility could be to produce few group assembly cross sections with SEAM using a single recondensation. The idea would be to conduct the normal fine group assembly level calculation, but then apply SEAM with which ever low order angular discretization and coarse group structure the whole core calculation will be using. An approach could be take like that of the SPH method. The SPH approach, however, requires a full iterative process to arrive at SPH factors since these modify the homogenized and condensed source as well as the total cross sections. On the other hand, SEAM only relies on the source of the converged fine group assembly solution, so the only cost of this calculation would be to converge the boundary fluxes of the coarse group problem. This would then produce a fully heterogeneous set of equivalent cross sections as that could simulate a high order angular, fine group whole core transport

calculation on the cheap. Discontinuity factors at the assembly, pin cell or even individual mesh level could also be coupled with this approach to provide better agreement with conditions in the core.

8.5.3 Combined SEAM and CMFD Iterations

Both SEAM and CMFD have their strengths and their weaknesses. With SEAM the level of spatial heterogeneity give it the advantage to accelerate the spatial convergence of a whole core transport solution. Unfortunately, this can't be coupled with a many or fine group calculation, otherwise the cost of the acceleration method would be just as high as the whole core transport calculation itself. CMFD, though, can homogenize either assemblies or pin cells to reduce the spatial degrees of freedom and allow it to focus more attention on the global spatial solution using better spectral resolution. Instead of the two methods competing over different parts of the problem, it may work even better to have both SEAM and CMFD work hand in hand. A single whole core transport sweep could be used to set up both the CMFD problem and the SEAM problem. The two could then iterate between each other with CMFD passing improved spectral information to SEAM and SEAM passing much improved spatial fluxes to CMFD. This seems to have the potential to further reduce the number of fine group whole core transport sweeps.

8.5.4 Unstructured and Heterogeneous Diffusion Acceleration Using SEAM

As was done in 1D, an equivalent fully heterogeneous diffusion problem could be constructed with SEAM in place of a lower order transport problem to further reduce the cost of solving the equivalent problem. Since even the cheapest transport calculations tested in this work were still 9-10 times more expensive than the diffusion solve used in CMFD, this has the potential to further reduce the cost of the SEAM calculation and increase speed up even further.

8.5.5 Application of Newton Methods to Recondensation Using SEAM

Lastly, it may be possible to recast the full 2D recondensation procedure using SEAM as a nonlinear system of equations and apply inexact Newton with preconditioning using the natural basis functions derived in this work. The quadratic convergence of the Newton method could

further reduce the number of fine group, whole core transport sweeps required to converge the solution, especially if used in conjunction with efficient preconditioners. This represents just a small fraction of what could be possible with the tools developed in this thesis.

9 References

1. Ahmed, N., Natarajan, T., & Rao, K. R. (1974). Discrete Cosine Transform. *IEEE Transactions on Computers*, C-23 (1), 90-93.
2. Boyd, W., Shaner, S., Li, L., Forget, B., & Smith, K. S. (In Press). The OpenMOC Method of Characteristics Neutral Particle Transport Code. *Annals of Nuclear Energy*, Manuscript submitted for publication
3. Brown, F.B., et al. (2003). MCNP – A General Monte Carlo N-Particle Transport Code, Version 5; Volume I: Overview and Theory. *LA-UR-03-1987*.
4. Chiba, G., Tsuji, M., Sugiyama, K., & Narabayashi, T. (2012). A note on application of superhomogeneisation factors to integro-differential neutron transport equations. *Journal of Nuclear Science and Technology*, 49 (2), 272-280.
5. Cho, J. Y., Joo, H. G., Park, S.Y., & Zee, S.Q. (2002). Consistent Group Collapsing Scheme for Multigroup MOC Calculations. *Proceedings from PHYSOR Conference*. Seoul, Korea.
6. Cho, N.Z. (2012). On Some Outstanding Problems in Nuclear Reactor Analysis. *Nuclear Engineering and Technology*, 44 (2).
7. Cho, N. Z., & Park, C. J. (2002). A Comparison of Coarse Mesh Rebalance (CMR) and Coarse Mesh Finite Difference (CMFD) Acceleration Methodologies for Neutron Transport Calculations. *NURAPT*.
8. Davidson, G. G., Evans, T. M., Jarrell, J. J. & Slaybaugh, R. N. (2011). Massively Parallel, Three-Dimensional Transport Solution for the k-Eigenvalue Problem. *Proceedings from International Conferences on Mathematics and Computational Methods Applied to Nuclear Science and Engineering*. Rio de Janeiro, RJ, Brazil.
9. Douglass, S., & Rahnema, F. Subgroup decomposition method. *Annals of Nuclear Energy*, 48, 84-101.
10. Everson, M., & Forget, B. (2012). A 1D Analysis of Two High Order MOC Methods. *Proceedings from PHYSOR Conference*. Knoxville, TN, USA.

11. Fridman, E., Duerigen, S., Bilodid, Y., Kotlyar, D., & Shwageraus, E. (2013). Axial discontinuity factors for the nodal diffusion analysis of high conversion BWR cores. *Annals of Nuclear Energy*, 62, 129-136.
12. Gibson, N., & Forget, B. (2012). Eliminating Flux Updates from the Discrete Generalized Multigroup Method. *Proceedings from American Nuclear Society – Winter Meeting*. San Diego, CA, USA.
13. Grundmann, U., & Mittag, S. (2011). Super-homogenisation factors in pinwise calculations by the reactor dynamics code DYN3D. *Annals of Nuclear Energy*, 38, 2111-2119.
14. Hebert, A. (1993). A Consistent Technique for the Pin-by-Pin Homogenization of a Pressurized Water Reactor Assembly. *Nuclear Science and Engineering*, 113, 227-238.
15. Hebert, A. (2009). *Applied Reactor Physics*. Montreal, QC, CA: Presses Internationales Polytechnique.
16. Herman, B. R., Forget, B., & Smith, K. (2013). Utilizing CMFD in OpenMC to estimate dominance ratio and adjoint. *Transactions of the American Nuclear Society*, 109, 1389-1392.
17. Herrero, J. J., Garcia-Herranz, N., Cuervo, D., & Ahnert, C. (2012). Neighborhood-corrected interface discontinuity factors for multi-group pin-by-pin diffusion calculations for LWR. *Annals of Nuclear Energy*, 46, 106-115.
18. Hong, S. G., Kim, K. S., & Song, J. S. (2010). Fourier Convergence Analysis of the Rebalance Methods for Discrete Ordinates Transport Equations in Eigenvalue Problems. *Nuclear Science and Engineering*, 164, 33-52.
19. Kim, K. S., & DeHart, M. D. (2011). Unstructured partial- and net-current based coarse mesh finite difference acceleration applied to the extended step characteristics method in NEWT. *Annals of Nuclear Energy*, 38, 527-534.
20. Larsen, E. The Suppression of Energy Discretization Errors in Multigroup Transport Calculations. *NNSA Integrated University Programs Final Report, NEUP 09-839*.

21. Lee, D. (2012). Convergence analysis of coarse mesh finite difference method applied to two-group three-dimensional neutron diffusion problem. *Journal of Nuclear Science and Technology*, 49 (9), 926-936.
22. Lee, M. J., Joo, H. G., Lee, D., & Smith, K. (2014). Coarse mesh finite difference formulation for accelerated Monte Carlo eigenvalue calculation. *Annals of Nuclear Energy*, 65, 101-113.
23. Li, L. (2013). A Low Order Acceleration Scheme for Solving the Neutron Transport Equation. (Unpublished Master of Science Thesis). Massachusetts Institute of Technology, Cambridge, MA, USA.
24. Neuman, C. P., & Schonbach, D. I. (1974). Discrete (Legendre) Orthogonal Polynomials – A Survey. *International Journal for Numerical Methods in Engineering*, 8 (4), 743-770.
25. Nuclear Energy Agency; Organisation for Economic Co-operation and Development. (2003). *Benchmark on Deterministic Transport Calculations Without Spatial Homogenization* (ISBN 92-64-02139-6). Danvers, MA: Copyright Clearance Center.
26. Park, R. P. & Cho, N.Z. (2004). Coarse-Mesh Angular Dependent Rebalance Acceleration of the Discrete Ordinates Transport Calculations. *Nuclear Science and Engineering*, 148, 355-373.
27. Park, R. P. & Cho, N.Z. (2004). Coarse-Mesh Angular Dependent Rebalance Acceleration of the Method of Characteristics in x-y Geometry. *Nuclear Science and Engineering*, 158, 154-163.
28. Rahnema, F., Douglass, S., & Forget, B. (2008). Generalized Energy Condensation Theory. *Nuclear Science and Engineering*, 160, 41-58.
29. Romano, P., & Forget, B. (2013). The OpenMC Monte Carlo particle transport code *Annals of Nuclear Energy*, 51, 274-281.
30. Sanchez, R. (2009). Assembly Homogenization Techniques for Core Calculations *Progress in Nuclear Energy*, 51, 14-31.
31. Smith, K. S. (1986). Assembly Homogenization Techniques for Light Water Reactor Analysis. *Progress in Nuclear Energy*, 17 (3), 303-335.

32. Smith, K. S., & Rhodes, J. (2002). Full-Core, 2-D, LWR core calculations with CASMO-4E. *Proceedings from PHYSOR Conference*. Seoul, Korea.
33. Xuhua, Z., Fu, L., & Dengying, W. (2009). Application of the discontinuity factor theory to accelerate routine rod worth calculations for modular HTRs. *Nuclear Engineering and Design*, 239, 281-288.
34. Yamamoto, A., & Endo, T. (2011). Overview of Core Simulation Methodologies for Light Water Reactor Analysis. *Nuclear Safety and Simulation*, 2 (1).
35. Yamamoto, A., Kitamura, Y., Ushio, T., & Sugimura, N. (2004). Convergence Improvement of Coarse Mesh Rebalance Method for Neutron Transport Calculations. *Journal of Nuclear Science and Technology*, 41 (8).
36. Yamamoto, A., Kitamura, Y., & Yamane, T. (2004). Cell Homogenization Methods for Pin-by-pin Core Calculations Tested in Slab Geometry. *Annals of Nuclear Energy*, 31, 825-847.
37. Yamamoto, A., Tabuchi, M., Sugimura, N., Ushio, T., & Mori, M. (2007). Derivation of Optimum Polar Angle Quadrature Set for the Method of Characteristics Based on Approximation Error for the Bickley Function. *Nuclear Science and Technology*, 44 (2), 129-136.
38. Yasseri, S., & Rahnema, F. (2013). Subgroup decomposition method in diffusion theory. *Annals of Nuclear Energy*, 60, 235-241.
39. Zhong, Z., Downar, T. J., & Xu, Y. (2008). Implementation of Two-Level Coarse-Mesh Finite Difference Acceleration in an Arbitrary Geometry, Two-Dimensional Discrete Ordinates Transport Method. *Nuclear Science and Engineering*, 158, 289-298.
40. Zhu, L., & Forget, B. (2010). A Discrete Generalized Multigroup Energy Expansion Theory. *Nuclear Science and Engineering*, 166, 239-253.
41. Zhu, L., & Forget, B. (2011). An Energy Recondensation Method Using the Discrete Generalized Multigroup Energy Expansion Theory. *Annals of Nuclear Energy*, 38, 1716-1727.
42. Zhu, L. (2012). Discrete generalized multigroup theory and applications. (Doctoral Dissertation). Retrieved from DSpace@MIT. (http://hdl.handle.net/1721.1/76503).

Appendix A

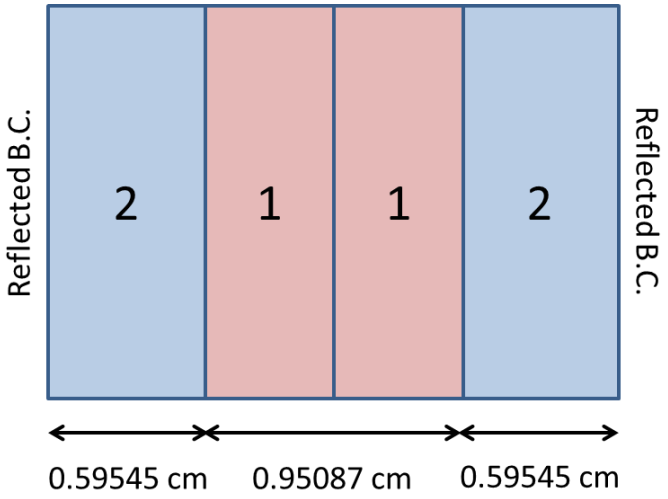

Figure 86 Specification of 1D test geometry

Table 21 4 Group Constants for Material 1

	$\Sigma_{T,g}^{1}$	$\nu\Sigma_{f,g}^{1}$	χ_{g}^{1}
g=1	0.2020132	0.0048928	0.9928300
g=2	0.4876502	0.0128748	0.0071700
g=3	0.7546648	0.1206095	0.0
g=4	1.2963174	0.2599271	0.0

Table 22 Scattering Cross Sections for Material 1

$\Sigma_{s,g\to h}^{1}$	g=1	g=2	g=3	g=4
h=1	0.1566631	0.0	0.0	0.0
h=2	0.0426728	0.4267787	0.0	0.0
h=3	0.0	0.0305304	0.4363064	0.3152042
h=4	0.0	0.0055881	0.2367508	0.8342739

Table 23 4 Group Constants for Material 2

	$\Sigma_{T,g}^2$	$\nu\Sigma_{f,g}^2$	χ_g^2
g=1	0.1807996	0.0	0.0
g=2	0.5365751	0.0	0.0
g=3	0.8700957	0.0	0.0
g=4	1.5457978	0.0	0.0

Table 24 Scattering Cross Sections for Material 2

$\Sigma_{s,g\to h}^2$	g=1	g=2	g=3	g=4
h=1	0.1188257	0.0	0.0	0.0
h=2	0.0617344	0.4765734	0.0	0.0
h=3	0.0	0.0494727	0.4651354	0.2681453
h=4	0.0	0.0094437	0.3942237	1.2664212

Appendix B

A 1D test geometry comprised of 7 BWR fuel assemblies of differing material composition is used to evaluate the ability of high order spatial recondensation to regain accuracy otherwise lost in moving from the fine group equations to the DGM equations. Two different assembly types were used to construct the core, one using an even distribution of low enriched fuel, the other containing an inner region of high enriched fuel. The thickness of the inner slab is set at 6.5024 cm and the outer slab thickness is set at 3.2512 cm. 1.1176 cm. spacings containing water are placed at the edges of the assemblies.

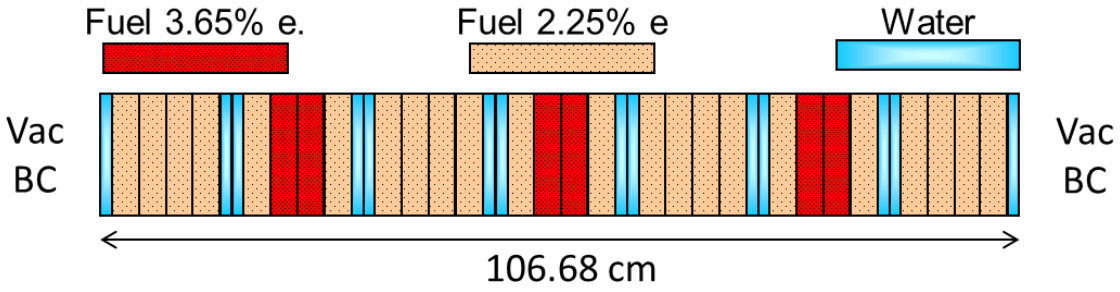

Figure 87 : 1D BWR core test geometry used for spatial recondensation testing

For each assembly the inner fuel region is evenly distributed between 3 cells, the outer fuel slabs are each split into 2 cells and 1 cell represents the water regions on each of the assembly edges. Vacuum boundary conditions are applied to the outer edges of the core. The same coarse group structure used in the assembly calculations is used in this analysis as well.