A GENERALIZED MULTIGROUP METHOD

A Thesis

by

ANDREW THOMAS TILL

Submitted to the Office of Graduate and Professional Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

Chair of Committee,	Marvin Adams
Co-Chair of Committee,	Jim Morel
Committee Member,	Jean-Luc Guermond
Head of Department,	Yassin Hassan

May 2014

Major Subject: Nuclear Engineering

ABSTRACT

The standard multigroup (MG) method for energy discretization of the transport equation can be sensitive to approximation in the weighting spectrum chosen for cross-section averaging. As a result, MG often inaccurately treats important phenomena such as self-shielding variations across a fuel pin in nuclear reactor simulations. From a finite-element viewpoint, MG uses a single fixed basis function (the pre-selected spectrum) within each group, with no mechanism to adapt to local spatial and angular solution realities. To address these issues, we introduce a Petrov-Galerkin finite-element multigroup (PG-FEMG) method, a generalization of the MG method that is related to the family of multiband (MB) methods. PG-FEMG uses integrals over several discontinuous energy domains within a group as its degrees of freedom, which allows PG-FEMG to be used in standard MG-based computer codes with changes to pre- and post-processing of the data only. We define a problem-wide effective total cross section as the basis of these discontinuous energy domains.

We implement the PG-FEMG method for several realistic pin-cell problems and find it to be significantly more accurate per degree of freedom than MG for several quantities of interest, including criticality eigenvalue and power profile shape. We find that PG-FEMG is much less sensitive to errors in weighting spectra compared to standard MG. We discuss straightforward generalizations to multi-dimensional problems of practical interest, including reactor depletion calculations.

DEDICATION

> "Do not be deceived, my beloved brethren. Every good thing given and every perfect gift is from above, coming down from the Father of lights, with whom there is no variation or shifting shadow" (James 1:16-17; NASB).

This thesis is dedicated to my parents, Thomas and Kimberly Till, who have always encouraged my natural curiosity, who instilled, at an early age, an indomitable Protestant work ethic, and who have unceasingly supported me in my academic endeavors, even when such endeavors led me far afield.

ACKNOWLEDGEMENTS

Though doing so in full would fill a volume larger than this already formidably sized document, I would like to acknowledge a few individuals and organizations who have made this work possible. First, to my co-chairs, Drs. Adams and Morel, who provided me with sufficient latitude in my research to investigate a topic outside the norm even before I received my fellowship and who selflessly provide continual guidance, thank you. Secondly, to my community at Texas A&M University, nuclear engineers, mathematicians and computer scientists alike, I have relished the opportunities and interactions you have fostered. Finally, to the Krell Institute and the Department of Energy Computational Science Graduate Fellowship program (DOE CSGF - grant number DE-FG02-97ER25308) for your funding and strong support of my professional development, I would like to offer particular gratitude.

NOMENCLATURE

BP	Band-preserving
DED	Discontinuous energy domain
DOF	Degree of freedom
FE(M)	Finite element (method)
k_{eff}	Fundamental-mode criticality eigenvalue
MB	Multiband
MG	Multigroup
MP	Moment-preserving
NKA	Nonlinear Krylov acceleration
ODF	Opacity distribution function
PG-FEMG	Petrov-Galerkin finite element multigroup
PT	Probability table
PWR	Pressured water reactor
QOI	Quantity of interest
SG	Subgroup

TABLE OF CONTENTS

Page

LIST OF FIGURES

FIGURE Page

LIST OF TABLES

TABLE Page

1. INTRODUCTION

1.1 Motivation

Safety, security, and reliability are vital aspects for systems that deal with radiation. Quantitative assessment of these aspects relies on characterizing the radiation: where is it, in what quantities, and in what forms? It has been shown the Boltzmann transport equation is a highly accurate model of how radiation moves through and interacts with a system; however, analytic solutions rarely exist. This puts the onus on using computational science to numerically simulate the radiation. We desire numerical schemes that accurately represent the underlying physics and that are computationally inexpensive.

We investigate the energy variable, which has interesting physics that are difficult to model efficiently. Radiation, be it neutrons, gamma rays, thermal photons, etc., may be treated as individual particles, where each particle has an associated kinetic energy. The interaction of these radiations with their environments, characterized by "cross sections," depends strongly on this energy. Nuclei and atoms possess resonances where the probability of interacting with a nearby radiation particle changes by many orders of magnitude for a small change in incident radiation energy, due to preferred energy levels in the compound nucleus or excited atom / molecule. While these resonances may be well-characterized, they are so numerous that brute-force resolution of the radiation intensity on the scale of the resonances is not a practical option in realistic problems today for deterministic transport codes. Instead, the cross sections are averaged over coarser intervals to create more tractable resolution that is amenable to computation. This coarsening, a type of discretization, carries with it approximations that can cause loss of fidelity.

Resonance shielding effects occur in the context of nuclear reactor simulations. Neutrons are born from fission with high energy and downscatter through the resonance region. As the neutrons stream through a resonant material, absorptions in the resonances cause depressed fluxes near resonance energies; such depressions increase in magnitude as the neutrons stream into a resonant material such as fuel, a phenomenon known as spatial self-shielding. This implies that the energy profile depends upon where the neutrons are and where they came from; the energy-dependent fluxes have spatial and angular dependencies. These effects are important because they influence reaction rates, which determine depletion rates, power profiles, and system criticality.

The multigroup (MG) method is the most employed numerical method for discretizing the energy variable in neutron and photon transport [1, 2]. In a MG solution of the transport equation, the energy variable upon which the solution depends continuously is divided into discrete intervals, with particles "grouped" according to the interval to which their energies belong. Cross sections are condensed via a flux-weighted average over each interval, using some approximate flux. MG is used because it is simple to implement and it can be accurate if the right condensing spectrum is used. It can conserve particles or energy, but not both.

The MG method's chief weakness, which limits its applicability and introduces substantial uncertainties into simulation results, is that its accuracy can be sensitive to an estimated spectral shape that may not be accurate for the problem at hand. When cross sections are averaged over each energy interval, an estimate of the energy spectrum of the flux must be used as the weighting function. MG itself offers no prescription for obtaining this function. Oftentimes, a spatially low-dimensional, high-energy-resolution calculation is done to estimate this shape. In practice, the same spectrum (shape) is used for all directions at a given position and for all posi-

tions in pre-defined spectral regions, which introduces error no matter what spectrum is used.

Spatial self-shielding causes the correct spectral shape to vary significantly across the spatial domain, leading to a loss of important details that may not be captured by the low-order calculation due to its fixed within-group spectral shape in space and angle. This can force the use of a prohibitively large number of groups to ensure accuracy. We seek a new energy discretization scheme that has a solid theoretical foundation, is computationally efficient, and allows the within-group spectral shape of the solution to adapt to local conditions in the current problem.

In this work, we propose, implement, and test a generalized-MG method that uses a novel energy discretization scheme to treat resonances in the cross sections. While we implement and assess the new method in the context of nuclear reactor simulations, we stress our method is sufficiently general as to apply to many photon-transport applications. The following sections characterize past work in the treatment of the energy variable, and introduce our generalized-MG method.

1.2 Past Work

The idea of subdividing an energy or frequency range within a group based on the total cross section or opacity dates back to Chandrasekhar in 1935 [3], when astrophysicists developed the opacity density function (ODF) method. Major contributors included Strom and Kurucz [4], Carbon [5], and Mihalas [6]. The premise of ODF was to turn the within-group opacity into a probability density function and divide the cumulative density function into bands. Particles in the group would have a given probability of interacting with the band and its associated opacity. When viewed in frequency space, the bands that were continuous in opacity space would become split into discontinuous ranges; for this reason the ODF method was also

called "picketing", as in a picket fence.[1] The ODF method allowed the effect of resonances to be incorporated. As temperatures, densities or materials changed, the lines would change, which would alter the ODF and thus the bands. The major drawback of the method was the difficulty and approximation required in figuring out how to take space-dependent band structures into account. Modern work on the ODF method includes the work by Auer and Lowrie [7], Ripoll and Wray [8], and Wray et al. [9], the last of which has some similarities with the method presented in this paper, though both were developed independently.

Since 1935 there has been a proliferation of methods designed to give fidelity within a group. Different authors independently reinvented the ODF idea under new names and applied it to their physics regime. Three such sets of methods are: multi-band (MB), which has seen application to both radiative transfer and neutronics; subgroup (SG), a new name to distinguish neutronics-specific applications of MB; and probability tables (PT), which were first introduced for the unresolved resonance range for neutronics. We briefly discuss each of these variants below.

Cullen [10] claims that Nikolaev [11, 12] and Stewart [13] were the first to propose using the MB framework for neutrons and photons, respectively, in the 1960's; Cullen himself did extensive work analyzing and implementing the method in the 1970's and 1980's [14, 15]. Modern work and implementation of the MB / SG method has been done by Yamamoto and Takeda [16–20], Huang et al. [21], Milošević [22], and Shilkov [23]. The last work is similar to the method presented in this paper, though both were developed independently. Due to its wide range of applicability in the literature and the fundamental similarities of all the methods, we will hereafter use MB to refer to any of the ODF, SG, PT, or MB methods generally.

[1] Most of the frequency range was associated with a low opacity allowing free-streaming — the open area in the fence analogy — while the rest would correspond to the high-opacity lines which had high absorptions — the wood in the fence analogy.

The probability table (PT) method was developed by Levitt [24] for treating neutrons in the unresolved resonance range. In this range, only the distribution of resonance parameters (partial line widths, magnitudes, etc.) are known. The standard technique is to build a realization — called a "ladder" — of the cross section by sampling from the resonance parameters. These cross sections are then collapsed, keeping correlation information among competing reactions, different temperatures and different nuclides. Several realizations are made and averaged together to treat self-shielding. Generally, the narrow-resonance approximation is made to allow analytic flux determinations; this approximation is acceptable in the unresolved resonance range, where resonances widths are small compared to the average energy lost per scatter. Cullen [14] generalized PT to the resolved resonance region. Ribon and Maillard [25] were major contributors to PT theory and introduced a new moments-based approach to determine the band parameters. Modern work has been done on the PT method by Hébert and Martin [26–29]. Since the output of the PT method is tables of correlated parameters that are straightforward to sample from, the probability table method sees substantial use in Monte Carlo calculations to treat the unresolved resonance range. Examples include codes such as MCNP from LANL [30], and RACER and MC21 from KAPL [31, 32]. Traditional implementations of the PT method have had negligible impact on the class of deterministic methods studied here.

MB methods can be grouped into three families: band-preserving (BP), moment-preserving (MP), and discontinuous energy domain (DED). BP MB methods split up each microscopic total cross section into bands such that the total cross section is split up by its magnitude. These bands are used to define nuclide-specific DEDs in the following sense. All energies that have total cross sections which are within a band (cross section range) are placed inside the DED corresponding to that band.

The total and partial cross sections are condensed over each DED to form band-wise microscopic cross sections. As each nuclide formally has its own set of DEDs, resonance interference and correlation issues arise when combining microscopic band cross sections into macroscopic band cross sections. Most early ODF, MB, and PT methods were BP.

MP MB methods use quadrature to preserve moments of the nuclide-specific total (and sometimes partial) cross section(s). The quadrature consists of a set of probabilities and (microscopic) total cross sections for each band. While this method has desirable characteristics with respect to positivity of coefficients, and accuracy and convergence of integration for a single nuclide, it suffers from the same correlation problem as BP MB and further correlation problems because partial cross sections may not be correlated to the total cross section. These correlation problems are overcome for Monte Carlo and collision probability methods by correlation matrices between nuclides / reaction types and conditional band probabilities. Most correlation-preserving techniques have been inapplicable to deterministic methods such as S_n and MOC, where band correlations are often ignored. More recent MB, SG, and PT methods have been MP.

DED MB methods are MG methods that use global DEDs. Such nuclide- and region-consistent DEDs resolve the mapping issues between regions and the correlation issues among nuclides within a region, but must be carefully constructed to preserve resonance features for all relevant nuclides / resonances in the problem. Oftentimes, a problem-wide effective total cross section is constructed to form bands in much the same way as BP MB. The difference is that these bands are not nuclide- or region-specific but global. These bands determine problem-wide DEDs over which to condense the cross sections, which are then used like normal MG cross sections. Our method, [9], and [23] are DED MB methods.

Modern reactor engineering calculations in the United States have used a hierarchical solution methodology. The scheme begins with pointwise cross section data which is processed into many groups with a narrow-resonance (NR) flux: $1/(\sigma_t(E) + \sigma_0)\ 1/E$. This is done for several values of the background cross section, σ_0. On a pin-cell level, cross sections are mixed and NR / IR (intermediate resonance) approximations with equivalence theory is used to determine the proper local background cross sections. One nuclide may use different background cross sections in different regions if it is mixed with different types and amounts of other nuclides in those regions. Cross sections are then collapsed into an intermediate number of groups using a lattice cell calculation, usually with a fixed source and reflective boundaries. This provides fluxes to condense the cross section to few-group values, which are used in a full-core calculation. Traditional references on this method include [1]. Modern references include [33–36].

The problem with the hierarchical approach is that resonance information is averaged out early in the hierarchy, even though the spectral shape of the correct final solution varies strongly with position and direction on the scale of individual resonances. This is especially problematic when spatial homogenization does not accompany energy condensation, because, in this case, cross sections that correspond to material averages are used in the heterogeneous geometry. Additionally, the standard hierarchy has trouble taking into account heterogeneity effects on scales larger than individual pin cells, including the effects of non-similar fuels in nearby pin cells or, for fast systems, resonances in structural materials.

Three novel treatments of the energy variable have been developed in the past 15 years. In 2002, Attieh and Pevey developed a generalized MG method that allowed an energy range to have partial membership in multiple groups using a combination of standard Heaviside functions and linear hat-functions [37, 38]. The motivation

was better energy spectra than piecewise-constant over a group available from MG. The result was a group spectral shape that was problem-adaptive. This method does not solve the energy discretization problem because within-group spectral shapes are not well approximated by polynomials of even moderate order.

Forget and Rahnema developed another generalized-MG method that expanded the neutron flux in terms of an orthogonal basis set [39, 40]. The 0th-order moment was shown to be equivalent to MG, producing the same integral quantities such as eigenvalues and reaction rates. Their technique allowed for estimation of the spectral shape within a group using a low-resolution energy (coarse-group) calculation; however, it suffered from an oscillatory interpolating polynomial that often went negative. Zhu and Forget later generalized this method to use discrete Legrendre orthogonal polynomials, which returned a fine-group flux estimate from a coarse-group calculation [41]. Further work developed the method as an iterative scheme using the unfolded flux from the current iteration to recondense the cross sections for the next iteration [42]. Under the constraints of using step-difference spatial finite difference, storing the angular flux, and computing cell-wise cross sections, the authors showed that their DGM (discrete generalized MG) method converged to the fine-group eigenvector and eigenvalues for several k-eigenvalue problems. More recently, Gibson and Forget have recast their method as producing fine-group accuracy with coarse-group work [35].

Finally, Douglass and Rahnema have a method they call the Subgroup Decomposition Method (SDM) [36]. Though they do not describe it in these terms, SDM is a high-order low-order (HOLO) method where the HO calculation is fine-group MG with a fixed source and the LO calculation is a k-eigenvalue problem with coarse-group MG and a consistency term from the HO calculation. The SDM method shares the same aims as the DGM method.

The DGM and SDM methods offer incremental improvements and do not address the need for improved approximations at the resolution scale of individual resonances to deal with spatial self-shielding. Both methods work within the reactor calculation hierarchy mentioned above and at best allow one level finer of energy resolution.

1.3 New Work

The motivation for this work is to develop a method that can treat the energy variable with high fidelity using substantially fewer energy unknowns ($\mathcal{O}(300)$) than would be required with the typical MG method ($\mathcal{O}(30,000)$ groups). Sharp resonances in the energy dependencies of the cross sections make standard polynomial-based finite-element methods — such as in [35, 37–42] — inadequate to this task. Our method is not a fine-group acceleration scheme (cf. [35, 36]) but a higher-fidelity method for treating the energy variable in transport calculations.

This work offers improvements to the solution of the neutron transport equation by introducing a new formalism for treating the energy variable in the family of MB methods. The proposed Petrov-Galerkin Finite-Element Multigroup Method (PG-FEMG or generalized-MG method) is a generalization of the MG method that uses multiple bases and weight functions within a group. Within this framework, the MG method can be thought of as a Petrov-Galerkin finite-element method (FEM) with one specified flux shape within each energy group; the weight function is unity and the basis function is the assumed flux weighting shape. With PG-FEMG, for each group, discontinuous energy intervals are created by grouping the energy domain according to the magnitude of the total cross section. Orthogonal Heaviside functions are generated that represent band membership. The weight (test) functions are these Heaviside functions; the basis (trial) functions are the product of these Heaviside functions and a spatial-region-specific normalized weighting flux. The full finite

element (FE) definition is discussed in Section 3.

For problems with more than one material or temperature, a generalization is employed to determine the global discontinuous energy intervals. An effective total cross section is used, where the total cross section of each user-defined spatial region is weighted with a flux parameter. This parameter may come from analytic expressions or a simplified transport calculation. In this work, we use region to mean a spatially contiguous set of cells that have approximately the same material composition. For the sake of this work, being composed of different materials or having different temperatures would differentiate regions.

The PG-FEMG method shares much with MB methods but offers important innovations. First and foremost, PG-FEMG is a FEM, whereas historically many MB implementations used band weights and probabilities that preserved various (flux-weighted) moments of the cross section within a group. While the primary strength of MB methods is that these moments capture the effect of both the background cross section and resonances of individual nuclides, the primary weakness of MB methods is that there are problems for which the true flux cannot be represented as a function of a single nuclide's microscopic total cross section, e.g., resonance interference effects.

In contrast, the PG-FEMG method divides the energy domain within a group into discontinuous intervals based on the magnitude of a problem-wide effective total cross section and uses a flux-weighted cross section in each of these intervals. Much like historical MB methods, this adds fidelity to the representation of the resonances themselves, as some intervals contain just the energy ranges corresponding to resonances. An advantage of PG-FEMG is that, as the number of bands (intervals) increases, the cross section variance within a band rapidly decreases. In this limit, the shape of the weighting spectrum is much less important than in the MG method.

Another distinction is the use of a global band structure. The majority of existing MB methods use local band parameters determined by the resonances of the local nuclide or group of nuclides. This introduces correlation issues, some of which were mentioned previously. Boundaries between two regions with different band definitions should involve a within-group band remapping, though this is rarely done. When in the presence of temperature or concentration gradients, rigorous treatment of the changing band structure would create additional complexity in the streaming term, though this too is often ignored. Resolving such difficulties would require that all the bands within a group be transported together so that no lagged information is used in the mapping or derivative terms. Resolving correlation issues requires the rewriting of existing MG codes. In contrast, using a global band structure allows the PG-FEMG method to use existing MG codes by adding only pre- and post-processing steps.

A final distinction is in the representation of the scattering transfer matrix. Both traditional MB and PG-FEMG methods are similar in that the non-differential cross sections need to be stored for each band in each group. Most MB methods have treated the scattering transfer probability function with low fidelity, if at all.[2] They may store a band-to-band scattering matrix for each group to treat within-group scatter but store only a group-to-group scattering matrix for across-group scatter. Any correlations between bands in disparate groups are lost and sometimes correlations among bands within a group are also lost. In contrast, the proposed method uses the complete band-to-band scattering matrix, both within- and across-group.

The proposed method has been implemented in a research code by altering the cross-section preparation step wherein the pointwise cross sections are collapsed into MG-like structures. This method preserves the format of the output cross sections —

[2]Many MB methods have been applied to radiative transfer problems with the scattering ignored.

one-dimensional cross sections and transfer matrices — but not the interpretation. A band/group index refers to a discontinuous set of energies instead of a continuous one. Otherwise, the same pre-processing is used, and the resulting cross sections can be used in existing MG codes. The output requires a quick post-processing step to collapse the bands into coarse groups or to expand the bands into their continuous basis function representations, but otherwise it fits into the standard calculational framework.

The method creates within-group scattering that looks like upscattering. The PG-FEMG scattering matrix in energy has upper-diagonal terms caused by within-group, across-band scattering. MG scattering energy matrices only contain upper-diagonal terms in the presence of upscattering, which does not occur in the resolved resonance regions. We refer to this phenomenon as "effective upscattering." The presence of the effective upscattering may necessitate the use of improved iterative algorithms, as MG solvers often assume upscattering only in the thermal groups. We use nonlinear Krylov acceleration (NKA) for this purpose, described in later sections.

Analysis was done on several reactor-relevant pin-cell problems that provided a rigorous basis for evaluating the efficacy and efficiency of the method. A series of tests of increasing difficulty were implemented on SCDT, a new one-dimensional, cyclindrical, S_n research code constructed for this work. The following models were evaluated:

1. A single-temperature, homogeneous mixture of low-enriched fuel (LEF) and water with white boundary conditions, to test energy fidelity decoupled from spatial and angular effects;

2. A typical PWR fuel pin of one-temperature LEF surrounded by water, to add spatial self-shielding, the phenomenon of interest;

3. A typical PWR fuel pin of multiple-temperature LEF surrounded by water, to broaden (but not shift) the resonances in a spatially-dependent manner;

4. A typical PWR fuel pin of one-temperature mixed-oxide fuel (MOX) surrounded by one-temperature low-enriched uranium (LEU) fuel surrounded by water, to add multiple important resonant nuclides in a heterogeneous arrangement; and

5. A cylindricized fuel lattice consisting of a typical PWR fuel pin of one-temperature MOX and water surrounded by fuel pins of LEU and water, to add additional spatial variation to the resonance interactions.

For each problem, a k-eigenvalue calculation was performed. The criticality eigenvalue and normalized power distributions were compared among standard MG, PG-FEMG with the same number of DOF, and a very-fine-group reference MG calculation for the benchmark.

This work required the creation and updating of several codes, including the debugging of a cross section converter code from NJOY (MATXS) output to SCDT or PDT [43] input, the creation of a condenser code to go from very-fine-group MG to coarser-group MG or coarser-group generalized-MG, the creation of a one-dimensional cylindrical S_n research code that solves the k-eigenvalue problem with NKA (SCDT), and a pre- / post-processing script framework that coordinates the inputs and outputs of these codes and performs analysis of the results. The end result of this work has been a clearer understanding of the strengths and weaknesses of the PG-FEMG method.

This work has inspired ideas for future research. The proposed PG-FEMG method only adds or modifies three stages in the journey from evaluated nuclear data to quantity of interest (QOI): definition of discontinuous energy domains from

a band boundary calculation, cross section condensation over discontinuous (previously continuous) energy domains, and post-processing. Many generalizations to the PG-FEMG method can be envisaged, such as using estimates of the solution to determine the band structure, calculating the band structure using a union of the band structures over all spatial regions, using ideas from the probability table method to generalize this method to the unresolved resonance range, and defining a band structure to meet a desired solution accuracy. Such ideas are beyond the scope of this work.

This work is laid out in the following sections. Section 2 gives an overview of our mathematical model, the Boltzmann transport equation, and presents the criticality-eigenvalue formulation, restrictions to cylindrical geometries, discretization in angle, and discretization in space. Section 3 presents the theoretical basis for the generalized-MG method by discussing the discretization in energy. Section 4 presents an overview of cross section preparation and of the workflow to implement and test the generalized-MG method. Section 5 discusses solution techniques for solving the discretized version of the transport equation, including discussion of NKA and its application. Section 6 provides an overview of the problems run for this work and of the verification problems performed. Section 7 presents results for the problems discussed in Section 6. Section 8 provides a discussion of the results, gives conclusions, and discusses possible future work. Appendices fill out the work by providing addenda to the sections: Appendix A corresponds to Section 2, Appendix B corresponds to Section 5, and Appendix C corresponds to Section 7.

2. THE TRANSPORT EQUATION

The physical model underlying the work done in this thesis is the transport equation. Sometimes referred to as the linearized Boltzmann transport equation, it describes the movement of neutral particles under non-equilibrium conditions and is one of the workhorse equations of the nuclear engineering discipline. This section describes the equation in its general form, gives some simplifications made, and then explains the discretizations made that allow the equation to be solved on modern computer architectures.

2.1 The Continuous Transport Equation

The Boltzmann transport equation, originally developed to model rarefied one-species gases, describes the density of neutral particles in a 6-dimensional phase-space of position (3-dimensional $\mathbf{r}$), direction (2-dimensional $\mathbf{\Omega}$) and energy (1-dimensional E), throughout time (t). The equation for neutrons is given by

$$\frac{1}{v(E)}\frac{\partial}{\partial t}\psi(\mathbf{r},E,\mathbf{\Omega},t)+\mathbf{\Omega}\cdot\boldsymbol{\nabla}\psi(\mathbf{r},E,\mathbf{\Omega},t)+\Sigma_t(\mathbf{r},E,t)\psi(\mathbf{r},E,\mathbf{\Omega},t)$$
$$=S\psi(\mathbf{r},E,\mathbf{\Omega},t)+F\psi(\mathbf{r},E,\mathbf{\Omega},t)+q_{\text{ext}}, \tag{2.1a}$$

$$S\psi(\mathbf{r},E,\mathbf{\Omega},t)=\int_0^\infty \mathrm{d}E' \int\limits_{4\pi} \mathrm{d}\mathbf{\Omega}'\ \ \Sigma_s(\mathbf{r},E'\rightarrow E,\mathbf{\Omega}\cdot\mathbf{\Omega}',t)\ \psi(\mathbf{r},E',\mathbf{\Omega}',t), \tag{2.1b}$$

$$F\psi(\mathbf{r},E,\mathbf{\Omega},t)=\int_0^\infty \mathrm{d}E'\ \ \frac{\chi(\mathbf{r},E'\rightarrow E,t)}{4\pi}\nu\Sigma_f(\mathbf{r},E',t)\int\limits_{4\pi}\mathrm{d}\mathbf{\Omega}'\ \ \psi(\mathbf{r},E',\mathbf{\Omega}',t), \tag{2.1c}$$

$$\psi(\mathbf{r},E,\mathbf{\Omega},t)=v(E)\,n(\mathbf{r},E,\mathbf{\Omega},t), \tag{2.1d}$$

$$\psi(\mathbf{r},E,\mathbf{\Omega},t)=f_{\text{in}}(\mathbf{r},E,\mathbf{\Omega},t)\qquad \mathbf{\Omega}\cdot\hat{\mathbf{n}}<0, \tag{2.1e}$$

$$\psi(\mathbf{r},E,\mathbf{\Omega},0)=f_0(\mathbf{r},E,\mathbf{\Omega}), \tag{2.1f}$$

where $\mathbf{r}$ is the position in Cartesian space (cm), $\mathbf{\Omega}$ is the direction of travel in Cartesian space normalized to a magnitude of unity, E is the lab-frame kinetic energy of the neutron (MeV), ψ is the angular flux (1/cm^2-s-MeV-ster), n is the neutron density (1/cm^3-MeV-ster), $v(E)$ is the velocity (cm/s), $\mathbf{\Omega}\cdot\boldsymbol{\nabla}$ is the streaming operator (1/cm), Σ_t is the total macroscopic cross section (1/cm), Σ_s is the doubly-differential macroscopic scattering cross section (1/cm-MeV-ster), $\nu\Sigma_f$ is the macroscopic fission-neutron production cross section (1/cm), $\chi/4\pi$ is the distribution of fission neutrons in energy and direction (1/MeV-ster), and q_{ext} is the doubly-differential extraneous

neutron source ($1/\text{cm}^3$-s-MeV-ster). Equations (2.1e) and (2.1f) are the boundary and initial conditions, respectively. The boundary conditions may be homogeneous (i.e., strictly linear in ψ) or not.

Equation (2.1) models neutrons in the absence of external fields (e.g., gravity), without spontaneous neutron decay, for isotropic materials (i.e., where scattering depends only $\mathbf{\Omega}' \cdot \mathbf{\Omega}$), neglecting neutron-neutron interactions and three-plus-body interactions. It describes how neutrons move in phase space. When integrated out over all phase space, Eq. (2.1) yields a balance equation for the conservation of neutrons.

Mathematically, Eq. (2.1) is well-posed. For positive cross sections, it can be shown that the solution to the equation is unique and exists. In a void, the system becomes hyperbolic, having a finite rate at which information travels through the system. For systems where the scattering plus fission sources, $S\psi + F\psi$, are always less than or equal to the total cross removal rate, $\Sigma_t\psi$, it can be shown that a steady-state solution exists with positive ψ. Finally, Eq. (2.1) is a linear system, which naturally lends the system to eigenvalue evaluation.

The fission source, Eq. (2.1c), is commonly simplified under the assumption that the fission spectrum is constant for all incident neutron energies. This is generally true except for high-energy ($\sim 1+$ MeV) incident neutrons. Since this work focuses on the resolved resonance region (< 1 MeV neutrons), we make this assumption, which yields:

$$F\psi(\mathbf{r}, E, \mathbf{\Omega}, t) = \frac{\chi(\mathbf{r}, E, t)}{4\pi} \int_0^\infty \mathrm{d}E' \;\; \nu\Sigma_f(\mathbf{r}, E', t) \int_{4\pi} \mathrm{d}\mathbf{\Omega}' \;\; \psi(\mathbf{r}, E', \mathbf{\Omega}', t). \tag{2.2}$$

2.2 The k-Eigenvalue Problem

Equation (2.1) may be cast as a steady-state, equivalent-critical eigenvalue problem. In this formulation, all non-homogeneous terms are set to zero and the fission source is scaled by eigenvalue k.

$$\begin{aligned} \boldsymbol{\Omega} \cdot \boldsymbol{\nabla} \psi(\mathbf{r}, E, \boldsymbol{\Omega}) + \Sigma_t(\mathbf{r}, E) \psi(\mathbf{r}, E, \boldsymbol{\Omega}) & \\ = S\psi(\mathbf{r}, E, \boldsymbol{\Omega}) + \frac{1}{k} F\psi(\mathbf{r}, E, \boldsymbol{\Omega}) & \quad (2.3a) \\ \psi(\mathbf{r}, E, \boldsymbol{\Omega}) = f(\psi, \mathbf{r}, E, \boldsymbol{\Omega}) \qquad \boldsymbol{\Omega} \cdot \hat{\mathbf{n}} < 0 & \quad (2.3b) \end{aligned}$$

where f is a linear mapping on ψ.

Though there are an infinite number of k and ψ that solve Eq. (2.3), generally it is desired to investigate the k that is largest in magnitude. It can be shown that this fundamental mode produces positive, real k and ψ [1, 2]. This k is referred to as k_{eff} and is the reactivity of the system.

The quantity k_{eff} determines the state of the system. For $k_{\text{eff}} < 1.0$, a system is called sub-critical; in the absence of an external source, the neutron population of such a system would decrease with time. For $k_{\text{eff}} > 1.0$, a system is called super-critical; if not identically zero, the neutron population of such a system would increase exponentially with time. For $k_{\text{eff}} = 1.0$, a system is called critical; in the absence of an external source, the neutron population of such a system would be steady-state (constant) and in the fundamental mode ψ.

Accurate determination of k_{eff} has important safety implications. Nuclear reactors need to maintain conditions of near-criticality for their entire operating cycle. Fresh nuclear reactor fuels need to have excess criticality to overcome the deleterious effects of fissile fuel burnup and neutron poison buildup. Materials with nonzero

fission cross sections should be stored in orientations that do not allow unintentional criticality. Failure to meet these criteria courts disaster.[1]

Computationally, calculation of k_{eff} is a common requirement. Not only does k_{eff} represent an important state of the system, it represents an integral quantity. Such quantities are less susceptible to discretization errors or input uncertainties. The solution technique is described in later sections.

2.3 Restriction to Cylindrically-Symmetric Geometries

The scope of the thesis included development of a one-dimensional transport code that could be used to simulate a periodic nuclear-reactor lattice. As fuel layout in modern US light water reactors (LWR's) consists of cylindrical fuel pins arrayed in a repeated grid, a cylindrically-symmetric geometry was chosen with reflective (white) boundary conditions to account for the grid.[2]

For cylindrical symmetry, the spatial domain is cast in cylindrical coordinates (r, z, θ) and the system is assumed to have dependence only on r. In this coordinate system, $\mathbf{\Omega}$ is defined such that the polar axis is the z axis and the $\omega = 0$ axis is the r axis. In the process of mapping $\mathbf{\Omega}$ onto a cylindrical coordinate system, the streaming operator picks up an angular derivative.

In curvilinear coordinates, $\mathbf{\Omega} \cdot \boldsymbol{\nabla}$ involves derivatives in both space (physical space) and direction (momentum space). As a particle streams, the unit vector from the axis of the cylinder to the location of the particle (the radial vector) changes as viewed in Cartesian space. Unless the particle is streaming directly toward or away from the polar axis, the azimuthal angle of the particle with respect to the radial vector changes as the particle streams. This gives rise to an angular streaming

[1] Cf. e.g., Godiva I.

[2] The effective angles of reflection off the boundary are wrong when a cylindrical approximation is made to model a square lattice [2]. White boundaries are used to average out these specular reflection-based errors.

derivative.

The only change to the transport equation for curvilinear coordinates comes in the $\nabla \cdot \mathbf{\Omega}$ term, which in "conservative form" is written as:

$$\mathbf{\Omega} \cdot \boldsymbol{\nabla} \psi = \frac{\mu}{r} \frac{\partial}{\partial r} \big(r \, \psi(r, \mathbf{\Omega}, E) \big) - \frac{1}{r} \frac{\partial}{\partial \omega} \big(\eta \, \psi(r, \mathbf{\Omega}, E) \big), \tag{2.4}$$

where r is the radial coordinate, ω is the azimuthal angle, and mu and η are angle cosines. Appendix A contains further details for the interested reader.

2.4 Discretization in Angle

A triangular Legendre-Chebyshev quadrature was used in angle. That is, Gauss-Legendre points and weights were used on the polar axis, and Gauss-Chebyshev points and weights were used on the azimuthal axis. This quadrature set is consistent with and commonly used for cylindrical geometries. Appendix A contains further details for the interested reader.

2.5 Discretization in Space

A linear discontinuous finite element method (DFEM) was used to model the spatial dependence within a cell. Since the method is globally discontinuous, it falls under the non-conforming FE (similar to finite volume) category. In 1D, LD has two degrees of freedom (DOF) per cell. Therefore, a 2x2 system needs to be solved per cell. DFEMs have become relatively standard spatial discretizations for the transport equation, the details of which are reserved for Appendix A.

2.6 Iterative Acceleration

A P_1-equivalent S_2 synthetic acceleration (S2SA) method was used to allow diffusive problems to be run. For the infinite-medium case, the factor by which the error in the scalar flux is reduced (the spectral radius) by source iteration alone is bounded

above by the scattering ratio ($c \equiv \sigma_{s,0,g\to g}/\sigma_{t,g}$) [2]. For optically thick regions with c arbitrarily close to unity, this means the method can and usually does converge arbitrarily slowly. It can be shown that the slowly-converging modes are those with slowly varying angular dependence.

To remedy this handicap, the one-group S2SA method is used. Optically thick problems with large scattering ratios are known to be diffusive and linearly anisotropic in angle (P_1). Theory says that accelerating with a P_1-equivalent method will eliminate the slowly varying angular error modes quickly while leaving the more rapidly-varying error modes alone, where source iteration performs well.

S2SA consists of solving the one-group, S_2-like system simultaneously for both angles. Since a matrix inversion and not source iteration is used, the convergence problem is circumvented. S2SA is a standard acceleration technique for one-dimensional transport equations; a complete description is left to Appendix A.

Appendix A contains further details on topics related to spatial and angular discretization for the interested reader.

3. DISCRETIZATIONS IN ENERGY

We have taken our general transport equation, restricted it to one-dimensional cylindrical geometries, formulated it in terms of a criticality eigenvalue, and discretized it in space and angle. We now approach the topic of this work and the one remaining variable to be discretized: the energy variable.

3.1 Definition of the Finite Element Method

Because the band structure is space- and angle-independent, without loss of generality we begin with the continuous k-eigenvalue formulation of the transport equation for neutronics with anisotropic scattering and white boundary conditions:

$$\begin{aligned}
\boldsymbol{\Omega}\cdot\nabla\psi(\mathbf{r},E,\boldsymbol{\Omega})+\Sigma_t(\mathbf{r},E)\psi(\mathbf{r},E,\boldsymbol{\Omega}) = & \\
\int_{4\pi}\mathrm{d}\Omega'\int_0^\infty \mathrm{d}E'\,\Sigma_s(\mathbf{r},E'\to E,\boldsymbol{\Omega}'\cdot\boldsymbol{\Omega})\,\psi(\mathbf{r},E',\boldsymbol{\Omega}') + & \\
\frac{\chi(\mathbf{r},E)}{4\pi\,k_{\text{eff}}}\int_0^\infty \mathrm{d}E'\,\nu\Sigma_f(\mathbf{r},E')\phi(\mathbf{r},E'), & \qquad (3.1a)\\
\phi(\mathbf{r},E)=\int_{4\pi}\mathrm{d}\Omega\;\psi(\mathbf{r},E,\boldsymbol{\Omega}), & \qquad (3.1b)\\
\psi(\mathbf{r},E,\boldsymbol{\Omega})=\frac{1}{\pi}\int_{\boldsymbol{\Omega}\cdot\hat{\mathbf{n}}>0}\mathrm{d}\Omega'\;\;|\boldsymbol{\Omega}'\cdot\hat{\mathbf{n}}|\,\psi(\mathbf{r},E,\boldsymbol{\Omega}') & \\
r\in\partial V,\;\boldsymbol{\Omega}\cdot\hat{\mathbf{n}}(\mathbf{r})<0. & \qquad (3.1c)
\end{aligned}$$

Standard MG splits the energy domain into non-overlaping regions called groups and assumes a constant spectral shape in each group. This is accomplished by using piecewise-constant weight functions that are unity in a group and zero outside it. Basis functions are equal to weight functions multiplied by the spectral shape within

a group. We have:

$$w_g(E) = \begin{cases} 1 & \text{if } E \in \Delta E_g, \\ 0 & \text{otherwise,} \end{cases} \tag{3.2a}$$

$$b_g(\mathbf{r}, E) = \begin{cases} f_g(\mathbf{r}, E) & \text{if } E \in \Delta E_g, \\ 0 & \text{otherwise.} \end{cases} \tag{3.2b}$$

PG-FEMG splits the energy domain into non-overlapping intervals (groups); however, it allows for within-group variation by using multiple basis functions within each group. The weight function is a band membership function, realized by a Heaviside function that is unity when the energy is in the energy range of the group and when the total cross section is within the total cross section of the band, and zero otherwise. Figure 3.3 has weight functions for an example three-band group. The basis functions are equal to the weight functions multiplied by the spectral shape within a band. We have:

$$w_{g,b}(E) = \begin{cases} 1 & \text{if } E \in \Delta E_{g,b}, \\ 0 & \text{otherwise,} \end{cases} \tag{3.3a}$$

$$b_{g,b}(\mathbf{r}, E) = \begin{cases} f_{g,b}(\mathbf{r}, E) & \text{if } E \in \Delta E_{g,b}, \\ 0 & \text{otherwise,} \end{cases} \tag{3.3b}$$

with

$$\Delta E_{g,b} \equiv \left\{ E \,\middle|\, E \in \Delta E_g \cap \Sigma_t(E) \in \Delta \Sigma_{t,g,b} \right\}. \tag{3.4}$$

Note that each $\Delta E_{g,b}$ corresponds to a unique, possibly discontinuous, set of energies and is orthogonal with respect to $\{\Delta E_{g,b}\}$. The union of the $\{\Delta E_{g,b}\}$ spans the

full energy domain. We discuss the $\{\Delta\Sigma_{t,g,b}\}$ later. In this framework, MG uses $\Delta E_{g,b} = \Delta E_g$, that is, one band per group.

It will be later shown that $\int_{\Delta E_{g,b}} \mathrm{d}E\ f_{g,b}(\mathbf{r}, E) = 1$ is a convenient normalization choice for Eq. (3.3b), and so we adopt this normalization.

3.2 Derivation of the PG-FEMG Transport Equation

We expand the angular and scalar fluxes using:

$$\psi(\mathbf{r}, E, \mathbf{\Omega}) \rightarrow \sum_{g,b} b_{g,b}(\mathbf{r}, E)\psi_{g,b}(\mathbf{r}, \mathbf{\Omega}), \tag{3.5a}$$

$$\Rightarrow \psi(\mathbf{r}, E, \mathbf{\Omega}) = f_{g,b}(\mathbf{r}, E)\,\psi_{g,b}(\mathbf{r}, \mathbf{\Omega}), \qquad E \in \Delta E_{g,b}, \tag{3.5b}$$

with $w_{g,b}(E)$ a weight function with local support defined above. Note that $w_{g,b}(E)$ is spatially global — the same $\{\Delta E_{g,b}\}$ are used throughout the domain. Equation (3.5a) is the *only* approximation that PG-FEMG makes.

To get the MG or PG-FEMG forms of the transport equation, simply integrate

over energy with the weight function and apply the basis function definitions:

$$\int_0^\infty \mathrm{d}E\; w_{k,n}(E)\Bigg\{\boldsymbol{\Omega}\cdot\nabla\left[\sum_{g,b} b_{g,b}(\mathbf{r},E)\psi_{g,b}(\mathbf{r},\boldsymbol{\Omega})\right] + \Sigma_t(\mathbf{r},E)\sum_{g,b} b_{g,b}(\mathbf{r},E)\psi_{g,b}(\mathbf{r},\boldsymbol{\Omega}) =$$
$$\sum_{l=0}^{L}\frac{2l+1}{4\pi}\int_0^\infty \mathrm{d}E'\;\Sigma_{s,l}(\mathbf{r},E'\to E)\sum_{m=-l}^{m} Y_{l,m}(\boldsymbol{\Omega})\sum_{g',b'} b_{g',b'}(\mathbf{r},E')\phi^{l,m}_{g',b'}(\mathbf{r}) +$$
$$\frac{\chi(\mathbf{r},E)}{4\pi\,k_{\mathrm{eff}}}\int_0^\infty \mathrm{d}E'\;\nu\Sigma_f(\mathbf{r},E')\sum_{g',b'} b_{g',b'}(\mathbf{r},E')\phi_{g',b'}(\mathbf{r})\Bigg\}, \tag{3.6a}$$

$$\phi_{g,b}(\mathbf{r}) = \int_{4\pi}\mathrm{d}\Omega\;\;\psi_{g,b}(\mathbf{r},\boldsymbol{\Omega}), \tag{3.6b}$$

$$\psi_{g,b}(\mathbf{r},\boldsymbol{\Omega}) = \frac{1}{\pi}\int_{\boldsymbol{\Omega}\cdot\hat{\mathbf{n}}>0}\mathrm{d}\Omega'\;\;|\boldsymbol{\Omega}'\cdot\hat{\mathbf{n}}|\,\psi_{g,b}(\mathbf{r},\boldsymbol{\Omega}')$$
$$r\in\partial V,\;\boldsymbol{\Omega}\cdot\hat{\mathbf{n}}(\mathbf{r})<0. \tag{3.6c}$$

Equation (3.6) uses the spherical harmonic addition theorem to expand the scattering source in terms of spherical harmonics, with definitions:

$$\phi^{l,m}_{g,b}(\mathbf{r}) \equiv \int_{4\pi}\mathrm{d}\Omega\;\;Y_{l,m}(\boldsymbol{\Omega})\,\psi_{g,b}(\mathbf{r},\boldsymbol{\Omega}), \tag{3.7a}$$

$$\Sigma_{s,l}(\mathbf{r},E'\to E) \equiv 2\pi\int_{-1}^{+1}\mathrm{d}\xi\;\;P_l(\xi)\,\Sigma_s(\mathbf{r},E'\to E,\xi), \tag{3.7b}$$

with $Y_{l,m}(\boldsymbol{\Omega})$ the real-valued spherical harmonic of degree l and order m, and $P_l(\xi)$ the Legendre polynomial of degree l. Most cross-section codes return $\sigma_{s,l}(E'\to E)$, not $\sigma_s(E'\to E,\boldsymbol{\Omega}'\cdot\boldsymbol{\Omega})$.

We re-arrange Eq. (3.6):

$$\begin{aligned}
&\boldsymbol{\Omega}\cdot\nabla\sum_{g,b}\psi_{g,b}(\mathbf{r},\boldsymbol{\Omega})\left[\int_0^\infty \mathrm{d}E\; w_{k,n}(E)\, b_{g,b}(\mathbf{r},E)\right]+\\
&\quad\sum_{g,b}\psi_{g,b}(\mathbf{r},\boldsymbol{\Omega})\left[\int_0^\infty \mathrm{d}E\; w_{k,n}(E)\Sigma_t(\mathbf{r},E)b_{g,b}(\mathbf{r},E)\right]=\\
&\qquad\sum_{l=0}^{L}\frac{2l+1}{4\pi}\sum_{m=-l}^{l}Y_{l,m}(\boldsymbol{\Omega})\sum_{g',b'}\phi^{l,m}_{g',b'}(\mathbf{r})\Bigg\{\\
&\qquad\quad\left[\int_0^\infty \mathrm{d}E'\; b_{g',b'}(\mathbf{r},E')\int_0^\infty \mathrm{d}E\; w_{k,n}(E)\Sigma_{s,l}(\mathbf{r},E'\to E)\right]\Bigg\}+\\
&\quad\frac{1}{4\pi\, k_{\mathrm{eff}}}\sum_{g',b'}\phi_{g',b'}(\mathbf{r})\left[\int_0^\infty \mathrm{d}E\; w_{k,n}(E)\chi(\mathbf{r},E)\right]\left[\int_0^\infty \mathrm{d}E'\;\nu\Sigma_f(\mathbf{r},E')b_{g',b'}(\mathbf{r},E')\right].
\end{aligned}\tag{3.8}$$

We use orthogonality:

$$\int_0^\infty \mathrm{d}E\; w_{g,b}(E)\, b_{g',b'}(\mathbf{r},E)=\delta_{g,g'}\delta_{b,b'}\int\limits_{\Delta E_{g,b}}\mathrm{d}E\; f_{g,b}(\mathbf{r},E)=\delta_{g,g'}\,\delta_{b,b'},\tag{3.9a}$$

to define spectral averages:

$$\begin{aligned}
\int_0^\infty \mathrm{d}E\; w_{g,b}(E)\, b_{g',b'}(\mathbf{r},E)\,\Sigma(\mathbf{r},E)&=\delta_{g,g'}\,\delta_{b,b'}\int\limits_{\Delta E_g}\mathrm{d}E\; f_g(\mathbf{r},E)\,\Sigma(\mathbf{r},E)\\
&=\delta_{g,g'}\,\delta_{b,b'}\,\Sigma_{g,b}(\mathbf{r}),\qquad(3.9b)\\
\Sigma_{g,b}(\mathbf{r})&\equiv\int_0^\infty \mathrm{d}E\;\; b_{g,b}(\mathbf{r},E)\,\Sigma(\mathbf{r},E)\\
&=\int\limits_{\Delta E_{g,b}}\mathrm{d}E\; f_{g,b}(\mathbf{r},E)\,\Sigma(\mathbf{r},E),\qquad(3.9c)
\end{aligned}$$

and spectral sums:

$$\chi_{g,b}(\mathbf{r}) \equiv \int_0^\infty \mathrm{d}E\; w_{g,b}(E)\,\chi(\mathbf{r},E) = \int\limits_{\Delta E_{g,b}} \mathrm{d}E\;\chi(\mathbf{r},E). \tag{3.9d}$$

Specifically, we define

$$\Sigma_{t,g,b}(\mathbf{r}) \equiv \int\limits_{\Delta E_{g,b}} \mathrm{d}E\; f_{g,b}(\mathbf{r},E)\,\Sigma_t(\mathbf{r},E), \tag{3.10a}$$

$$\chi_{g,b}(\mathbf{r}) = \int\limits_{\Delta E_{g,b}} \mathrm{d}E\;\chi(\mathbf{r},E), \tag{3.10b}$$

$$\nu\Sigma_{f,g',b'}(\mathbf{r}) \equiv \int\limits_{\Delta E_{g',b'}} \mathrm{d}E'\; f_{g',b'}(\mathbf{r},E')\,\nu\Sigma_f(\mathbf{r},E'), \tag{3.10c}$$

$$\Sigma_{s,l,(g',b')\to(g,b)}(\mathbf{r}) \equiv \int\limits_{\Delta E_{g',b'}} \mathrm{d}E'\; f_{g',b'}(\mathbf{r},E') \int\limits_{\Delta E_{g,b}} \mathrm{d}E\;\Sigma_{s,l}(\mathbf{r},E'\to E). \tag{3.10d}$$

Note that the choice of basis functions in energy for $\psi(\mathbf{r},E,\mathbf{\Omega})$ determines the MG or PG-FEMG cross sections. It is important to be consistent with these bases: they are used both for condensation of the cross sections and for expansion of the flux.

Because the group and band structures are spatially global, there are no spatial dependencies on the $\Delta E_{g,b}$ and hence the $w_{g,b}(E)$. Contrast this with standard (BP) MB, where different regions in space have different $\Delta E_{g,b}$ and hence different $w_{g,b}(\mathbf{r},E)$, which necessitates mapping bands between regions. With PG-FEMG, no mapping is needed between regions. This allows existing MG codes to use PG-FEMG cross sections.

Correlation is automatically enforced by using consistent energy domains for the entire problem. Contrast this with (MP) MB, which requires extra machinery and data to enforce band-to-band correlation among disparate nuclides, temperatures,

and spatial regions.

Recall that our spectral shapes, $f_{g,b}(E)$, are pdfs (normalized to unity):

$$\int_{\Delta E_{g,b}} \mathrm{d}E \; f_{g,b}(\mathbf{r}, E) = 1, \quad \forall\, \mathbf{r}. \tag{3.11}$$

This simplifies the streaming term, as the integral over $\Delta E_{g,b}$ removes the spatial dependence of $f_{g,b}(\mathbf{r}, E)$.

Also note that orthogonality implies:

$$\psi_{g,b}(\mathbf{r}, \mathbf{\Omega}) = \int_0^\infty \mathrm{d}E \;\; w_{g,b}(E)\, \psi(\mathbf{r}, E, \mathbf{\Omega}) = \int_{\Delta E_{g,b}} \mathrm{d}E \;\; \psi(\mathbf{r}, E, \mathbf{\Omega}). \tag{3.12}$$

Equation (3.12) gives us the physical interpretation that the $\psi_{g,b}$ are simply the integrals of the angular flux over $\Delta E_{g,b}$.

Upon application of the orthogonality, spectral averages, and spectrum sums, we find the k-eigenvalue formulation of the MG or PG-FEMG transport equation for neutrons with isotropic scattering and white boundary condition can be written as:

$$\begin{aligned}
\mathbf{\Omega} \cdot \nabla \psi_{g,b}(\mathbf{r}, \mathbf{\Omega}) + \Sigma_{t,g,b}(\mathbf{r}) \psi_{g,b}(\mathbf{r}, \mathbf{\Omega}) = & \\
& \sum_{l=0}^{L} \frac{2l+1}{4\pi} \sum_{g',b'} \Sigma_{s,l,(g',b')\to(g,b)}(\mathbf{r}) \sum_{m=-l}^{l} \phi^{l,m}_{g',b'}(\mathbf{r}) Y_{l,m}(\mathbf{\Omega}) + \\
& \quad \frac{\chi_{g,b}(\mathbf{r})}{4\pi\, k_{\text{eff}}} \sum_{g',b'} \nu\Sigma_{f,g',b'}(\mathbf{r}) \phi_{g',b'}(\mathbf{r}), && \text{(3.13a)} \\
\phi_{g,b}(\mathbf{r}) = & \int_{4\pi} \mathrm{d}\Omega \;\; \psi_{g,b}(\mathbf{r}, \mathbf{\Omega}), && \text{(3.13b)} \\
\psi_{g,b}(\mathbf{r}, \mathbf{\Omega}) = & \frac{1}{\pi} \int_{\mathbf{\Omega}\cdot\hat{\mathbf{n}}>0} \mathrm{d}\Omega' \;\; |\mathbf{\Omega}' \cdot \hat{\mathbf{n}}|\, \psi_{g,b}(\mathbf{r}, \mathbf{\Omega}') \\
& \qquad r \in \partial V, \;\; \mathbf{\Omega} \cdot \hat{\mathbf{n}}(\mathbf{r}) < 0. && \text{(3.13c)}
\end{aligned}$$

for $\psi_{g,b}$ the PG-FEMG angular flux, $\phi_{g,b}$ the scalar flux, $\mathbf{r}$ the spatial coorinate, $\mathbf{\Omega}$ the direction of travel, g the energy group, b the energy band, $\Sigma_{t,g,b}$ the local macroscopic total cross section, $\Sigma_{s,(g',b')\to(g,b)}$ the scattering transfer macroscopic cross section, $\chi_{g,b}$ the fission distribution function (in energy), $\nu\Sigma_{f,g,b}$ the neutron production from fission macroscopic cross section, and k_{eff} the eigenvalue. It is trivial to adapt Eq. (3.13) to use a fission matrix.

If in Eq. (3.13) we replace all (g, b) with g, and likewise with primes, then we obtain the standard MG transport equation. In Eq. (3.13), the integrals to perform the averages and sums are over discontinuous energy regions, while standard MG has them over continuous ranges. This is the only difference. We can use Eq. (3.13) to solve both the MG and PG-FEMG systems. This is one of the strengths of the method: the solve step can be carried out with existing MG codes.

The problems in this thesis used isotropic scattering kernels and sums of dyadic fission kernels for the scattering and fission sources. Each component (e.g., Pu-239, U-235) was assumed to have a dyadic fission matrix, but the macroscopic fission matrix was not assumed dyadic (but rather a sum of dyads). The fission rate is given as

$$R_f(\mathbf{r}, E, \mathbf{\Omega}) = \frac{1}{4\pi\, k_{\text{eff}}} \sum_i \chi_i(\mathbf{r}, E) \sum_{g',b'} \int_{\Delta E_{g',b'}} \mathrm{d}E'\; N_i(\mathbf{r})\, \nu\sigma_{i,f}(\mathbf{r}, E')\, b_{g',b'}(\mathbf{r}, E')\, \phi_{g',b'}(\mathbf{r}), \tag{3.14}$$

with index i running over the nuclides.

3.3 Definition and Generation of the Bands

The PG-FEMG method utilizes global band memberships, which stem from one set of problem-wide $\Delta\Sigma_{t,g,b}$. To choose these band cross sections, we simply divide

up the problem-averaged total cross section over the coarse group:

$$\Delta\Sigma_{t,g,b} = \left\{ \Sigma_t(E) \,\middle|\, \Sigma_t(E) \in [\Sigma_{t,g,b-1}, \Sigma_{t,g,b}) \right\} \quad b = 1, \ldots, B_g, \tag{3.15}$$

for B_g bands in group g, where $\Sigma_{t,g,0}$ and Σ_{t,g,B_g} are the minimum and maximum values of $\Sigma_t(E)$ in coarse interval g. Notice that this effective $\Sigma_t(E)$ has no space, angle or time dependence: it is an effective average over the entire problem. We use a number- and flux-weighted average to determine this $\Sigma_t(E)$ from the $\sigma_t(\mathbf{r}, E)$, where the flux weighting comes from region-averages of the reference calculation for "high-fidelity" weighting cases or is unity for "low-fidelity" weighting cases. We discuss this more below.

We choose the band divisions by a simple affine recurrence relation:

$$\Sigma_{t,g,0} = \min_{E \in \Delta E_g} \Sigma_t(E), \tag{3.16a}$$

$$\Sigma_{t,g,B_g} = \max_{E \in \Delta E_g} \Sigma_t(E), \tag{3.16b}$$

$$\Sigma_{t,g,b} = \beta\, \Sigma_{t,g,b-1} + \alpha, \quad b = 1, \ldots, B_g. \tag{3.16c}$$

These imply:

$$\Sigma_{t,g,B_g} = \beta^{B_g} \Sigma_{t,g,0} + \alpha \left(1 + \beta + \ldots + \beta^{B_g - 1}\right), \tag{3.16d}$$

$$= \beta^{B_g} \Sigma_{t,g,0} + \alpha \frac{1 - \beta^{B_g}}{1 - \beta}, \tag{3.16e}$$

$$= \beta^{B_g} \Sigma_{t,g,0} + \gamma (1 - \beta^{B_g}) \Sigma_{t,g,0}, \tag{3.16f}$$

$$\Sigma_{t,g,B_g} = \Sigma_{t,g,0} \left(\gamma + (1 - \gamma)\beta^{B_g}\right), \tag{3.16g}$$

$$\beta = \left[\frac{\left(\Sigma_{t,g,B_g} / \Sigma_{t,g,0}\right) - \gamma}{1 - \gamma} \right]^{1/B_g}, \tag{3.16h}$$

for some choice of $\gamma \in (-\infty, 1)$. Notice that PG-FEMG simply reverts to MG in the case of $B_g = 1$.

Using the definition of γ yields α:

$$\gamma \equiv \frac{\alpha}{\Sigma_{t,g,0}\,(1-\beta)}, \tag{3.17}$$

$$\Rightarrow \alpha = \gamma\,(1-\beta)\,\Sigma_{t,g,0}. \tag{3.18}$$

This scheme interpolates the bands between $\min \Sigma_t(E)$ and $\max \Sigma_t(E)$ in the coarse group. The user must supply a value for γ, with a value of zero implying straight logarithmic division, and a positive value giving increasing spacing: $\Sigma_{t,g,b+1}/\Sigma_{t,g,b} > \Sigma_{t,g,b}/\Sigma_{t,g,b-1}$.

We used a value of $\gamma = 0.80$ in order to put more bands at the lower cross sections and fewer bands near the higher cross sections. This seemed a reasonable way to keep fidelity where most of the data was located (in a small resonance), while still maintaining fidelity in the large resonances.

In the case that the above method failed to produce a band that contained data, a linear interpolation of the indices of the sorted $\Sigma_t(E)$, $E \in \Delta E_g$, was used to determine the $\Sigma_{t,g,b}$. For example, if there were 10 data points in coarse group g and 2 bands were desired, then,

$$\Sigma_{t,g,0} = \Sigma_t[0], \tag{3.19a}$$

$$\Sigma_{t,g,1} = \Sigma_t[5], \tag{3.19b}$$

$$\Sigma_{t,g,2} = \Sigma_t[9], \tag{3.19c}$$

given $\Sigma_t[0:9]$ sorted ascendingly. In practice we found that this was invoked only for the energy resolution studies where asymptotically large numbers of groups/bands

were desired.

3.4 Examples of Total Cross Section and Band Generation

Figure 3.1 gives an example effective total cross section calculation. The microscopic cross sections (a) are combined in a weighted sum to produce a problem-wide, effective total cross section. Different weighting choices produce different effective total cross sections (b,c,d). Obviously the band boundaries are sensitive to this choice.

For problems with multiple regions, each region containing a constant macroscopic cross section, Fig. 3.1 may again be used, except with the microscopic cross sections replaced with region-wise macroscopic cross sections. A scheme to average these macroscopic cross sections must be employed.

For this work, two different weighting spectra were compared for each problem: a low-fidelity choice and a high-fidelity choice. In the context of the effective total cross section, the low-fidelity choice was to use the unweighted average of region-specific macroscopic total cross sections as the effective total cross section; the high-fidelity choice was to use a reference-flux-weighted average of the region-specific macroscopic total cross sections. Each region was defined to have a constant macroscopic cross section: different materials and/or different temperatures created different regions. When the reference flux was used, it was region-averaged in space, and integrated in angle and energy. While these choices may not have been optimal, we believe that using a weighted sum of region-local macroscopic total cross sections provides a natural and straightforward choice that is correct in the limit of one region. The effective total cross section algorithm is an area for future investigation.

Figures 3.2 and 3.3 give an example of how the band boundaries were chosen. Once an effective total cross section had been formed (Fig. 3.2a), and energy group

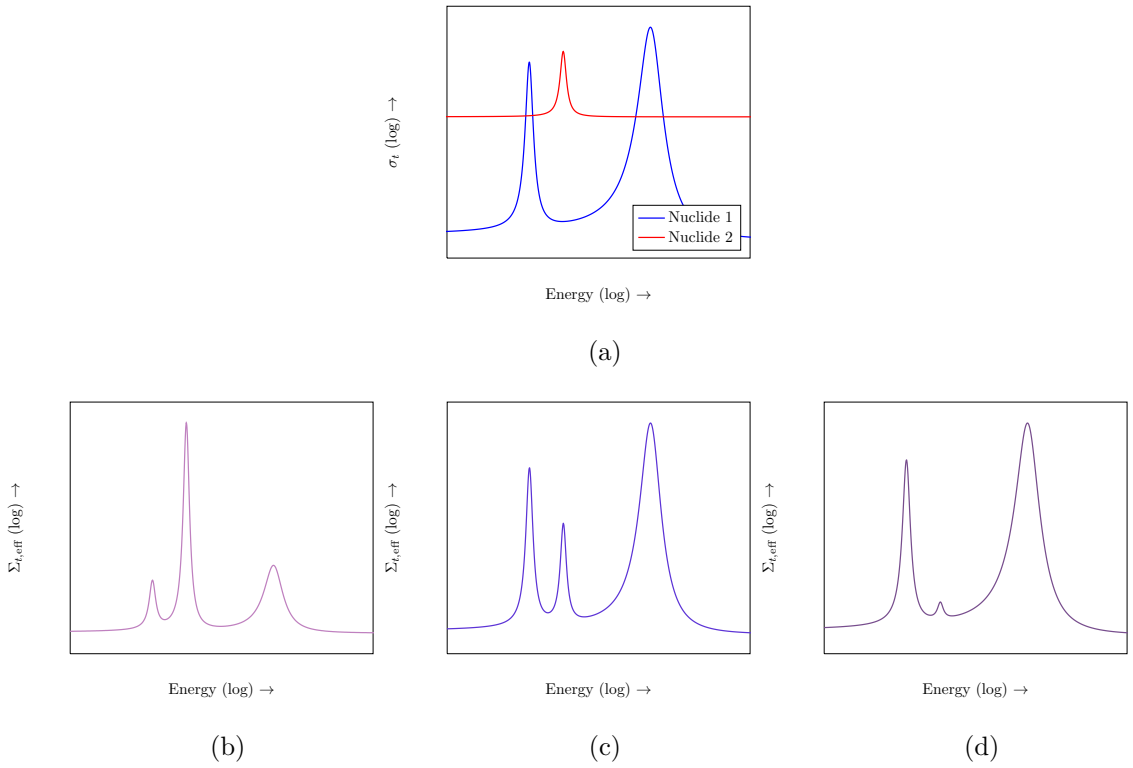

Figure 3.1: Example effective total cross section calculation. Microscopic total cross sections (a) are combined with different relative weightings into different problem-wide effective macroscopic total cross sections (b,c,d).

boundaries were chosen (Fig. 3.2b), the minimum and maximum effective total cross section for that group were determined (Fig. 3.2c). Equation (3.16) was used to determine the band boundaries (Fig. 3.2d), which uniquely defined the $\Delta E_{g,b}$ (Figs. 3.2e and 3.2f), which in turn uniquely defined the weight functions (Fig. 3.3). The example in Fig. 3.2 has three bands and one group, so there are three (orthogonal) weight functions in Fig. 3.3.

Notice that physical downscatter in energy would cause a particle to go from the green band to the blue band to the red band, back to the blue band, back to the green band, etc., in Fig. 3.2f. Upon applying Eq. (3.10d), no matter how the bands were ordered, there would be scattering from a band of higher number to a band of lower number in the scattering matrix. In standard MG, upper-diagonal terms in the scattering matrix are non-zero when upscattering occurs. When using a standard MG solver with the PG-FEMG cross sections, it would look like the PG-FEMG method produced artificial upscattering (artificial because the neutrons do not physically upscatter at resonance energies). This is an unavoidable consequence of the method when using energy group sizes that are much wider than resonance widths.

Figures 3.4 and 3.5 give another example, this time with two groups and two bands per group. The points to take away from this example are that the weight functions are global, orthogonal in energy, and spatially invariant; and that each group has its own band structure.

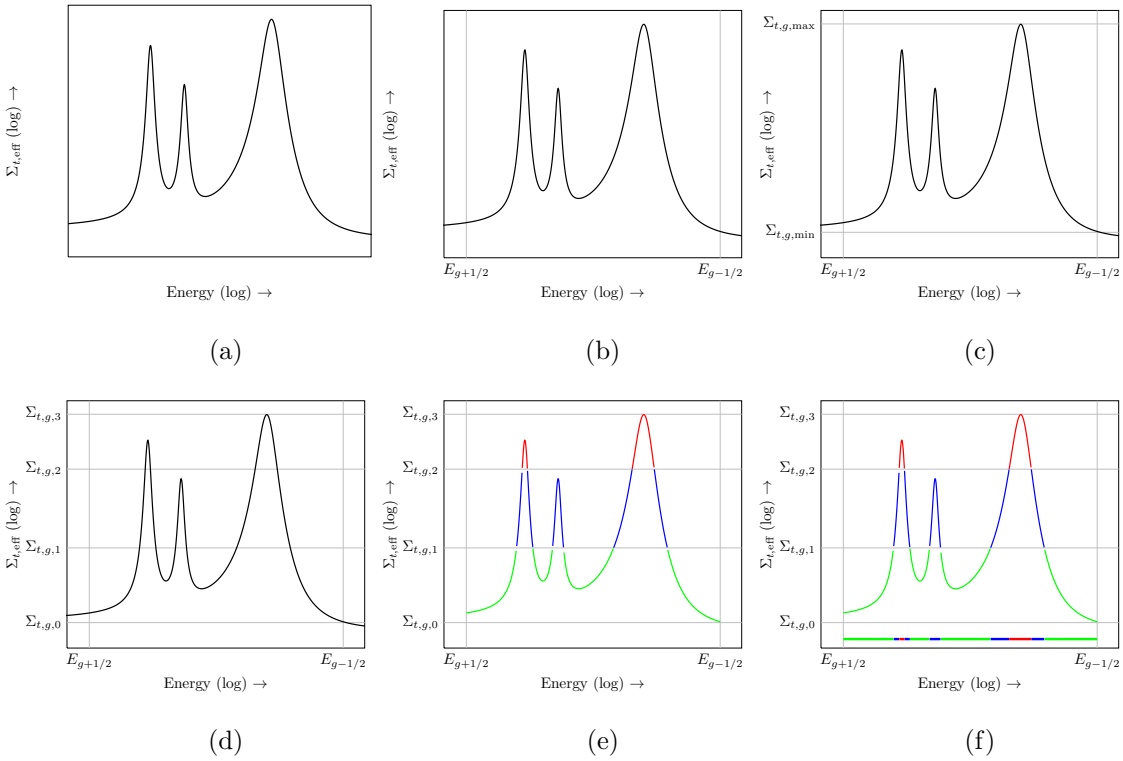

Figure 3.2: Example band boundary determination for one group and three bands per group. Begin with a macroscopic cross section (a), determine energy / group (b) and cross section / band bounds (c,d), and split into bands (e,f).

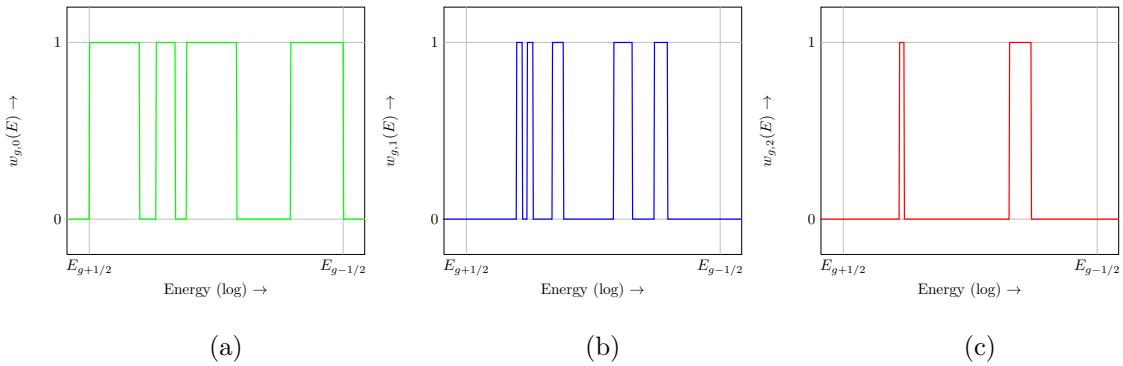

Figure 3.3: The three orthogonal weight functions corresponding to the band boundary calculation in Fig. 3.2.

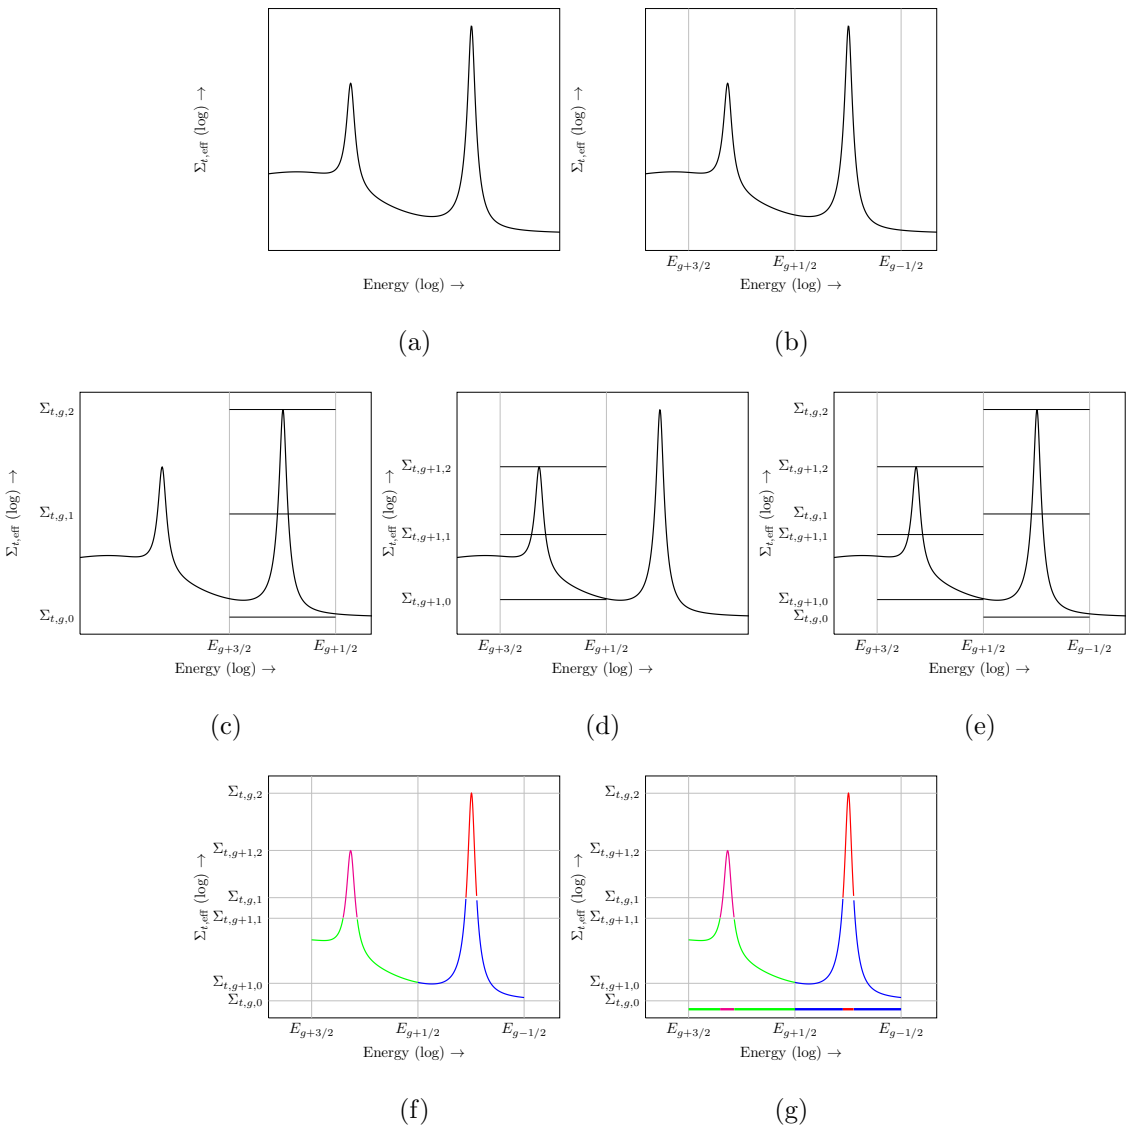

Figure 3.4: Example band boundary determination for two groups and two bands per group. Begin with a macroscopic cross section (a), determine energy (b) and cross section bounds (c,d,e), and split into bands (f,g).

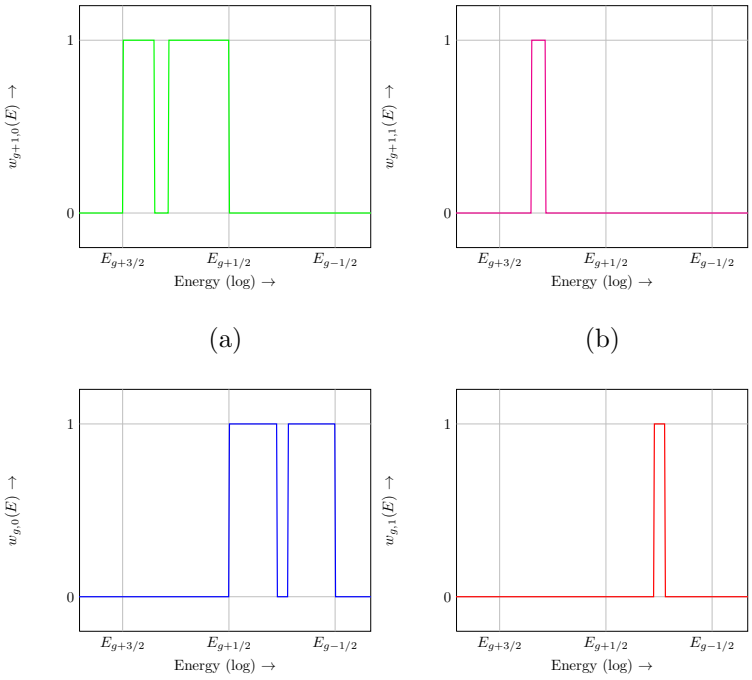

Figure 3.5: The four global weight functions corresponding to the band boundary calculation in Fig. 3.4.

4. WORKFLOW AND CROSS SECTION PREPARATION

Cross sections are parameters in the transport equation that represent reaction rates per unit flux or differential quantities such as fission energy spectra. Such quantities are measured directly or indirectly by experiments. Theory is used to interpolate datapoints and statistics are used to combine the results of disparate studies. The result is the creation of libraries that contain the values of cross sections at specific energy values — called pointwise cross sections. These cross sections were processed by the LANL code NJOY to produce ultra-fine-multigroup cross sections that describe pointwise cross sections properly averaged out over energy groups which were small with respect to the size of the resonances. A home-spun (C++) code was used to determine the band structure from the ultra-fine-multigroup cross sections and collapse these cross sections into the PG-FEMG and coarse-group MG cross sections necessary for this research.

A recent report from LANL listed cross section preparation as the (often overlooked) weakest link in the calculation process for the slowing down of neutrons in water using Monte Carlo codes [44]. The myriad options available to the user of NJOY, along with the obscurity and complexity of the output format, make this result unsurprising. We circumvent this issue by only using NJOY to create ultra-fine-group cross sections and by making all comparisons to a reference solution.

4.1 Pointwise Libraries

The Cross Section Evaluation Working Groups (CSEWG), coordinated through the National Nuclear Data Center at Brookhaven National Laboratory, handles the ENDF libraries and formats. These libraries contain the evaluated partial cross sections, resonance parameters, transfer matrices, etc. for the United States. Such

data has been verified and validated for numerous cases, including nuclear reactor analysis. Uncertainties in the cross sections are given in the form of covariance matrices; using such data is beyond the scope of this thesis.

At the time of writing, the most recent cross section evaluation was ENDF/B-VII.1, released in December 2011. Some of these data files proved to be inpractically large, especially for fissile nuclides. The slightly older but highly similar ENDF/B-VI.8 data were used for this work. This does not pose a complication to the work, as the same cross sections were consistently used for the reference, MG, and PG-FEMG calculations.

4.2 NJOY

Los Alamos National Laboratory (LANL) developed the NJOY code to process the ENDF/B libraries into a useful form [45]. It is built of many modules that perform tasks such as reconstructing the cross sections on a common energy grid in a point-wise format (RECONR), broadening the resonances of the cross sections correctly to account for material temperature (BROADR), treating the thermal cross sections — especially scattering — through $S(\alpha, \beta)$ and other treatments (THERMR), treating the unresolved-resonance self-shielding (UNRESR), converting point-wise cross sections into MG cross sections (GROUPR), and converting those MG cross sections into an easily readable format (WIMSR or MATXSR) [46]. NJOY was used to take the ENDF/B-VI.8 data and convert it to the MATXSR format.

Listed below are several relevant input parameters for the NJOY runs. An infinite background cross section (σ_0) was used for simplicity; this choice had no effect on the ultra-fine-group cross sections, as that group structure resolved the resonances[1]. The

[1]It may have influenced the unresolved resonance region, but this region was not of interest to the present work. Further, such influences were consistent for reference, MG, and PG-FEMG cross sections.

built-in LWR spectrum (IWT=5) was used for flux-weighting the ultra-fine group cross sections; these weightings were saved for future use. The ideal gas thermal treatment of the cross sections was used, which assumes each nuclei exists as a mono-atomic gas; this was done for simplicity and because the more advanced thermal treatments, e.g., hydrogen in liquid water, required addition checks not present in the MATXSR to PDT-XS converter, such as a recalculation of the total cross section in the thermal region for consistency.

Several cross sections were ignored in this work. In the scattering kernel, inelastic scattering, $(n, 2n)$, etc., and anisotropic scattering were ignored. When considering absorption reactions, only capture, (n, γ), and fission, (n, f), were included. We feel this offered sufficient similarity to and fidelity of real reactor problems while still providing a simplified model on which to do our research. Note that the cross sections used in the reference solution were always consistent with those used in the MG and PG-FEMG solutions, so our exclusions affected all problems equally.

The ultra-fine-group energy structure is described in future sections. There are two items of note, the first being that the structure was only ultra-fine in the resonance region; the second being that only the lower-energy portion of the resolved resonance region was assumed to be resolved. Energies higher than the cutoff (given below) were assumed to be in the unresolved resonance region. Again, consistency was maintained among the reference, MG, and PG-FEMG solutions; equally importantly, these restrictions allowed fidelity to be maintained in the resonance region, as the reference solution was memory-limited.

4.3 Code Flow

Running a single problem required the complex orchestration of several players. NJOY — the first player — created the reference ultra-fine-group cross section files.

The inputs to this stage were ENDF/B-VI.8 cross sections from the standard ENDF libraries. These were processed by NJOY from a script created by the Python master script. The results were MATXS files, which were passed through the XsConverter converter — the second player — into PDT- and SCDT-readable cross sections.

These fine-group ($\mathcal{O}(1000)$ groups) cross sections were passed to SCDT — the third player — for a reference calculation. The NJOY processing and reference calculations accounted for the majority of the time spent doing the calculations. However, they only needed to be done once per problem. The reference fluxes were averaged over spatial regions for possible future use. (Recall that a spatial region has a constant macroscopic cross section.)

The cross sections in each spatial region were condensed by XsCondenser — the fourth player — using either a region-averaged reference or a generic flux. The generic flux was the flux already present in the cross section files from the NJOY calculation, a "mid-life PWR" flux (IWT=5) with infinite background cross section. The reference flux was meant to be high-fidelity and the generic-flux was a lower bound on low fidelity. The condensation process was used to produce both MG and PG-FEMG cross sections. To ensure fair comparisons, the total degrees of freedom (DOF) were the same for both MG and PG-FEMG condensations. A desired resolution and number of bands to use in the resonance region were specified and group boundaries were calculated from the equal DOF constraint.

The energy region was split into five regions: fast ($2\times10^7-3\times10^4$ eV), unresolved resonance ($3\times10^4-1\times10^3$ eV), resolved resonance ($1\times10^3-3\times10^0$ eV), epithermal ($3\times10^0-6\times10^{-1}$ eV), low-lying Pu resonance ($6\times10^{-1}-1\times10^{-1}$ eV) and thermal ($1\times10^{-1}-1\times10^{-5}$ eV). The resolved resonance region was the region of interest for the PG-FEMG method; it was here that more than one band per group was used. Note that this region does not extend throughout the full resolved resonance

region, but merely the lower-energy component. We were limited by the the number of groups required for the reference calculation and the associated group-squared memory and computation requirements.[2] Group energies go from high to low to follow the neutron lifetime (born fast and then slow down). Doing so makes the MG scattering matrix lower-diagonal for all but thermal energies.

The newly condensed cross sections were given again to SCDT — reprising its role as the fifth player — for the MG and PG-FEMG calculations. Because PG-FEMG cross sections are readable with standard MG codes, no modifications were made to SCDT for these calculations.

Once the calculations were done, a post-processing step was done to expand the MG and PG-FEMG fluxes back to the reference spectrum — the sixth and final player. Since the condensation method was known, expansion was trivial (Eq. (3.5b)).

Several QOI were investigated, including k-eigenvalue, the power shape, and absorption rates in each spatial region. Errors in these QOI were computed by comparing the MG and PG-FEMG solutions to the reference solution.

Figure 4.1 shows the flowchart for the work done in this thesis, from ENDF to NJOY, to SCDT, to XsCondenser, to Python.

[2]Recall we are not using anisotropic cross sections for simplicity. If we were, the angular moments required to accurately reproduce the scattering kernel would increase along with the group count, resulting in at least group-cubed memory and computation requirements. We hope to address this in future work.

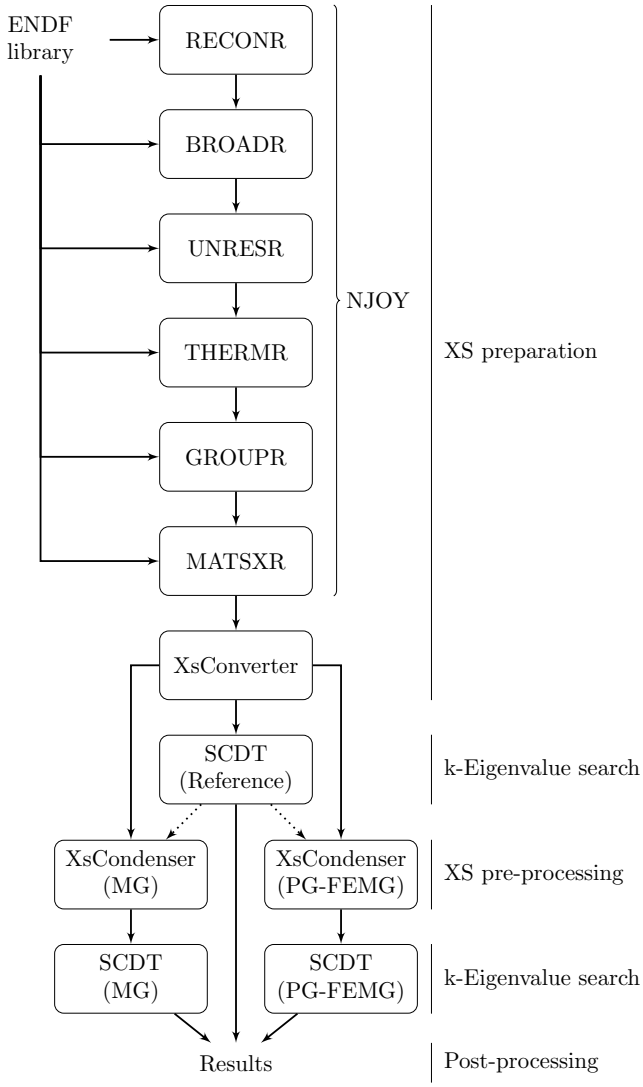

Figure 4.1: Thesis work flowchart. System progresses from top to bottom. Dotted lines indicate possible information flow; solid lines indicate information flow.

This process was repeated for several energy resolutions and for high- and low-fidelity weighting schemes for each problem studied. For a given run, the same flux-weighting was used to produce both MG and PG-FEMG cross sections.

The above paragraphs but summarize the intricacies of this thesis, especially with reference to the cross section preparation process. While the interesting parts of this

work may be distilled into a novel cross section condensation scheme, doing so ignores the substantial infrastructure required to orchestrate the whole performance and produce useful results. A Python pre-processing script creates the myriad input decks for the cross section creation, k-eigenvalue search, cross section condensation, and post-processing codes. These are tied together by an automatically generated master script which calls each part at its proper time and provides it with its proper inputs as their expected names. The algorithms of the cross section condensation code were described in the previous section. The inputs to the cross section creation code (NJOY) were described in this section. The k-eigenvalue search code is described in the following section.

5. SOLUTION TECHNIQUES FOR THE TRANSPORT EQUATION

The discrete system created by discretizing the transport equation must be able to be solved efficiently if it is to be useful. A germane question in this thesis is whether the proposed PG-FEMG method provides more reduction in error than it requires in additional work. We define what we mean by work below.

Efficiency of an algorithm may be measured in various ways, including time to solution, memory usage, required floating point operations (FLOP), or by some representative proxy. The research codes used for this work were only moderately optimized for speed, memory usage, or FLOP usage. Further, these metrics are often either system dependent, compiler dependent, or both. We instead introduce required sweeps as our proxy. It has been observed that, especially for multi-dimensional problems, most of the compute time and effort is contained within the sweep. MG and PG-FEMG differ only in the pre- and post-processing of the group-dependent fluxes. The solve steps and hence sweep steps are identical; they simply use different cross section data. For this reason, we feel sweeps in this code are a valid and accurate indication of work. We ignore the cost of pre- and post-processing because it is needed for both schemes and because it may be done ahead of time or afterward, respectively. Future work may look into quantifying this cost.

We define a sweep as the inversion of the left-hand side of the transport equation over angle and space for one energy group. That is, the application of $(\nabla\cdot\Omega+\Sigma_{t,g,b})^{-1}$ for a given group g and, possibly, band, b. While it may be more common to call one sweep the above process over all energy groups (bands), that does not provide the required granularity when comparing methods with highly disparate numbers of groups. All figures showing "required sweeps" will clearly declare if they are averaged

(typical definition) or summed (our definition) over the number groups.

Efficiency incorporates algorithmic design, choice of programming language, targeted hardware, and effective expression of the algorithmic idea within this language for this hardware. The problem of interest is an eigenvalue problem, which may be solved with either general nonlinear or eigenvalue-specific iterative algorithms. The algorithms used are described in the following pages. The transport equation solver was written in C++ for its properties of speed, object-oriented design, and popularity in the scientific community. The targeted hardware was a single-threaded CPU; parallelizing the transport equation with its serial dependencies in space (the streaming), angle (the angular derivative), and energy (the down- and up-scattering) is a non-trivial task beyond the scope of this thesis. To the extent feasible, best practices were used with respect to method encapsulation, use of libraries for well-defined tasks such as input parsing and nonlinear solves, and use of standard data types. The author does not claim to have written the optimal transport code for this problem, though substantial thought and effort was put into the design of the code.

Investigation of the generalized-MG method required writing a code that could solve the k-eigenvalue formulation of the transport equation. The following pages describe the algorithmic techniques used to solve this formulation. This section is split into three main topics. The first topic is traditional schemes to solve the k-eigenvalue problem: power iteration of the fission source as an outer iteration, and source iteration of the scattering source as an inner iteration. The second topic is a new alternative to the first: using NKA to converge both sources simultaneously. The third topic regards the code that actually implements these algorithms: the Serial, Cylindrical Deterministic Transport solver (SCDT).

Full descriptions of the iterative methods employed are located in Appendix B.

5.1 Traditional Iterative Schemes and the Need for NKA

The traditional iterative technique is to use nested iterations: an outer iteration on the fission source and eigenvalue, and several layers of inner iterations on the scattering source. These inner iterations are generally Gauß-Seidel (GS) in energy to converge between-group scattering outside of an iteration to converge within-group scattering. Preconditioning and Krylov formulations are often done within a group [47], though sometimes there is another outer Krylov application to converge between-group scattering in the case of upscattering (GS is an exact inverse for downscattering). Traditionally, for each level of the nested iteration scheme, multiple inner iterations are done for each outer iteration. Because the eigenvalue update is nonlinear, there will always be at least two nested levels using traditional schemes: an outer nonlinear update with Picard or Newton, and an inner linear update involving Krylov.

The insight of modern iterative methods, as epitomized by Warsa (e.g., [49]), is that this nested scheme can be inefficient, as it relies on many applications of an expensive innermost iteration — a transport sweep with DSA preconditioning in this case — for every update of an outer iterate quantity — here the eigenvalue or upscattering source. Warsa looked for algorithms and iterative techniques that would break nesting, including the nesting of linear iterations within nonlinear iterations. Nonlinear Krylov acceleration (NKA), developed by Carlson and Miller [48], allowed the combination of linear and nonlinear iterations in a Krylov-like framework. Using only one inner iteration per outer iteration at each level of the previously nested scheme broke nesting.

NKA, as implemented by Calef et al. [49], is advantageous over traditional iterative schemes in the presence of the effective upscattering introduced by the PG-

FEMG method for two reasons. The first is a consequence of broken nesting: less work is done between upscattering source updates, which allows the upscattering source to be converged faster. The second is a consequence of both broken nesting and the expense of the transport sweep: Krylov updates are inexpensive compared to sweeps [47]; performing more Krylov updates per sweep, as is realized from broken nesting, means more of the work is given to the less expensive Krylov update. These ideas, combined with the possibility of completely broken nesting using NKA, will allow faster convergence for k-eigenvalue problems with effective upscattering.

There is one significant caveat: broken nesting is advantageous only if the inner iterations are reasonably rapidly convergent and all inner iterations converge approximately equally rapidly. For example, if one group has a high within-group scattering ratio and no DSA preconditioning, it can be more efficient to iterate on the within-group scattering source for that one group before moving on to the next group.

5.2 Overview of Traditional Schemes

The transport equation may be treated in more generality using operator notation. The notation may refer to both the differential transport equation

$$\mathcal{L}\psi = \left(\mathbf{\Omega}\cdot\boldsymbol{\nabla} + \Sigma_t(\mathbf{r}, E)\right)\psi(\mathbf{r}, E, \mathbf{\Omega}), \tag{5.1a}$$

$$\left(\mathcal{S}\phi\right)_\kappa = \int_0^\infty \mathrm{d}E' \;\; \Sigma_{s,\kappa}(\mathbf{r}, E' \to E)\,\phi_\kappa(\mathbf{r}, E'), \tag{5.1b}$$

$$\left(\mathcal{F}\phi\right)_\kappa = \delta_{0,\kappa}\,\chi(\mathbf{r}, E)\int_0^\infty \mathrm{d}E' \;\; \nu\Sigma_f(\mathbf{r}, E')\,\phi_0(\mathbf{r}, E') \tag{5.1c}$$

— where here κ is the index of the $Y_{l(\kappa)}^{m(\kappa)}(\mathbf{\Omega})$ or $P_{l(\kappa)}(\mu_0)$ moment — or to the discretized one.

Using the standard moment-to-discrete and discrete-to-moment mappings (Ap-

pendix A), the transport equation simplifies to

$$\mathcal{L}\psi = \mathcal{MSD}\psi + q, \tag{5.2a}$$

where

$$q = \frac{1}{k_{\text{eff}}}\mathcal{MFD}\psi \tag{5.2b}$$

in the case of the k-eigenvalue problem.

The traditional schemes for solving the k-eigenvalue problem use a nested set of iterations. The outer iterations converge the fission source, $\mathcal{MFD}\psi$, and eigenvalue, k_{eff}; the inner iterations converge the scattering source, $\mathcal{MSD}\psi$, for a given fission source (given q). As the usual method to update the eigenvalue is power iteration, the entire scheme is sometimes referred to as power iteration, though this is a bit of a misnomer, as it technically only applies to the outer iteration.

The iteration indices used will be l and u for the inner iterations (for convergence of the within-group and across-group scattering sources, respectively), and z for the outer iterations (for convergence of the fission source).

5.2.1 Convergence of the fission source and eigenvalue through power iteration

Equation (5.2) may be cast as a convergence of the fission source problem by lagging the fission source to the previous iteration

$$\mathcal{L}\psi^z = \mathcal{MSD}\psi^z + \frac{1}{k^{z-1}}\mathcal{MFD}\psi^{z-1} \tag{5.3a}$$

Equation (5.3a) may be cast in standard eigenvalue form by inverting the transport-plus-scattering operator. This is functionally equivalent to converging the scattering

source.

$$(\mathcal{L} - \mathcal{MSD})\,\psi^z = \frac{1}{k^{z-1}}\mathcal{MFD}\psi^{z-1} \tag{5.3b}$$

$$k^{z-1}\psi^z = (\mathcal{L} - \mathcal{MSD})^{-1}\,\mathcal{MFD}\psi^{z-1} \tag{5.3c}$$

The k-eigenvalue is converged through power iteration, a common eigenvalue-solving algorithm [2]. This is a method of solving an eigenvalue system that picks out the largest eigenvalue in magnitude; as previously mentioned, this is k_{eff}. The update for k_{eff} is motivated by keeping the fission source the same size and is given by

$$1 = \frac{\frac{1}{{k_{\text{eff}}}^z} W^{\text{T}}\mathcal{F}\phi^z}{\frac{1}{{k_{\text{eff}}}^{z-1}} W^{\text{T}}\mathcal{F}\phi^{z-1}} \tag{5.4}$$

where W^{T} is an integral over all phase space, namely

$$W^{\text{T}}(\cdot) = \sum_{g=1}^{G} \frac{1}{4\pi} \int_{4\pi} \mathrm{d}\Omega \int_0^R \mathrm{d}r \;\; r\,(\cdot) \tag{5.5}$$

Equations (5.3c) and (5.4) provide the necessary update equations for eigenvector and eigenvalue, respectively.

True power iteration occurs when Eq. (5.3c) is solved to a tight tolerance. This can be inefficient when the fission source is far away from the true value, as an exact inversion of the transport-plus-scattering operator is not necessary for initial convergence. For this reason, the initial tolerance on the scattering convergence can be loose, and tightened as the eigenvalue and eigenvector approach the solution. While such schemes are beyond the scope of this work, NKA performs a similar role in reducing iteration time.

The convergence of power iteration is determined by a parameter known as the dominance ratio, which is the ratio of the penultimate to ultimate eigenvalue magnitudes [2]. For systems with small dominance ratios, power iteration is efficient. Examples of such systems include systems in which leakage rates are not small relative to absorption rates. Power iteration converges slowly for systems with large dominance ratio. Systems in which leakage rates are small relative to absorption rates generally have large dominance ratios. A common example of such a system is the full core of a LWR. For such situations, it is desired to accelerate the eigenvalue solve. This is the purpose of NKA, described in future pages.

5.2.2 Convergence of the across-group scattering source through GS

Typical convergence of the across-group scattering source is done through a Gauß-Seidel (forward substitution) process by energy group. As most of the scattering goes from higher neutron energies to lower neutron energies, the highest-energy group is solved first; the other groups are solved in order of highest to lowest energies.

As the GS process yields the exact inverse for a lower-diagonal matrix, no convergence criterion is needed for problems without upscattering, as the inversion will be exact. For processes with upscattering, two options exist. The first is to ignore upscattering in the inner iterations and converge it through the outer iterations (fission source iterations). This works well for small amounts of upscattering or when accelerations are applied to the system (e.g., NKA).

The second option is to repeat the GS process in the upscattering groups. In the case of PG-FEMG, the set of upscattering groups includes both thermal groups that possess true upscattering and epithermal groups with the band structure that creates artificial upscattering within a coarse group due to resonances (non-monotonic cross sections). The GS process may be done normally or in reverse — i.e., from lowest to

highest energies. The latter may be efficacious when the upscattering approximately equals the downscattering in magnitude.

5.2.3 Convergence of the within-group scattering source through SI

Equation (5.2) may be cast as a convergence of the within-group scattering source problem by lagging the within-group scattering source and q to the previous iteration

$$\mathcal{L}\psi^l = \mathcal{MSD}\psi^{l-1} + q, \tag{5.6a}$$

for $\mathcal{S}$ representing only within-group scattering. In-scattering from other groups and the fission source are represented by q.

To update the system, the action of $\mathcal{L}^{-1}$ must be applied:

$$\psi^l = \mathcal{L}^{-1}\left(\mathcal{MSD}\psi^{l-1} + q\right). \tag{5.6b}$$

This action for a single group or band within a group is referred to as a sweep and is described below.

Upon convergence of the within-group scattering source, the satisfies

$$\psi = \left(\mathcal{L} - \mathcal{MSD}\right)^{-1} q \tag{5.6c}$$

5.2.4 Inversion of the streaming operator through sweeps over angle and space

The inversion of the streaming operator, $\mathcal{L}$, may be done exactly by the process of a sweep. Both the angular[1] and spatial derivatives in $\mathcal{L}$ are first-order and physically represent advection / streaming. When put in a matrix form, they constitute block-lower-diagonal matrices that may be inverted without the need of iterations; thus,

[1]Recall that angular derivatives arise from a cylindrical coordinate system representation.

no convergence criteria are needed.

The ordering of the sweeps has a loop over angle outside a loop over spatial cells. Physically, particles advect along their direction of travel ($\mathbf{\Omega}$). For a given angle, the sweep direction in space is therefore well-defined and it is natural to put the loops in space inside the loops in angle.

Boundary conditions (BC) define the starting point of a sweep. For angular sweeps, these conditions must be generated by special sweeps over the starting-direction. For spatial sweeps, these conditions are given in terms of incident angular fluxes. For positive μ_λ, this means beginning at $r = 0$ and sweeping to $r = R$ with the level-isotropic center flux as the BC to set the upwinded flux. For negative μ_λ, this means sweeping from $r = R$ to $r = 0$ with a BC at $r = R$ to set the upwinded flux.

5.2.5 Acceleration of the within-group scattering source through S2SA

After every sweep, the within-group scattering source convergence may be accelerated through S_2 synthetic acceleration (S2SA). A sweep performs the following update for a given source

$$\mathcal{L}\psi^{l+1/2} = \mathcal{MSD}\psi^{l} + q. \tag{5.7a}$$

If the system were converged at iteration $l + 1$, it would satisfy

$$\mathcal{L}\psi^{l+1} = \mathcal{MSD}\psi^{l+1} + q. \tag{5.7b}$$

Taking the difference between these two equations yields

$$\mathcal{L}\left(\psi^{l+1} - \psi^{l+1/2}\right) = \mathcal{MSD}\left(\psi^{l+1} - \psi^{l+1/2} + \psi^{l+1/2} - \psi^{l}\right), \tag{5.7c}$$

or

$$\mathcal{L}\,\delta\psi = \mathcal{MSD}\,\delta\psi + q, \tag{5.8a}$$

$$\delta\psi = \psi^{l+1} - \psi^{l+1/2}, \tag{5.8b}$$

$$q = \mathcal{MSD}\left(\psi^{l+1/2} - \psi^{l}\right). \tag{5.8c}$$

Equation (5.8a) is just as hard to solve as the transport equation. To form a preconditioner, moments of Eq. (5.8a) are taken and the high-order operator $\mathcal{L}$ is replaced with the low-order diffusion operator $\mathcal{L}_{\mathcal{D}}$. For this problem, this is accomplished by assuming a linear trial space for $\delta\psi$ in $\mathbf{\Omega}$. Two unknowns are needed to match the two degrees of freedom of the space. The 0^{th} and 1^{st} moments of Eq. (5.8a) are taken and the system is transformed into the ψ^{-} and ψ^{+} basis. This yields the so-called P_1-equivalent S2SA. More details are provided in Appendix A.

The S2SA system is hepta-diagonal and may be solved with little memory overhead using direct methods such as LU-decomposition. No convergence criteria are necessary for this inversion. See Section B.3.5 for more details.

Once the system is solved, the solution is transformed once again into the $\delta\phi$ and δJ basis. These values are added to the current estimates to complete the acceleration step

$$\phi^{l+1} = \phi^{l+1/2} + \delta\phi. \tag{5.9}$$

The S2SA accelerates the convergence of the scattering source, which depends on the moments of the flux only. Since we are using an isotropic scattering kernel, we do not need to update (or store) the currents, except on the boundaries.

Our problems of interest — pin cells — are modeled with white boundaries. To

accelerate these implicit boundary conditions and to preserve balance, the angular fluxes at the boundary are updated as well.

5.3 Overview of NKA

A nonlinear Krylov acceleration (NKA) scheme was used to accelerate the non-linear k-eigenvalue problem. In addition to the typical acceleration of the fission source and eigenvalue, this method was also used to accelerate the scattering source. Notationally, once again z will be used for outer iterations and l for inner iterations. Inner iterations are the same as above (acceleration of the within-group scattering source inside a GS iteration over groups). The standard NKA method uses only one inner iteration per outer iteration, diverging sharply from the fully-resolved inner iterations of the power iteration scheme.

NKA bears a striking resemblance to Krylov methods, which have gained popularity in recent years as a means of stabilizing and hastening linear solves both in nuclear engineering and more generally. NKA may also be found in the literature as non-linear GMRES or Anderson Mixing. Appendix B provides a background on Krylov methods for the uninitiated.

5.3.1 Anderson Mixing method

The NKA method used in this work was developed by Carlson and Miller [48, 50]. Calef et al. [49] showed that NKA is a variant of Anderson Mixing [51]. Anderson Mixing splits the update of the solution vector into components in the vector space of the saved residuals, and components orthogonal to that space. For the former, Anderson Mixing uses information gained over successive iterates to compute the action of the inverse of the Jacobian. For the latter, having no other information, it simply uses a fixed-point update (that component itself). A fuller description of the NKA algorithm may be found in Appendix B.

The current gold standard[2] for Krylov-accelerated k-eigenvalue calculations is Jacobian-Free Newton Krylov (JFNK) [49, 52]. This consists of a Newton's method combined with a linear Krylov solver like GMRES. The name comes from the fact that the Jacobian matrix itself is never formed; only its action is required to solve for its inverse. GMRES is usually restarted after a set number of iterations, bounding its memory use.

NKA is distinct primarily in two ways. First, the vector iterated upon is always an approximation to the solution. This is in contrast to popular Krylov-based acceleration schemes like GMRES, where the vector iterated upon is a residual. This means it is easier to use lagged information or set-to-zero methods. Second, the nonlinear update is performed once per sweep. This is in contrast to the standard power iteration scheme applied to find the k-eigenvalue, where the scattering source is mostly converged between outer (fission source and eigenvalue) iterations.

NKA presents different strengths and weaknesses than JFNK. For M saved states (restart period in GMRES parlance), NKA memory requirements scale as order $2M$, while JFNK requirements scale as order M. For a low-to-typical M of 10 and including a typical overhead of 3 state saves, JFNK uses 56% of the memory required by NKA, which is substantial. JFNK consists of outer Newton iterations wrapped around inner GMRES iterations. At the beginning of a new Newton iteration, all the subspace information from the previous GMRES iterations are lost. When the method is near the true solution, this represents unnecessary redundancy. In contrast, NKA uses information from all saved previous states.

[2]Traditionally, a gold standard meant that gold was used to back up a paper currency. In the modern vernacular, a gold standard is the best or most commonly used method or idea. It is this second meaning that is used here.[3]

[3]See R. McClarren's dissertation for more helpful footnotes.

5.3.2 Simultaneous convergence of the fission and scattering sources through NKA

NKA was used to simultaneously converge the fission and scattering sources. Using one inner iteration per outer iteration, the NKA method without the nonlinear update would look like

$$\mathcal{L}\psi^z = \left(\mathcal{MSD} + \frac{1}{k^{z-1}}\mathcal{MFD}\right)\psi^{z-1}. \tag{5.10}$$

Algebraic manipulation yields

$$\psi^z = \mathcal{L}^{-1}\left(\mathcal{MSD} + \frac{1}{k^{z-1}}\mathcal{MFD}\right)\psi^{z-1}, \tag{5.11a}$$

$$\phi^z = \mathcal{DL}^{-1}\mathcal{M}\left(\mathcal{S} + \frac{1}{k^{z-1}}\mathcal{F}\right)\phi^{z-1}. \tag{5.11b}$$

If, instead of iterative form, Eq. (5.11b) were put in residual form, it would be

$$f_\phi = \phi - \mathcal{DL}^{-1}\mathcal{M}\left(\mathcal{S} + \frac{1}{k}\mathcal{F}\right)\phi, \tag{5.12}$$

where f_ϕ is the residual on the eigenvector.

The basis for the residual on the eigenvalue is the power iteration update for k_{eff}, namely

$$1 = \frac{\frac{1}{k^z}W^{\mathrm{T}}\mathcal{F}\phi^z}{\frac{1}{k^{z-1}}W^{\mathrm{T}}\mathcal{F}\phi^{z-1}}. \tag{5.13a}$$

Some care must be taken when forming a residual from this quantity, as the equation is trivially satisfied when ${k_{\text{eff}}}^z = {k_{\text{eff}}}^{z-1}$ and $\phi^z = \phi^{z-1}$, as would occur when removing the iteration index. Instead, realize that Eq. (5.11b) defines what ϕ^z

should be. Inserting that into the update equation yields the desired result

$$1 = \frac{\frac{1}{k^z} W^{\mathrm{T}} \mathcal{F} \left(\mathcal{D}\mathcal{L}^{-1} \mathcal{M} \left(\mathcal{S} + \frac{1}{k^{z-1}} \mathcal{F} \right) \phi^{z-1} \right)}{\frac{1}{k^{z-1}} W^{\mathrm{T}} \mathcal{F} \phi^{z-1}}, \tag{5.13b}$$

$$f_k = 1 - \frac{W^{\mathrm{T}} \mathcal{F} \left(\mathcal{D}\mathcal{L}^{-1} \mathcal{M} \left(\mathcal{S} + \frac{1}{k} \mathcal{F} \right) \right) \phi}{W^{\mathrm{T}} \mathcal{F} \phi}. \tag{5.13c}$$

The residual on the eigenvalue is then

$$f_k = \frac{W^{\mathrm{T}} \mathcal{F} f_\phi}{W^{\mathrm{T}} \mathcal{F} \phi}. \tag{5.14}$$

Equations (5.12) and (5.14) give the residual equations. A residual vector is made by appending f_k on f_ϕ; this vector is handed to the black-box NKA solver described above. To ensure that both eigenvector and eigenvalue are converged equally, their magnitudes were kept similar by initializing all the ϕ's as unity in regions with a fission source. Note that the (memory) size of ϕ is generally much smaller than the size of ψ and depends on how many scattering moments need to be kept.

As the NKA solve is a nonlinear method, there is no guarantee that the system will converge to the fundamental mode. However, the update for the eigenvalue is consistent with power iteration, which does converge to the fundamental mode. Calef et al. [49] found that initializing the system by doing several power iterations was sufficient to ensure convergence quickly and to the fundamental mode. We repeat this process and find similar results.

The NKA method consists of doing one full sweep of the scattering source per group and then calculating the residual functions, Eqs. (5.12) and (5.14). These residuals are handed to the NKA solver, which hands back a δ vector. The eigenvector

and eigenvalue are updated with this δ vector for the next iteration:

$$\phi^z = \phi^{z-1} - \delta_\phi, \tag{5.15a}$$

$$k_{\text{eff}}{}^z = k_{\text{eff}}{}^{z-1} - \delta_k. \tag{5.15b}$$

This process is repeated until convergence is achieved.

Other residuals may be used instead of the ones provided in Eqs. (5.12) and (5.14). In the context of the NKA method, the sweep plays the role of a preconditioner. Within this paradigm, additional preconditioning can be applied. For example, after the sweep, S2SA can be applied to accelerate the scattering source. If the scattering source is still the problem, multiple source iterations ("inner" iterations) may also be done per NKA ("outer") iteration.

If ϕ^{z-1} is the state of the system at the beginning of an iteration and $\phi^{z-1/2}$ is the result of the action of the preconditioner on ϕ^{z-1}, then a more generalized residual form of the eigenvector is

$$f_\phi = \phi^{z-1/2} - \phi^{z-1} \tag{5.16}$$

Equation (5.14) remains the same but uses the new eigenvector residual.

5.4 SCDT

SCDT, short for *Serial Cylindrical Deterministic Transport*, is the code that was developed by the author to perform the required eigenvalue calculations for this work. It was written in C++ in an object-oriented framework. More information may be found in Appendix B.

6. CODE VERIFICATION AND PROBLEM SPECIFICATIONS

6.1 Overview of Cases Run

This section gives an overview of the the testing and analysis performed on the generalized-MG method in this thesis. The next section gives results. Three independent parameters were varied: problem, cross section weighting, and test type. Problems refer to the geometries and materials used; weights refer to the weighting spectrum used to collapse the (ultra-)fine-group cross sections into the coarse-group MG and PG-FEMG cross sections; test types refer to the energy group structure or set of structures used for each problem and each weight. Sets of structures were used for energy resolution studies.

Five problems were investigated. They were (1) a homogeneous, one-temperature pin cell; (2) a heterogeneous, one temperature pin cell; (3) a heterogeneous, four-temperature pin cell; (4) a heterogeneous, one temperature, multi-fuel pin cell; and (5) a heterogeneous, one-temperature, multi-fuel cylindricized fuel lattice. For each problem, three calculations were done: a fine-group reference solution, a standard MG coarse-group solution, and the PG-FEMG solution. The MG solution had the same DOF as the PG-FEMG solution, to compare fidelity per DOF. Note that the MG and the generalized-MG required the same memory allocation.

Two sets of weights were used. The first was the reference solution averaged over each spatial region; the second was a generic mid-life PWR spectrum from NJOY without discernible resonances in the resolved resonance region used in this work. These spectra were meant to provide bounds in fidelity for what might actually be done in real-world calculations. For some problems, region-averaging the reference solution was found to significantly degrade its accuracy as a weighting spectrum; the

following sections have more information. Recall that a spatial region is a set of cells with constant macroscopic cross section.

For each problem after the first (verification) problem, a standard series of tests were run with various energy structures to investigate the effect of energy structure on accuracy and required work. The tests were a base problem with unrefined but representative energy group structure and two energy resolution tests. Each test used the same energy grid structure for each problem and each weighting.

Both resolution studies investigated the effect of increasing the number of coarse groups and bands within the resonance region on required work as measured in sweeps and error in the QOIs with respect to the reference solution. The first study always resolved the non-resonance regions and changed resolution only within the resonance region. Order 3 - 70 DOF were used in the resonance region to investigate accuracy in the under-resolved limit.

The second study increased resolution in all regions and focused on asymptotic effects. It investigated cases with DOF spanning many orders of magnitude. The highest resolution of this study had MG / PG-FEMG DOF equal to the reference DOF, mainly for verification purposes; viz., to ensure the QOI error would be near the solver tolerance, which was around 10^{-8}.

6.2 Problem Specifications

Tables 6.1 and 6.2 give an overview of the geometry and materials, respectively, used in the five problems. Problem 1 was homogenized and had only 1 radius; problems 2 and 3 had only 1 fuel type (MOX) surrounded by moderator; problem 4 had 2 fuel types (MOX, then LEU) surrounded by moderator; problem 5 had an inner fuel (MOX), moderator, and an outer fuel (LEU). The next section gives more details for each problem, including cartoons of the geometry.

Table 6.1: Geometry specifications for the problems. All distances are in cm.

Problem	Inner fuel radius	Moderator radius	Outer fuel radius
1	10.0	N/A	N/A
2	0.470	0.722	N/A
3	0.470	0.722	N/A
4	0.332	0.722	0.470
5	0.470	1.193	1.444

Table 6.2: Material specifications used in the problems.

Material	Problems used	Density (1/b-cm)	Components (atom %)		
Slurry	1	0.07167	Pu-239 (0.8)	U-238 (23.3)	H-1 (75.9)
Moderator	2,3,4,5	0.07513	H-1 (66.7)	O-16 (33.3)	
MOX Fuel	2,3,4,5	0.07342	Pu-239 (1.7)	U-238 (31.7)	O-16 (66.6)
LEU Fuel	4,5	0.07342	U-235 (1.7)	U-238 (31.7)	O-16 (66.6)

As previously noted, the energy domain was divided into physics-based regions, including fast, unresolved, resolved, epithermal, Pu-239 low-lying resonance (epithermal II), and thermal. Our focus was primarily on the resolved resonance region, which was shrunk from the full resolved resonance region due to memory limitations on the reference solution. The purpose of this work was to capture many resonances in the reference solution, not to capture all the resonances or to get the perfect real-world answer. A large portion of the resolved resonance region was kept and resolved for the reference solution (up to 10^3 eV; see below).

Tables 6.3 and 6.4 give the details on the energy boundaries and DOF for the tests, including the refinement strategy. For each region in the energy domain, the Ref / BCG (RF) triplet gives the reference DOF (Ref), the base number of coarse groups for PG-FEMG for the test (BCG), and the refinement factor for the test if it was a resolution study (RF). A value of Ref / BCG (RF) = $F/N(R)$ would indicate the reference solution used F energy groups in the region and the PG-FEMG solution used $N \max(1, R\,r)$ coarse groups in the region for a resolution factor of $r > 1$. The MG solution would have used a number of groups equal to the PG-FEMG coarse group count multiplied by the number of bands in the region, b. Another `max` was used to ensure the PG-FEMG and MG DOF were less than or equal to the reference DOF in each region. The DOF for both MG and PG-FEMG in a region was $\min(F, b\,N \max(1, R\,r))$ for $r > 1$, and $\min(F, b\,N)$ for $r = 1$ or for base tests.

Two types of energy resolution tests were run. Res I is the resonance-only energy resolution study; Res II is the full energy resolution study. For the Res I study, refinement factors of $r = \{1, 2, 4, 8, 16, 64, 128, 256\}$ were used and $b = \{1, 2, 3, 4\}$ bands were used in the resonance region. For the Res II test, refinement factors of $r = \{1, 2, 4\}$ were used and $b = \{1, 2, 3, 4, 5, 6\}$ bands were used in the resonance region. All base problems used $r = 0$ and $b = 4$ bands in the resonance region.

Table 6.3: First half of the energy grid specifications used in the problems and their tests. Ref, BCG and RF stand for reference, base coarse groups and refinement factor, respectively. Energies denote upper bounds for the region, with the lowest energy bound being 1×10^{-5} eV.

Test name	Problem(s)	Fast (2×10^7 eV) Ref / BCG (RF)	Unresolved (3×10^4 eV) Ref / BCG (RF)	Resolved (1×10^3 eV) Ref / BCG (RF)
Base	1	1 / 1	1 / 1	1000 / 4
Base	2,3,4,5	1 / 1	1 / 1	1880 / 3
Res I	2,3,4,5	1 / 1 (0)	1 / 1 (0)	1880 / 1 (5)
Res II	2,3,4,5	1 / 1 (0)	1 / 1 (0)	1880 / 1 (2)

Table 6.4: Second half of the energy grid specifications used in the problems and their tests. Epithermal II encompasses the low-lying resonance of Pu-239. The epithermal regions include upscattering.

Test name	Problem(s)	Epithermal (3×10^0 eV) Ref / BCG (RF)	Epithermal II (6×10^{-1} eV) Ref / BCG (RF)	Thermal (1×10^{-1} eV) Ref / BCG (RF)
Base	1	20 / 2	30 / 4	1 / 1
Base	2,3,4,5	20 / 1	30 / 3	1 / 1
Res I	2,3,4,5	20 / 1 (0.5)	30 / 2 (0.5)	1 / 1 (0)
Res II	2,3,4,5	20 / 20 (0)	30 / 30 (0)	1 / 1 (0)

6.3 Quantities of Interest

Several QOI were investigated, including k-eigenvalue, the power shape, and absorption rates in each spatial region. Errors in these QOI were computed by com-

paring the MG and PG-FEMG solutions to the reference solution.

The power QOI was actually the neutron production rate, our eigenvector. The neutron production rate is related to power produced and may be considered a power proxy. Because the normalization condition for the eigenvector is on its magnitude, only differences in spatial shapes were compared, not the magnitude itself.

The absorption QOI was the absorption rate in each spatial region (e.g., LEU fuel), not component, (e.g., U-235). While it was incapable of measuring nuclide-specific data, and hence incapable of giving detailed information on resonance interference fidelity, it did give this in an averaged sense. The absorption QOI may be considered a depletion proxy.

The three QOI are, mathematically:

$$k_{\text{eff}} = \frac{\int_0^\infty \mathrm{d}E \ \int_D \mathrm{d}^3\mathbf{r} \ \ \nu\Sigma_f\phi}{\int_0^\infty \mathrm{d}E \ \int_D \mathrm{d}^3\mathbf{r} \ \ (\mathrm{div}\mathbf{J} + \Sigma_a\phi)}, \qquad \text{Criticality eigenvalue,} \tag{6.1a}$$

$$P(\mathbf{r}) = \int_0^\infty \mathrm{d}E \ \ \nu\Sigma_f(\mathbf{r},E)\,\phi(\mathbf{r},E), \qquad \text{Power shape, and} \tag{6.1b}$$

$$R_{a,i} = \int_0^\infty \mathrm{d}E \ \int_{D_i} \mathrm{d}^3\mathbf{r} \ \ \Sigma_a(\mathbf{r},E)\,\phi(\mathbf{r},E), \qquad \text{Absorption rates,} \tag{6.1c}$$

where D is the entire domain, D_i is the domain for region i, $\mathbf{J}$ is the net current density (whose contribution is zero due to white BC), and other definitions are given in Section 2.

Relative errors in these QOIs were computed by comparing the MG and PG-FEMG solutions to the reference solution. The power-shape error is defined in terms of spatial integrals using a fine spatial mesh with the linear discontinuous spatial discretization:

$$E_P = \sqrt{\frac{\int_D \mathrm{d}^3\mathbf{r} \ \ \left[P(\mathbf{r})_{\text{PG-FEMG}} - P(\mathbf{r})_{\text{ref}}\right]^2}{\int_D \mathrm{d}^3\mathbf{r} \ \ \left[P(\mathbf{r})_{\text{ref}}\right]^2}}. \tag{6.2}$$

6.4 Verification Problems

Every code must be verified and validated to an application prior to use. Verification was done in two stages. As SCDT was being built, informal unit tests were performed on every newly added component to ensure they worked. This included simple things such as reading in several simple materials and ensuring that they were mixed to yield the macroscopic cross sections correctly, or verifying that the spatial grid was set up correctly.

Having confidence in many of the individual pieces, further verification was done on the completed code using model problems. SCDT was tested in the context of a neutron transport class. Many fixed-source tests were performed, including verification of the second-order accuracy of the DFEM method, verification of balance with and without S2SA, verification of the infinite-medium solution with isotropic and anisotropic scattering, comparison to a diffusive solution, and investigation of boundary layer effects.

Additional verification tests were made for the k-eigenvalue solver. The analytic solution for the homogeneous, infinite-cylinder problem is straightforward to determine and yields k_∞ and its associated eigenvector. This verified the cross section handling and energy matrices. NKA was verified by comparing it to the power iteration results.

No proper validation steps were taken for the k-eigenvalue functionality of SCDT, as the validation would likely reflect more the sensitivity of k-eigenvalue to cross section group structure than to other discretization choices, as energy is often the limiting variable for fidelity. This was deemed acceptable, as SCDT was meant as a research tool to explore numerical methods and not a production-level code that needed to accurately represent the real world. Further, in every case a reference

calculation was performed, allowing a clean evaluation of the PG-FEMG method in isolation.

6.4.1 k-Eigenvalue verification via k-infinity calculations

As there are no known analytic or semi-analytic transport k-eigenvalue solutions for cylindrical systems, spatially-infinite systems were instead used. That is, k-infinity eigenvalues and eigenvectors were investigated.

Three problems were used for verification. All were infinite-medium (reflecting boundary conditions).

The first problem had one group and parameters

$$\Sigma_{t,g} = \{1.0\} \quad \text{cm}^{-1}, \tag{6.3a}$$

$$\chi_g = \{1.0\}, \tag{6.3b}$$

$$\nu\Sigma_{f,g} = \{0.0001\} \quad \text{cm}^{-1}, \tag{6.3c}$$

$$\Sigma_{s,g'\to 1} = \{0.999\} \quad \text{cm}^{-1}. \tag{6.3d}$$

The second problem had two groups and parameters

$$\Sigma_{t,g} = \{10.0, 10.0\} \quad \text{cm}^{-1}, \tag{6.4a}$$

$$\chi_g = \{0.9, 0.1\}, \tag{6.4b}$$

$$\nu\Sigma_{f,g} = \{4.0, 9.0\} \quad \text{cm}^{-1}, \tag{6.4c}$$

$$\Sigma_{s,g'\to 1} = \{5.0, 0.0\} \quad \text{cm}^{-1}, \tag{6.4d}$$

$$\Sigma_{s,g'\to 2} = \{1.0, 5.0\} \quad \text{cm}^{-1}. \tag{6.4e}$$

The third problem had three groups and parameters

$$\Sigma_{t,g} = \{10.0, 10.0, 20.0\} \quad \text{cm}^{-1}, \tag{6.5a}$$

$$\chi_g = \{0.9, 0.1, 0.0\}, \tag{6.5b}$$

$$\nu\Sigma_{f,g} = \{1.0, 2.0, 6.0\} \quad \text{cm}^{-1}, \tag{6.5c}$$

$$\Sigma_{s,g'\to 1} = \{1.0, 0.0, 1.0\} \quad \text{cm}^{-1}, \tag{6.5d}$$

$$\Sigma_{s,g'\to 2} = \{7.0, 1.0, 4.0\} \quad \text{cm}^{-1}, \tag{6.5e}$$

$$\Sigma_{s,g'\to 3} = \{1.0, 6.0, 10.0\} \quad \text{cm}^{-1}. \tag{6.5f}$$

Note that this third problem has up-scattering.

Table 6.5 shows the results of the three problems described above. The tolerance set in SCDT was 10^{-8}. In all cases, the normalization was $\phi_1 = 1$. Errors show that the k-eigenvalue problem was properly implemented, at least in the spatially-independent case.

Runs of SCDT with large radius in mean-free-paths (small leakage) were done to ensure that the eigenvalue and eigenvector were close to their infinite-medium counterparts. They were qualitatively the same and became closer as the radius was increased. Further, the eigenvalue was smaller for the finite-medium case. All these were expected and add qualitative verification to the code.

Table 6.5: Comparison of analytic and SCDT solutions for the infinite medium case.

Problem	Analytic				Error		
	k_∞	ϕ_1	ϕ_2	ϕ_3	k_∞	ϕ_2	ϕ_3
1	0.10000000	1			3.50×10^{-13}		
2	1.22400000	1	0.31111111		$< 10^{-16}$	8.92×10^{-16}	
3	0.95543278	1	1.25833333	0.85500000	4.68×10^{-9}	3.99×10^{-9}	3.52×10^{-9}

6.4.2 Condensation verification

A cross section condensation verification problem was performed for an infinite-medium of H-1, Pu-239 and U-238. Since the reference weighting spectrum was used and there was no spatial variation, the MG and PG-FEMG solutions should exactly reproduce the reference solution when expanded. This was found to be the case to within solver tolerances. Details are provided in the first pages of Section 7.

7. RESULTS

This section has the following structure:

- First division (subsection) $\leftrightarrow$ problem (1, 2, 3, 4, or 5)
- Second division (subsubsection) $\leftrightarrow$ cross section weighting type (reference or generic)
- Third division (subsubsubsection) $\leftrightarrow$ test type (base, resonance-resolution study, or full-resolution study)

In each subsubsubsection, figures are presented that characterize the given runs and their results. Selected component runs of the resolution studies may be found in Appendix C.

Matplotlib (a Python library) was used to generate the figures. Axes can sometimes be confusing with regard to scalings and offsets. The string "`1e-9+1.296`" near the axis means that, for an axis value of y, the true value on the plot is $10^{-9}y + 1.296$; that is, the data is offset by 1.296 and scaled by 10^{-9}.

7.1 Problem 1: Condensation Verification and Homogeneous Pin Cells

Problem 1 contains one region composed of H-1, Pu-239 and U-238. A white boundary was used. Figure 7.1 gives a cartoon of the geometry.

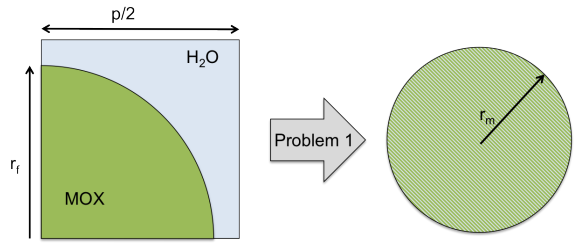

(a) Problem 1

Figure 7.1: Cartoon of the geometry and its simplification to 1D cylindrical geometry for problem 1.

Table 7.1 gives a description of each region for problem 1. Regions are defined as spatially contiguous groups of cells that share a common cross section and are ordered such that larger region numbers correspond to larger radii in the problem. When spatial averaging is done, it is over a region. Unlisted temperatures are 400 K. Tables 6.1 and 6.2 give further geometry and material information, respectively.

Table 7.1: List of spatial regions for problem 1.

Region Number	Description
1	MOX and H-1 slurry

7.1.1 Reference-calculation weighting

In this subsection, the region-averaged reference solution was used to flux-weight the (ultra-fine-group) cross sections during the condensation processes to produce the MG and PG-FEMG cross sections The same DOF were used for both MG and PG-FEMG runs.

7.1.1.1 Base problem

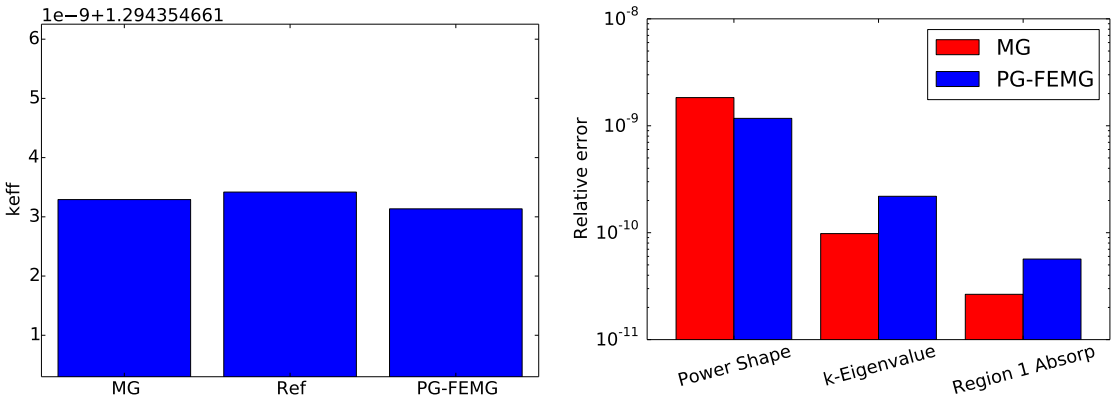

(a) k-Eigenvalues (b) Relative Errors

Figure 7.2: Comparison of eigenvalues and relative errors for problem 1 with high-fidelity weight. Values are given for the reference, MG and PG-FEMG methods.

Fig. 7.2 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. 7.2a gives the k-eigenvalues for each of the calculations. Fig. 7.2b gives the error in k-eigenvalue, power shape, and absorption rates. All errors are within iterative convergence tolerances. This is expected for the infinite-medium problem without spatial or angular dependence and with reference weighting.

Fig. 7.3a gives the energy dependence of the reference solution source. As the so-

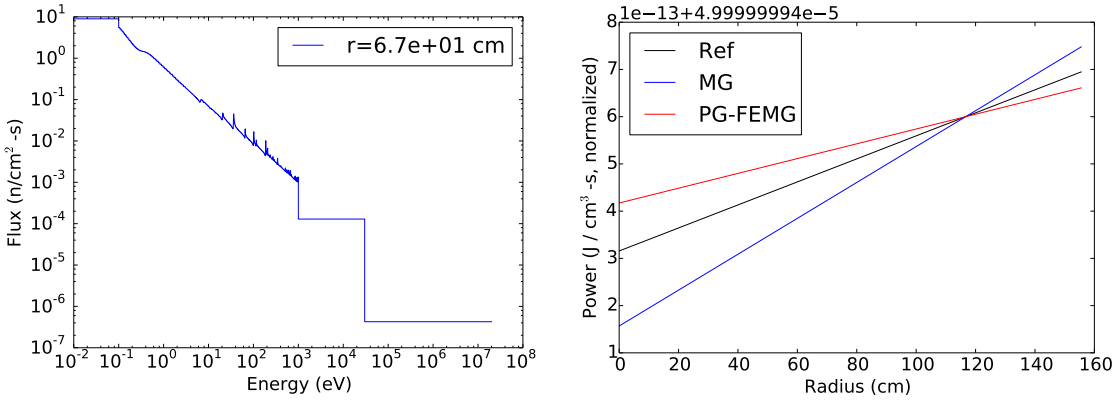

(a) Reference solution neutron source

(b) Power profile

Figure 7.3: Total source and power profile as functions of position for problem 1 with high-fidelity weight. (a) Total source (scattering plus fission) for the reference solution as a function of energy for select radial points. (b) Power profile as a function of position for the reference, MG and PG-FEMG solutions.

lution to the first problem has no spatial or angular dependence, this is the spectrum for all positions and angles.

Fig. 7.3b gives the power profile shapes in space for the reference, MG, and PG-FEMG solutions. They are all flat to within iterative convergence tolerance.

Fig. 7.4a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. 7.4b gives the effective total cross section and the weighting spectrum used. For this calculation, a high-fidelity weighting spectrum was used (note resonances in b).

7.1.2 Generic weighting

In this subsection, a generic, low-fidelity solution was used to weight the cross sections during the condensation process, namely the iwt=5 option in NJOY (mid-life PWR spectrum with O-16 resonances). The same group structure was used as

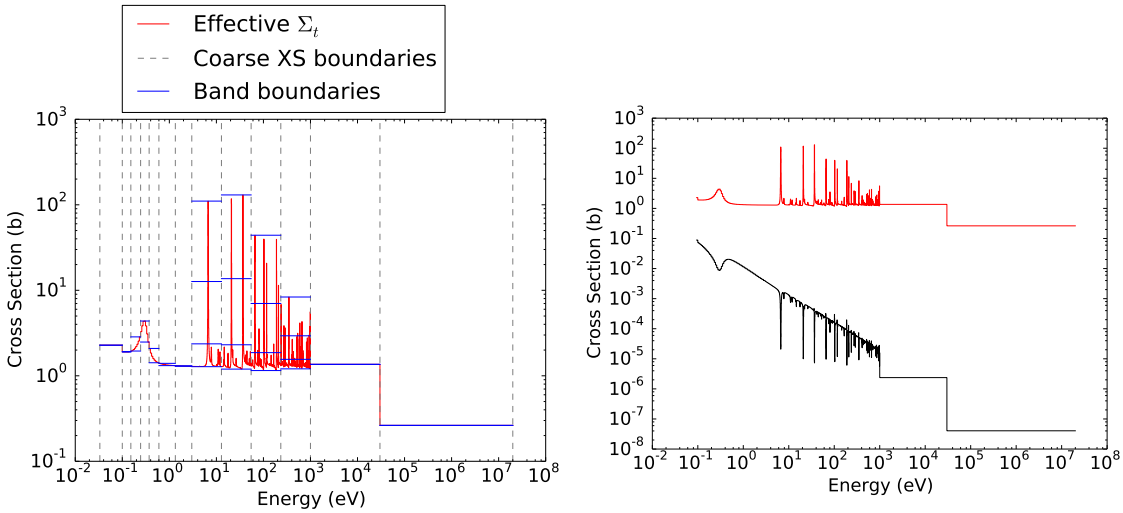

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure 7.4: Effective total cross sections for the PG-FEMG band boundary calculation for problem 1 with high-fidelity weight.

in the reference-weighted results.

7.1.2.1 Base problem

Fig. 7.5 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. 7.5a gives the k-eigenvalues for each of the calculations. Fig. 7.5b gives the error in k-eigenvalue, power shape, and absorption rates. Both PG-FEMG and MG capture the flat spatial shape of the power in this infinite-medium problem. The k-eigenvalue and absorption rate errors are smaller for PG-FEMG than MG by over a factor of 10.

Fig. 7.6a gives the energy dependence of the reference solution source. As the solution to the first problem has no spatial or angular dependence, this is the spectrum for all positions and angles.

Fig. 7.6b gives the power profile shapes in space for the reference, MG, and PG-FEMG solutions. They are all flat to within iterative convergence tolerance.

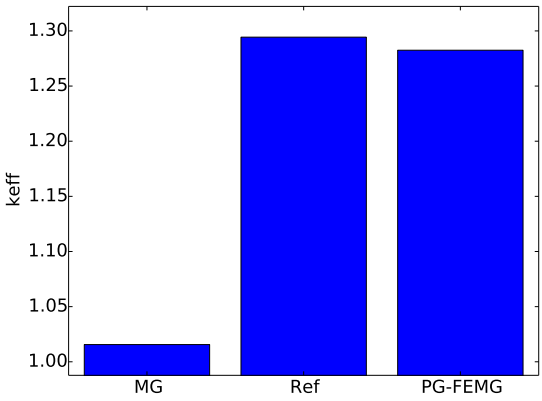

(a) k-Eigenvalues

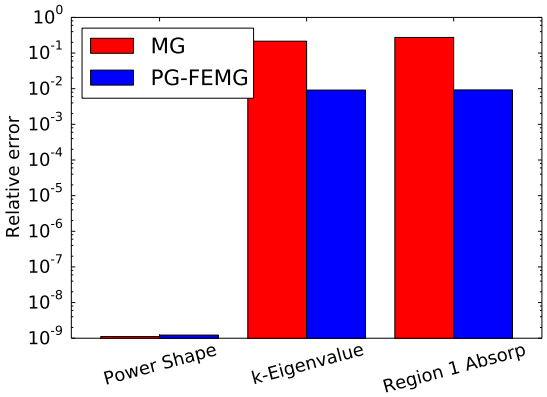

(b) Relative Errors

Figure 7.5: Comparison of eigenvalues and relative errors for problem 1 with low-fidelity weight. Values are given for the reference, MG and PG-FEMG methods.

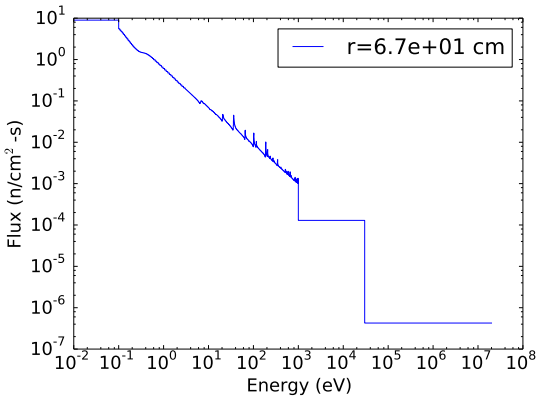

(a) Reference solution neutron source

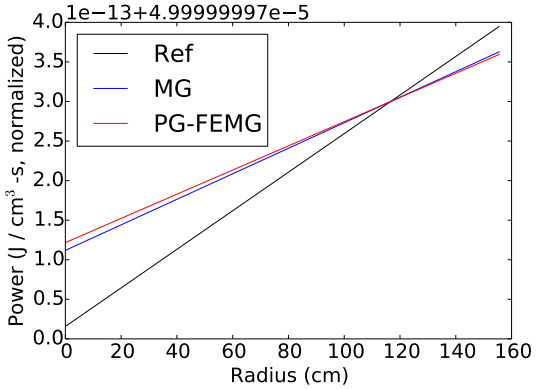

(b) Power profile

Figure 7.6: Total source and power profile as functions of position for problem 1 with low-fidelity weight. (a) Total source (scattering plus fission) for the reference solution as a function of energy for select radial points. (b) Power profile as a function of position for the reference, MG and PG-FEMG solutions.

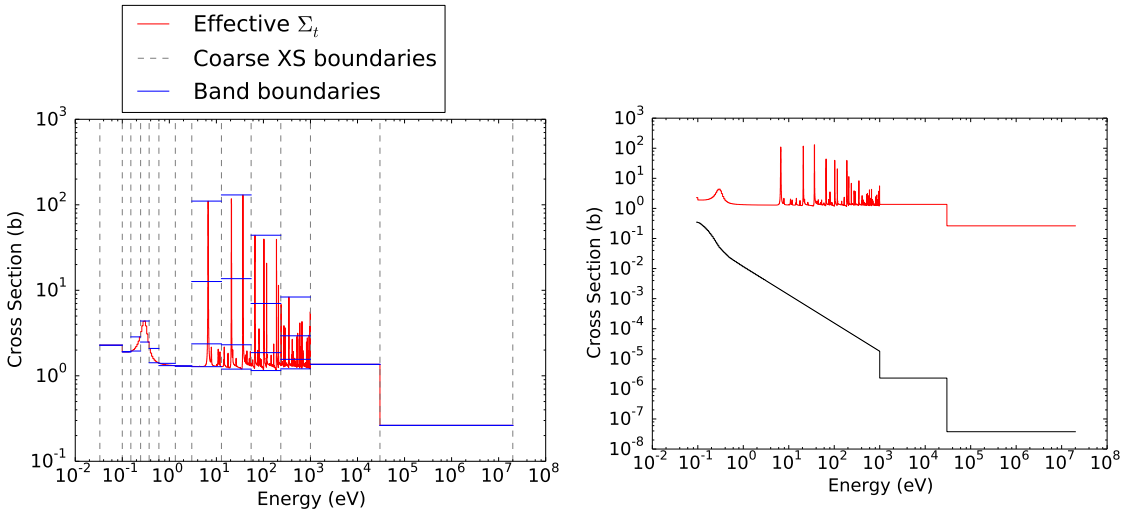

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure 7.7: Effective total cross sections for the PG-FEMG band boundary calculation for problem 1 with low-fidelity weight.

Fig. 7.7a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. 7.7b gives the effective total cross section and the weighting spectrum used. For this calculation, a low-fidelity weighting spectrum was used (note smoothness in b).

7.2 Problem 2: Heterogeneous, 1-T Pin Cells

Problem 2 is a simplified two-region, one-temperature PWR fuel-cell consisting of fuel (U-238, Pu-239, and O-16), and moderator (O-16 and H-1). A white boundary was used. Figure 7.8 gives a cartoon of the geometry.

Table 7.2 gives a description of each region for problem 2. Regions are defined as spatially contiguous groups of cells that share a common cross section and are ordered such that larger region numbers correspond to larger radii in the problem. When spatial averaging is done, it is over a region. Unlisted temperatures are 400

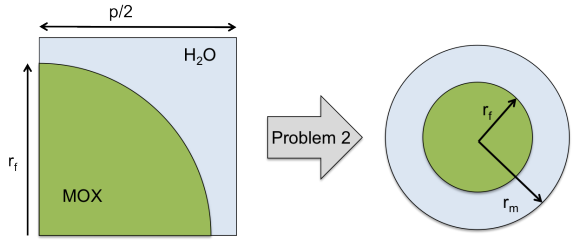

(a) Problem 2

Figure 7.8: Cartoon of the geometry and its simplification to 1D cylindrical geometry for problem 2.

K. Tables 6.1 and 6.2 give further geometry and material information, respectively.

Table 7.2: List of spatial regions for problem 2.

Region Number	Description
1	MOX
2	Light water

7.2.1 Reference-calculation weighting

In this subsection, the region-averaged reference solution was used to flux-weight the (ultra-fine-group) cross sections during the condensation processes to produce the MG and PG-FEMG cross sections The same DOF were used for both MG and PG-FEMG runs.

7.2.1.1 Base problem

Fig. 7.9 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. 7.9a gives the k-eigenvalues for each of the calculations. Fig. 7.9b gives the error in k-eigenvalue, power shape, and absorption rates. Both

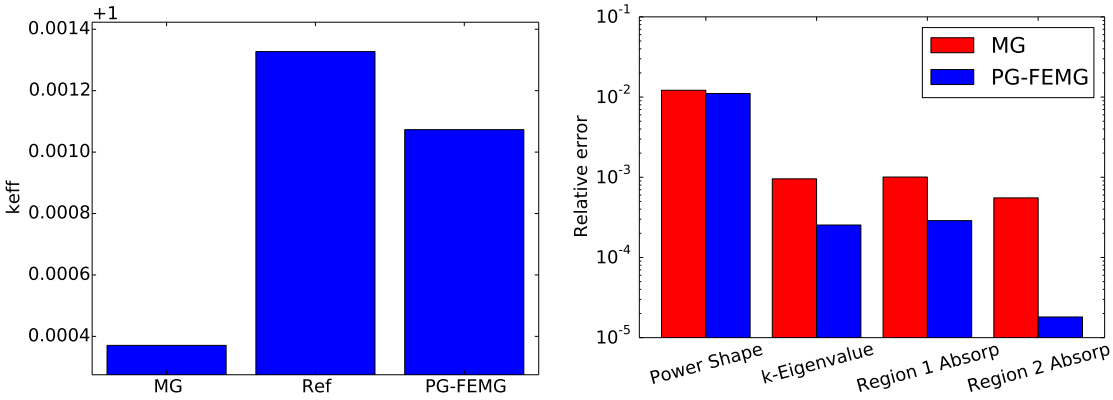

(a) k-Eigenvalues (b) Relative Errors

Figure 7.9: Comparison of eigenvalues and relative errors for problem 2 with high-fidelity weight. Values are given for the reference, MG and PG-FEMG methods.

PG-FEMG and MG do approximately equally well with the power shape error. PG-FEMG has a lower error for the other QOI, beating MG by factors of approximately 3, 3, and 30 for k-eigenvalue, absorption in the fuel, and absorption in the moderator, respectively.

Fig. 7.10a gives the energy dependence of the reference solution source at several equally-spaced radial positions. It shows the source to be affected by the resonances. As the radius is increased, the source becomes smoother.

Fig. 7.10b gives the power profile shapes in space for the reference, MG, and PG-FEMG solutions. Note that the normalization was on this parameter, so they all have the same integral, but may have different spatial shapes. In both cases, the reference solution shows stronger self-shielding effects near the edge of the fuel. The MG and PG-FEMG shapes track the reference shape well until the outside of the fuel, where the reference shows a sharper slope. The error near the edge of the fuel is approximately 4% for both cases.

Fig. 7.11 (ff.) shows the energy dependence of the reference, MG and PG-FEMG

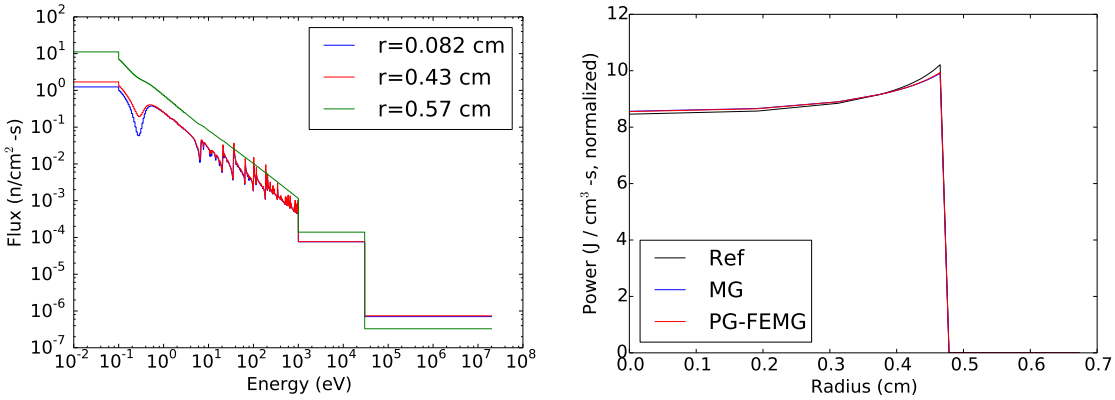

(a) Reference solution neutron source

(b) Power profile

Figure 7.10: Total source and power profile as functions of position for problem 2 with high-fidelity weight. (a) Total source (scattering plus fission) for the reference solution as a function of energy for select radial points. (b) Power profile as a function of position for the reference, MG and PG-FEMG solutions.

solutions averaged over spatial regions and for the exiting partial current. They show the detailed resonance structure captured by the reference solution and the fidelity of the MG and PG-FEMG reconstructions. Note that the flux dips are space- and direction-dependent. Since a reference solution was used for both cross section condensation and for flux reexpansion, we expect and find both the MG and PG-FEMG solutions to have region-averaged reconstructions accurate to the reference solution.

Fig. 7.12 (ff.) shows the energy dependence of the reference, MG and PG-FEMG solutions at selected spatial regions in the fuel. They show the detailed resonance structure captured by the reference solution and the fidelity of the MG and PG-FEMG reconstructions. Even though a reference solution was used for both cross section condensation and for flux reexpansion, the within-group MG spectrum cannot adapt in space within a region and can only be accurate in at best an average sense.

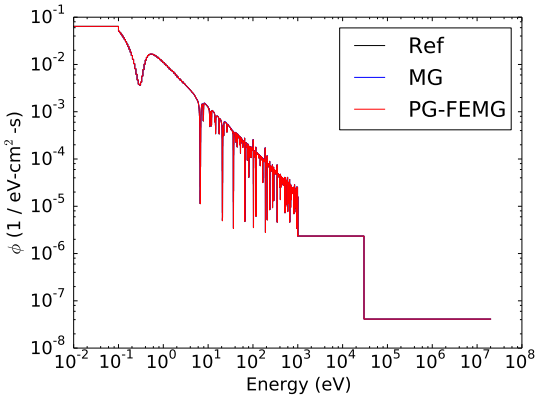

(a) Region 1 average flux

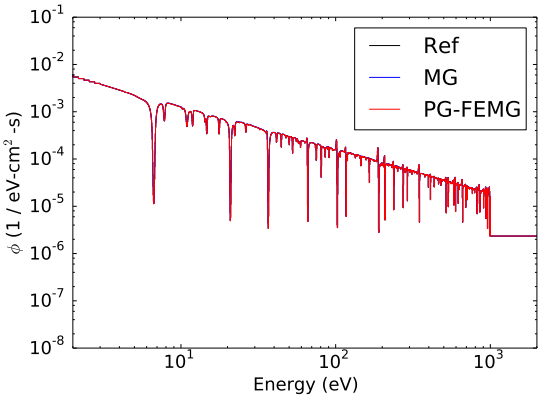

(b) Region 1 average flux (zoom)

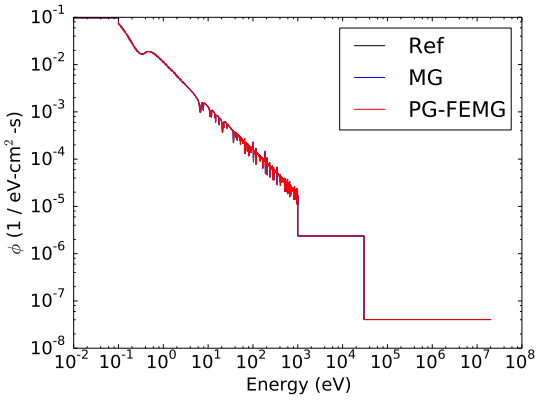

(c) Region 2 average flux

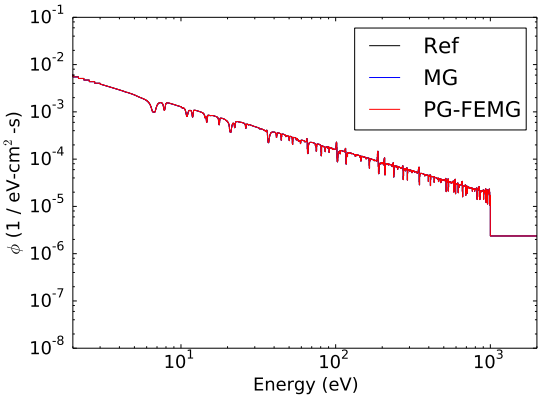

(d) Region 2 average flux (zoom)

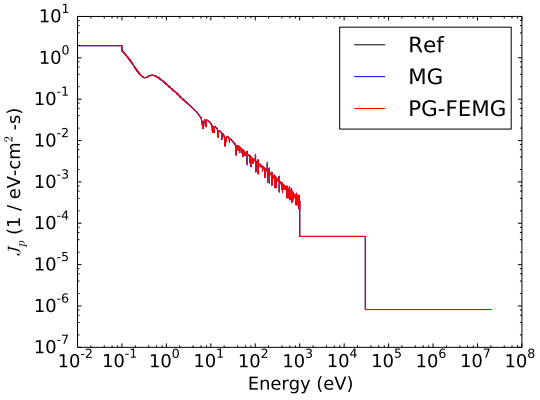

(e) Outgoing partial current

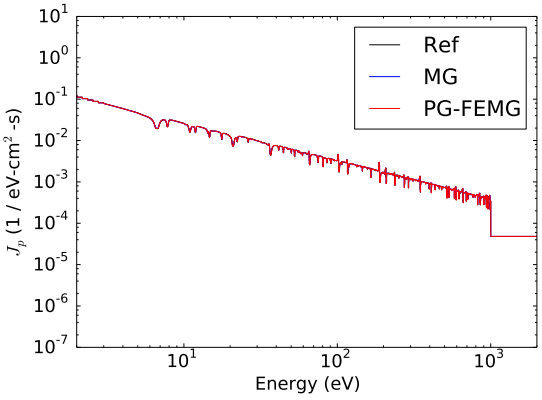

(f) Outgoing partial current (zoom)

Figure 7.11: Energy spectra averaged over spatial regions and at problem boundary for problem 2 with high-fidelity weight.

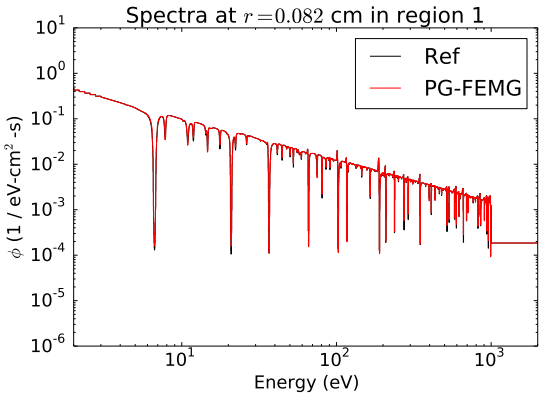

(a) Flux at point 1 for PG-FEMG

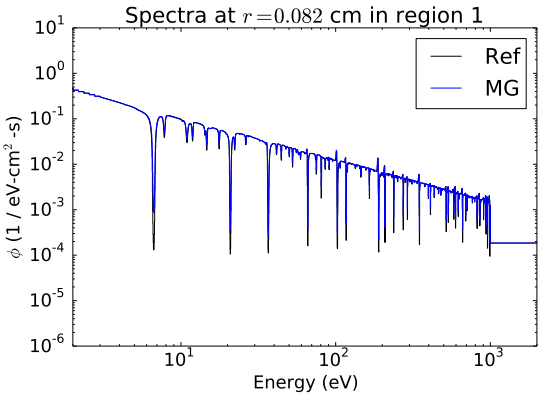

(b) Flux at point 3 for MG

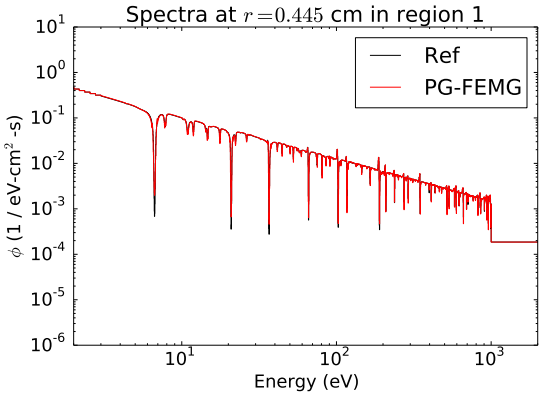

(c) Flux at point 5 for PG-FEMG

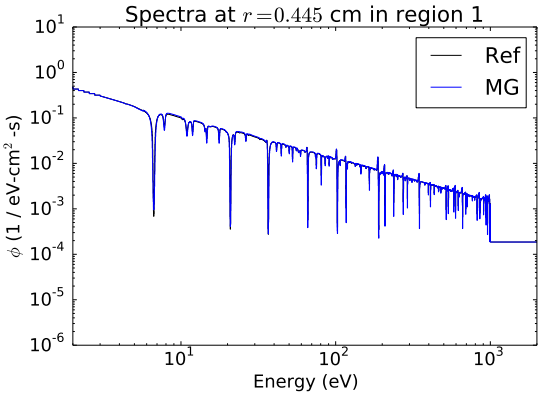

(d) Flux at point 7 for MG

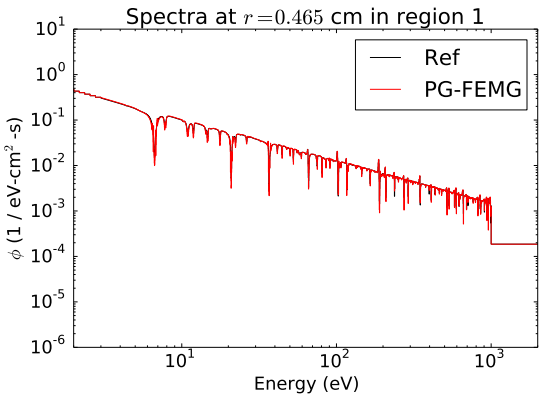

(e) Flux at point 9 for PG-FEMG

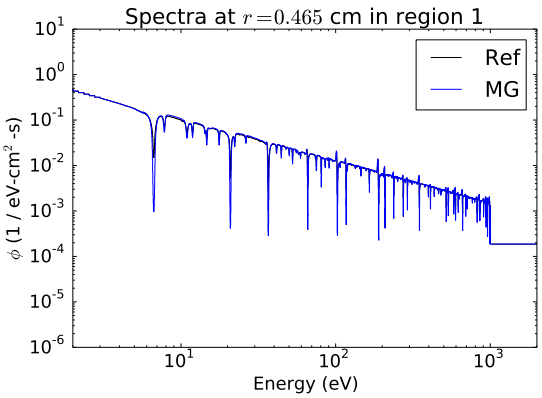

(f) Flux at point 11 for MG

Figure 7.12: Energy spectra at selected spatial points for problem 2 with high-fidelity weight.

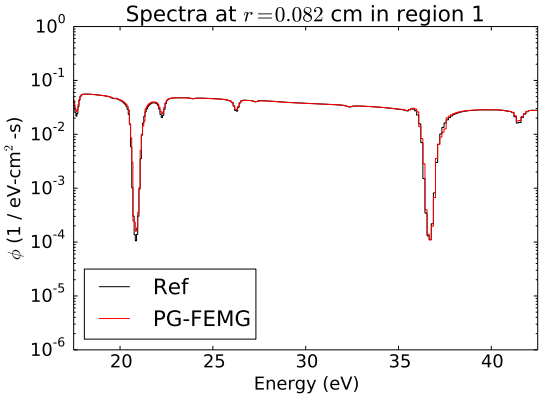

(a) Flux at point 2 for PG-FEMG (zoom)

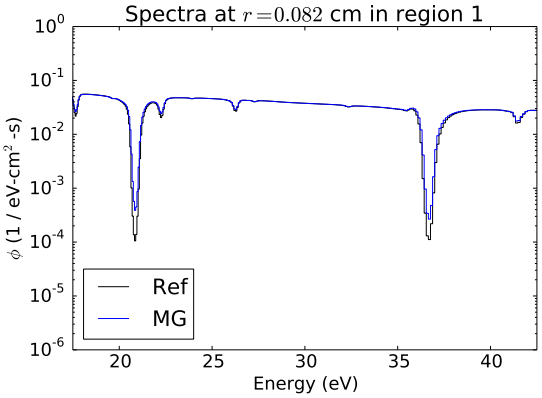

(b) Flux at point 4 for MG (zoom)

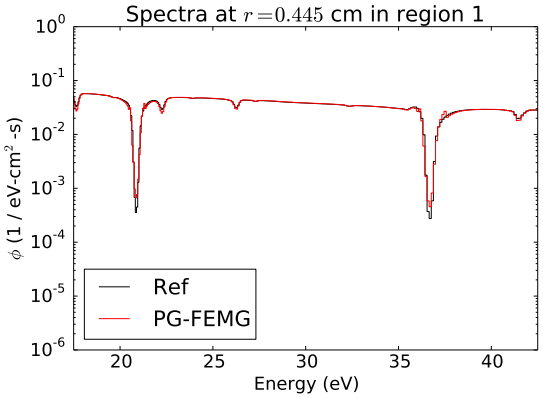

(c) Flux at point 6 for PG-FEMG (zoom)

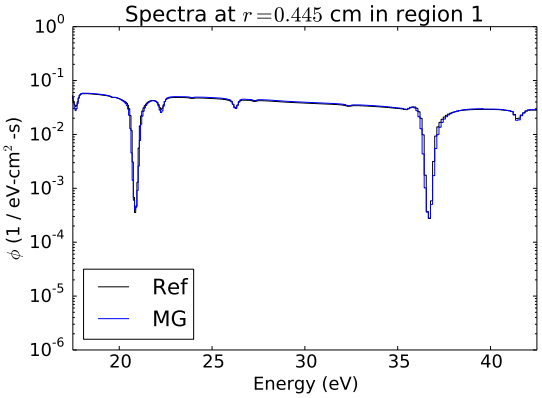

(d) Flux at point 8 for MG (zoom)

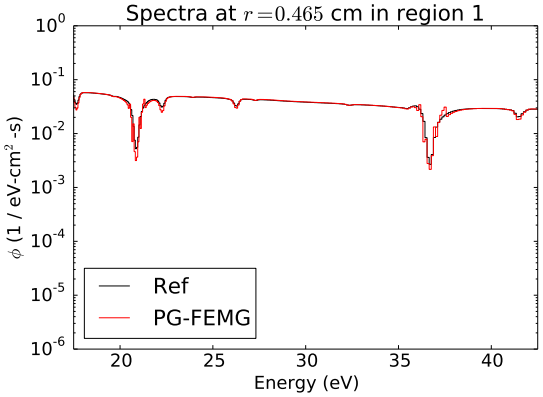

(e) Flux at point 10 for PG-FEMG (zoom)

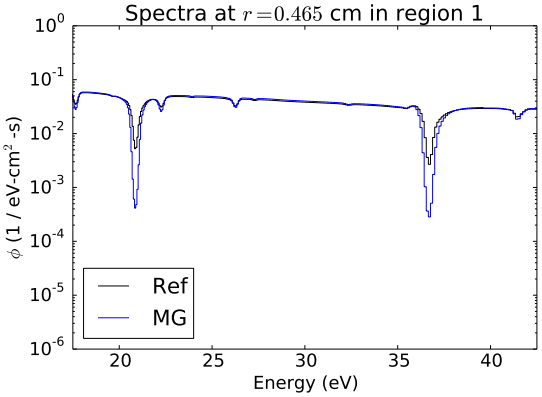

(f) Flux at point 12 for MG (zoom)

Figure 7.13: Energy spectra at selected spatial points for problem 2 with high-fidelity weight (cont.).

The PG-FEMG method has DOF at the resonances and so its spectrum can locally adapt to the correct shape insofar as the DOF capture the varying portions of the spectrum.

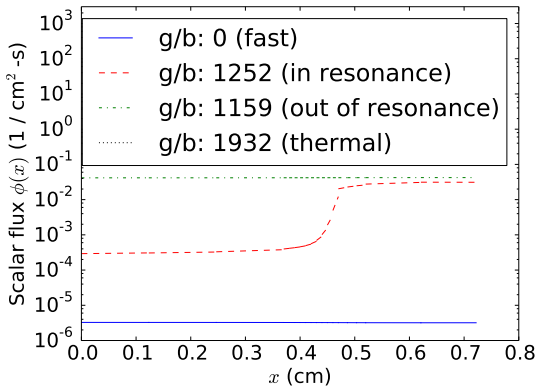

(a) Reference

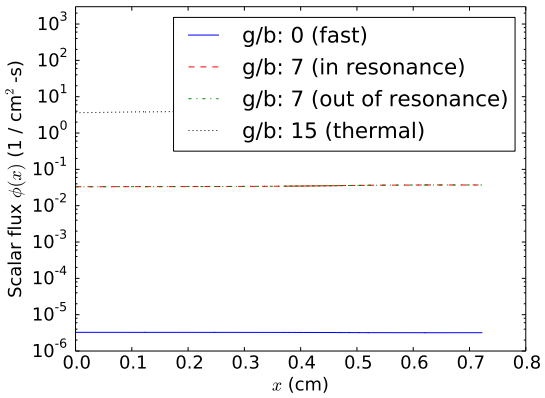

(b) MG

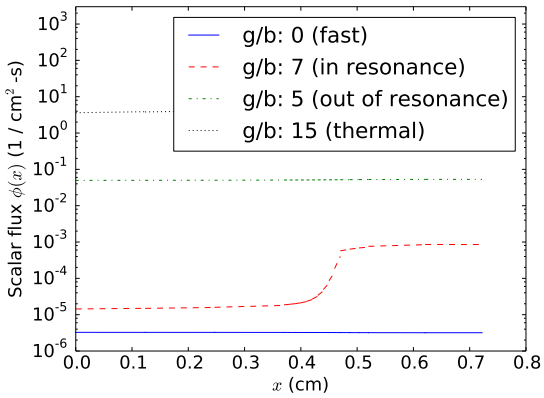

(c) PG-FEMG

Figure 7.14: Spatial dependence for several groups of interest for problem 2 with high-fidelity weight. The group corresponds to the same energy values in all cases. Since no expansion has been done on the multigroup or multiband fluxes, they should only be compared qualitatively to the reference fluxes.

Fig. 7.14 gives the unprocessed group / band fluxes as functions of position for the reference, MG, and PG-FEMG solutions. A set of representative energies were

used, the most interesting of which were the in- and out-of- resonace cases. These show the effect of adding DOF to the resonances with PG-FEMG, where fidelity is preserved, or MG, where features are lost. Note that boundary layers are preserved with PG-FEMG.

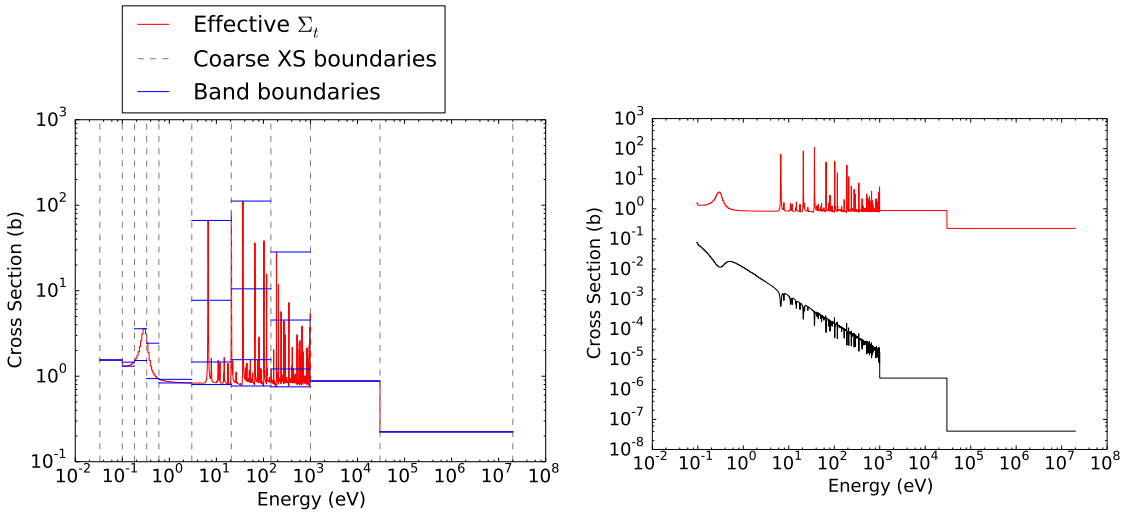

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure 7.15: Effective total cross sections for the PG-FEMG band boundary calculation for problem 2 with high-fidelity weight.

Fig. 7.15a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. 7.15b gives the effective total cross section and the weighting spectrum used. For this calculation, a high-fidelity weighting spectrum was used (note resonances in b).

Fig. 7.16 gives the group-to-energy mappings for the MG (a) and reference (b) solutions. Notice the dramatic reduction in DOF from reference to MG and recall MG and PG-FEMG have the same DOF. As it does not make sense to map band

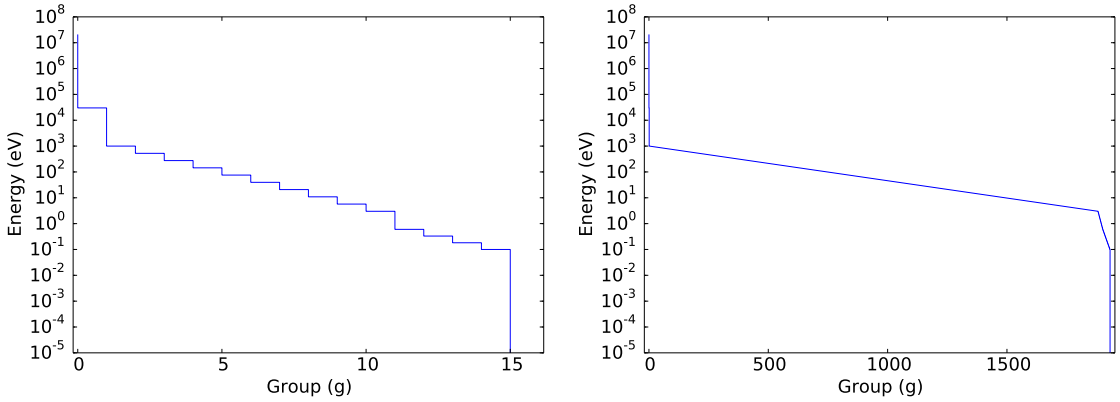

(a) Multigroup group mapping

(b) Reference group mapping

Figure 7.16: Group boundaries and mappings for problem 2 with high-fidelity weight.

boundaries to energies, no figure is provided for the PG-FEMG case.

Fig. 7.17 ff. give the PG-FEMG (MG, left) and MG (right) cross sections as functions of energy. For the PG-FEMG cases, there may be more than one XS value per coarse group. In an attempt to demonstrate this, lines at each of the band values were given with light shading between them to indicate the cross sections range (cf. the picket fence idea of bands).

Fig. 7.19 ff. give the sparsity pattern of the scattering kernel when represented as a matrix with rows corresponding to out-going group and columns corresponding to incoming group. Notice how PG-FEMG creates artificial upscatter in the resonance region, while standard MG is lower-triangular for all but the thermal energies. Ordering is standard: a high group number implies a low energy. Within the figures, brighter colors indicate larger magnitudes. Grey is within-group scattering, blue is down-scattering, and red is up-scattering. Dotted lines, if any, indicate how many bands belong to a PG-FEMG coarse group. Blue and red colors are relative to the given figure; grey colors can be compared between figures, as they indicate the

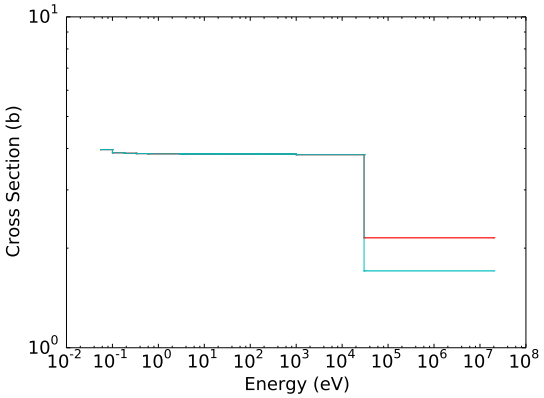

(a) Cross sections for O-16 (PG-FEMG)

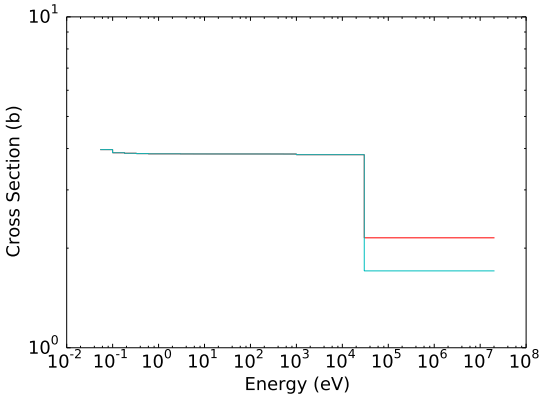

(b) Cross sections for O-16 (MG)

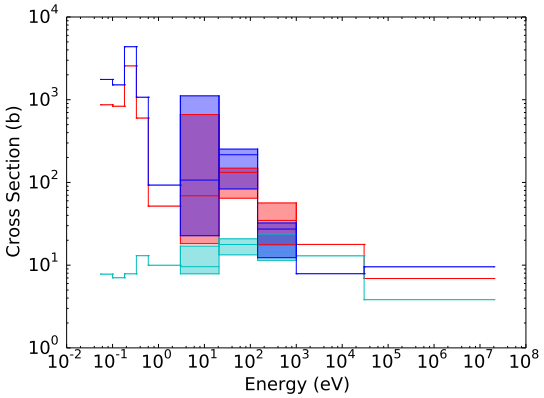

(c) Cross sections for Pu-239 (PG-FEMG)

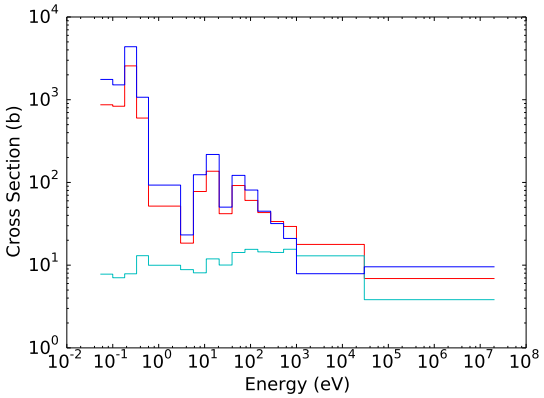

(d) Cross sections for Pu-239 (MG)

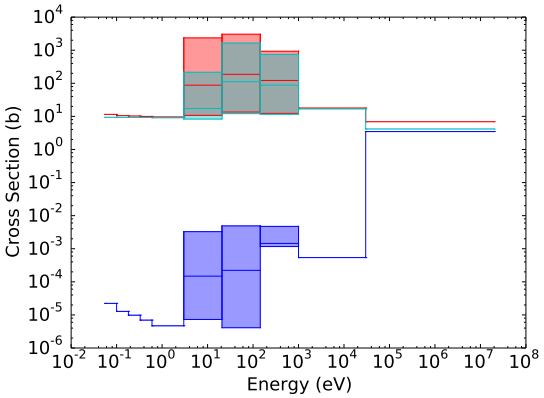

(e) Cross sections for U-238 (PG-FEMG)

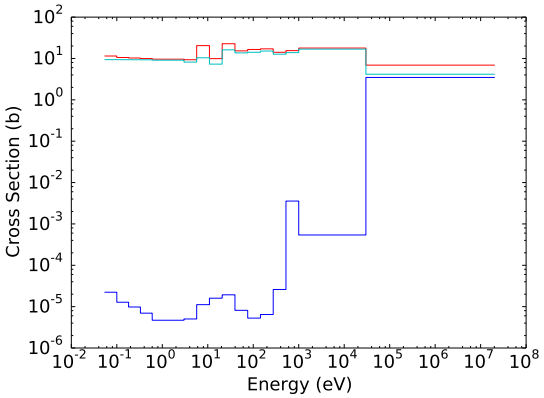

(f) Cross sections for U-238 (MG)

Figure 7.17: Various cross sections as functions of energy for region 1 and problem 2 with high-fidelity weight. Absorption (σ_a), total (σ_t), and fission production ($\nu\sigma_f$) cross sections are given.

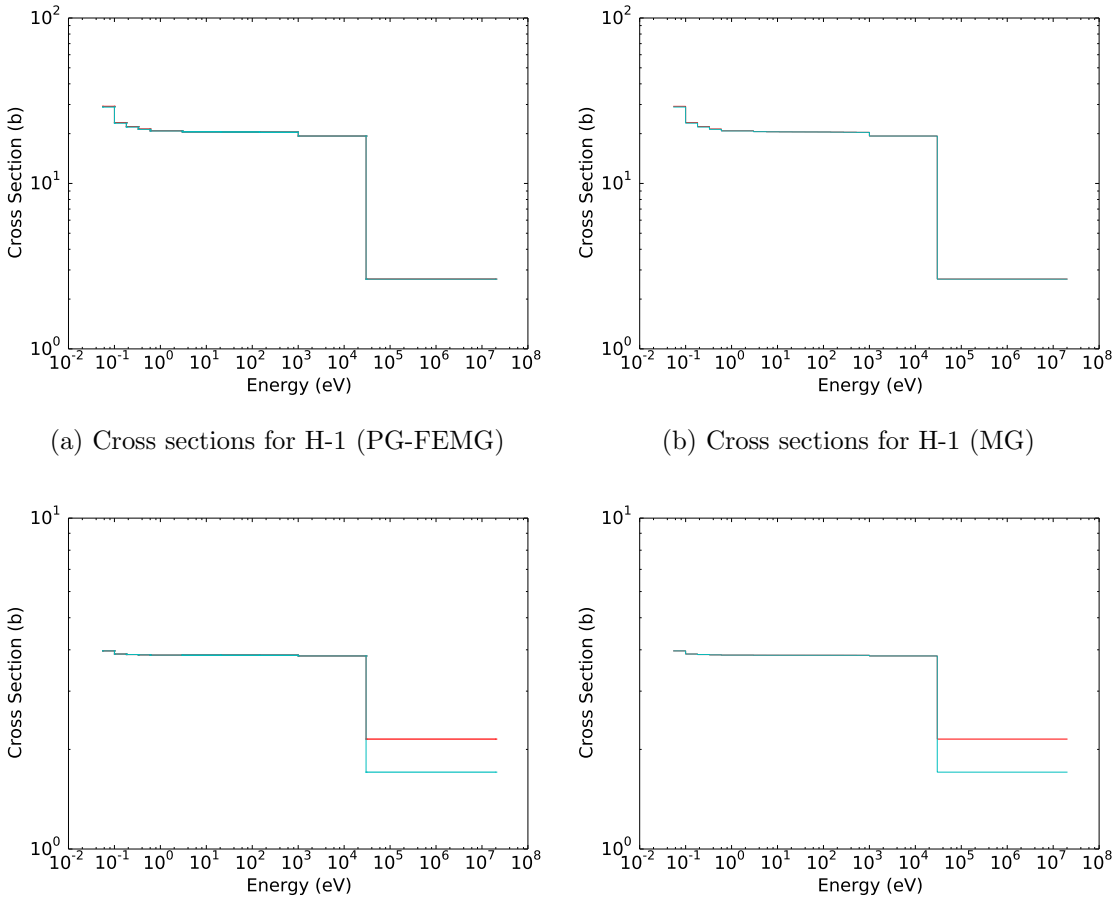

(a) Cross sections for H-1 (PG-FEMG)

(b) Cross sections for H-1 (MG)

(c) Cross sections for O-16 (PG-FEMG)

(d) Cross sections for O-16 (MG)

Figure 7.18: Various cross sections as functions of energy for region 2 and problem 2 with high-fidelity weight. Absorption (σ_a), total (σ_t), and fission production ($\nu\sigma_f$) cross sections are given.

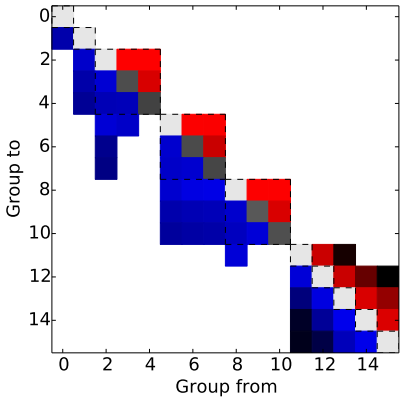

(a) Scattering matrix for O-16 (PG-FEMG)

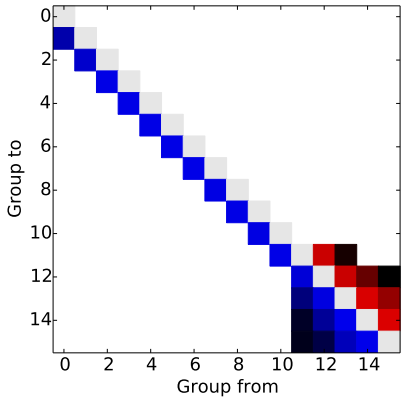

(b) Scattering matrix for O-16 (MG)

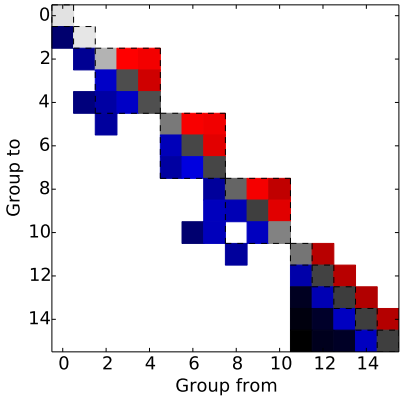

(c) Scattering matrix for Pu-239 (PG-FEMG)

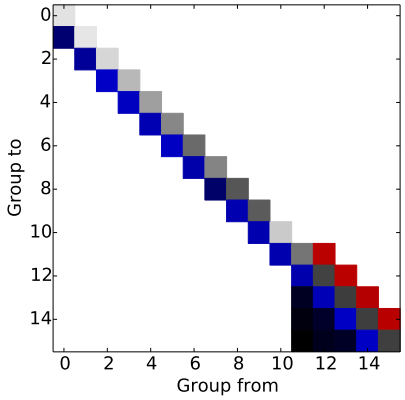

(d) Scattering matrix for Pu-239 (MG)

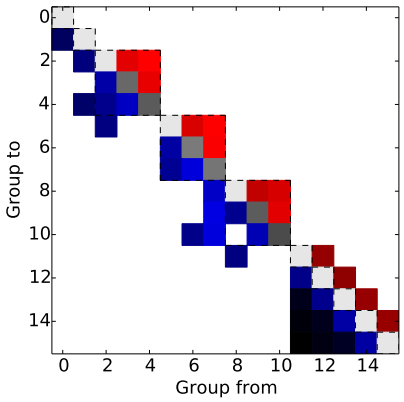

(e) Scattering matrix for U-238 (PG-FEMG)

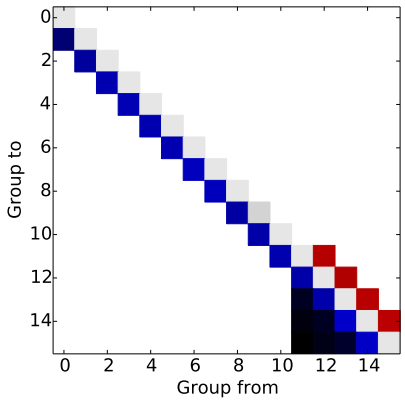

(f) Scattering matrix for U-238 (MG)

Figure 7.19: Scattering matrices for region 1 and problem 2 with high-fidelity weight. PG-FEMG matrices (left) are compared to MG matrices (right).

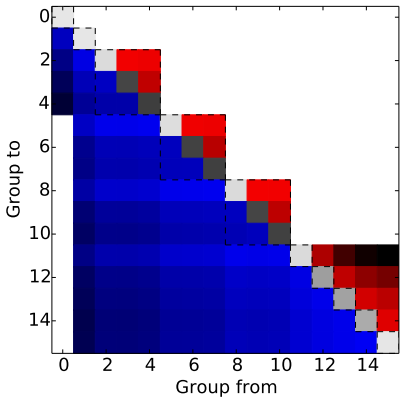

(a) Scattering matrix for H-1 (PG-FEMG)

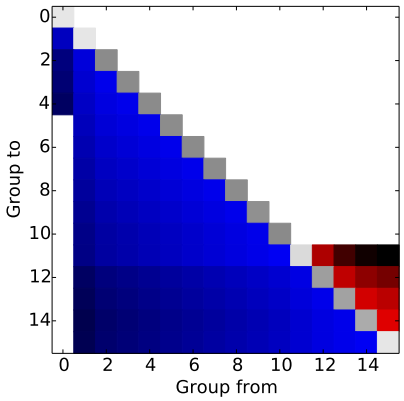

(b) Scattering matrix for H-1 (MG)

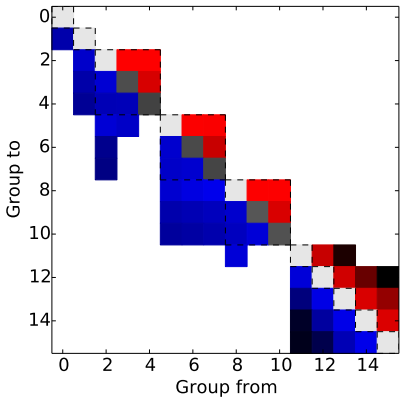

(c) Scattering matrix for O-16 (PG-FEMG)

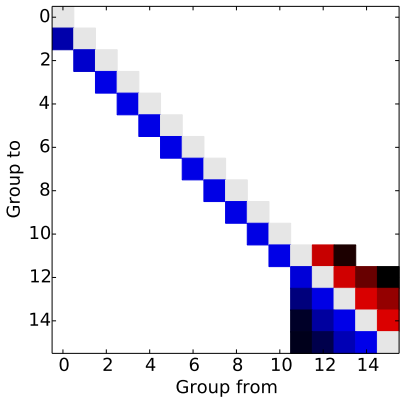

(d) Scattering matrix for O-16 (MG)

Figure 7.20: Scattering matrices for region 2 and problem 2 with high-fidelity weight. PG-FEMG matrices (left) are compared to MG matrices (right).

within-group scattering ratio.

7.2.1.2 Resonance-only energy resolution study

An energy resolution study was done to determine how various QOIs changed as a function of increasing energy fidelity (DOF) for moderate DOF. The energy group structure was always fully resolved in the non-resonance regions and changed only in the resolved-resonance region. All errors are relative comparisons between the reference solution and the MG or PG-FEMG solutions.

Figure 7.21 gives results from the energy resolution study. The number of coarse groups and the bands per coarse group in the resonance region were varied to determine their effect on error of the QOI and their effect on required work to solve the system, as measured in sweeps. The abscissa is DOF, which for MG is the number of groups and for PG-FEMG is the sum of the bands per group over all groups. This resolution study resolved (to the reference level) all non-resonance energies and used varying yet moderate numbers of DOF within the resolved resonance region.

Figure 7.21a-c give work measurements to answer the questions: how efficient is PG-FEMG compared to standard MG, what are the work savings compared to the reference solution, and how do band structures affect sweeps per DOF. The required sweeps per DOF is proportional to the difficulty of inverting the scattering matrix. This matrix had coupled within-coarse-group scattering when PG-FEMG was used.

All runs were done with the NKA solver. Up to 3 steps of the traditional iteration scheme — with its requisite inner iterations, including up to 2 iterations of upscattering per (outer) power iteration — were used as an initial condition for the NKA solver. Failure to do this initialization often resulted in poor or non-convergent NKA performance. This should be expected, as nonlinear solvers are only guaranteed to work in the neighborhood of the converged solution. We initialized NKA with power

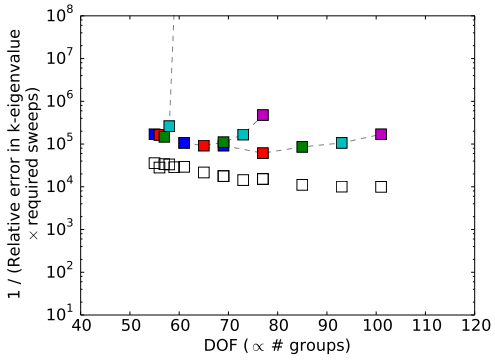

(a) k-Eigenvalue efficiency

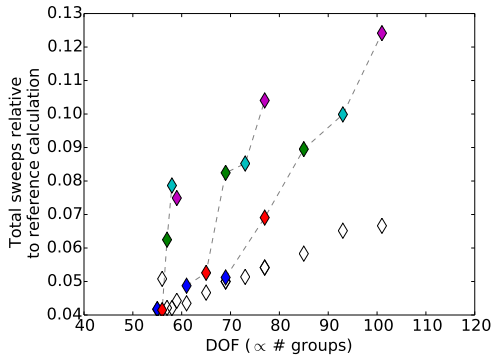

(b) Relative transport sweeps summed over bands / groups (work measurement)

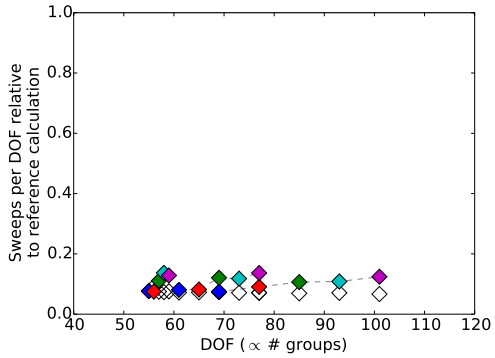

(c) Relative transport sweeps per DOF

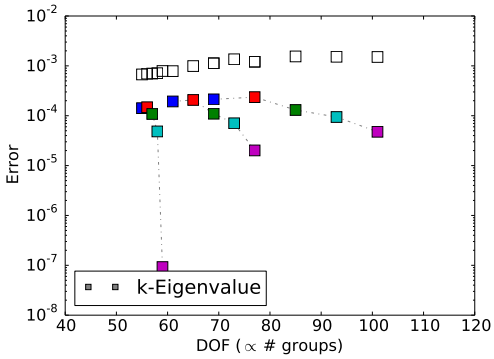

(d) k-Eigenvalue error

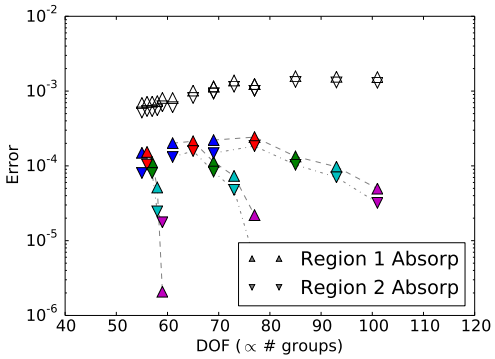

(e) Region absorption rate errors

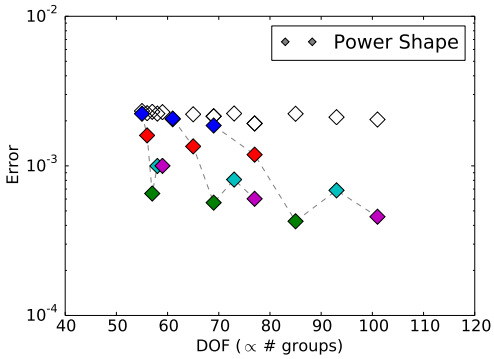

(f) Power shape

Figure 7.21: Error in various QOI relative to the reference solution for several energy group structures and numbers of bands for problem 2 with high-fidelity weight, for the resonance-only energy resolution study. The MG solution corresponds to unfilled markers; the PG-FEMG solutions are colored according to how many bands per group were used in the resonance region, be that two, three, four, five, or six. Dotted lines connect multiband cases with the same coarse group structure.

iteration for all problems studied.

Figure 7.21a shows efficiency using the figure of merit 1/(work *times* error). Higher figures of merit indicate higher efficiencies. Figure 7.21b gives work per (number of) DOF, again relative to the reference solution. Figure 7.21b shows PG-FEMG required approximately the same amount of work per DOF as MG. Both MG and PG-FEMG required substantially less work per DOF than the reference solution (horizontal line at 1.0). Figure 7.21c shows total work scaled with the number of groups, which was expected, as we defined one sweep as the inversion of the streaming-plus-total-collision operator for one group; the full inversion of the left-hand side would require a number of sweeps equal to the number of groups. The PG-FEMG method did indeed take more total work than the MG method, but not by a wide margin, less than a factor of two. Note all work was scaled relative to the reference solution. (It takes MG and PG-FEMG DOF on the order of reference group DOF to have MG or PG-FEMG to be more expensive than the reference solution.)

Figure 7.21d-f show how error in the QOI was affected by changing energy resolution. This resolution study did not use the extreme numbers of DOF that were used in the full energy resolution study and thus the MG energy resolution did not enter the asymptotic regime. For this reason, errors stayed approximately constant as the number of groups were increased. For the k-eigenvalue calculation, going from 1 band per group (MG) to 2 bands per group decreased the error by a factor of 5. Increasing from 2 bands per group to 6 bands per group decreased the same error by another factor of 5-2000. Using one group for the entire resonance region produced the factor of 2000. Similar decreases occured for the absorption rates. Increasing from 1 to 2 bands per group had negligible influence on the power shape error. Increasing from 2 to 6 bands per group decreased the same error by an additional factor of 2.5-8 depending on group structure. Increasing bands per group monotonically

decreased QOI error except for 5 bands per group and the power shape error, which always increased. It is difficult to point to a specific cause for this behavior, but it may be due to poor band boundaries.

7.2.1.3 Energy resolution study

An energy resolution study was done to determine how various QOIs changed as a function of increasing energy fidelity (DOF) for large DOF. The energy group structure was resolved in all energy regions for this study. All errors are relative comparisons between the reference solution and the MG or PG-FEMG solutions.

Figure 7.22 gives the results from the energy resolution study. The number of groups and bands per group in the resolved resonance region were varied to determine their effect on QOI error. This resolution study focused on extremely large DOF to investigate asymptotic effects.

Figure 7.22a-c give work measurements. Figure 7.22a gives efficiency of the k-eigenvalue calculation (1/(work $\times$ error)), Figure 7.22b gives work per DOF, and Figure 7.22c shows total work scaled with the number of groups. Figure 7.22b shows PG-FEMG required more work per DOF than MG or the reference solution (horizontal line) in general, but this ratio was approximately constant over all numbers of DOF considered.

Figure 7.22d-f show how error in the QOI was affected by changing energy resolution. They give a sense of the extreme number of groups required to enter the asymptotic regime when only part of the resolved resonance region was included, to say nothing of the high-energy end of the resolved region where even tighter group spacing would be required. Increasing from 1 band per group (MG) to 2 bands per group decreased the error by a factor of approximately 3-7 for k-eigenvalue and absorption rates, depending on group structure. Increasing from 2 to 4 bands per group

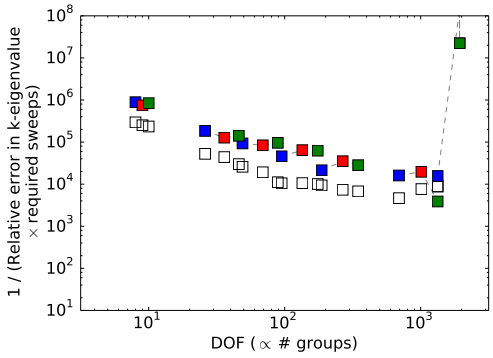

(a) k-Eigenvalue efficiency

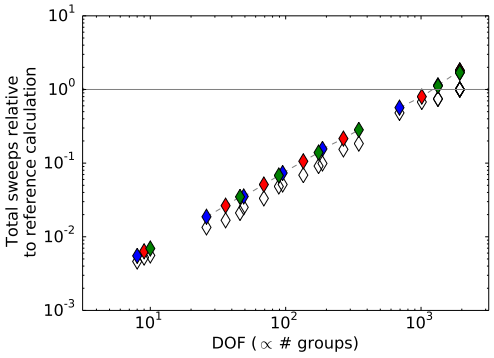

(b) Relative transport sweeps summed over bands / groups (work measurement)

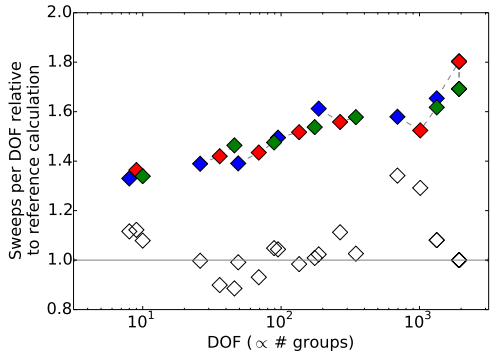

(c) Relative transport sweeps per DOF

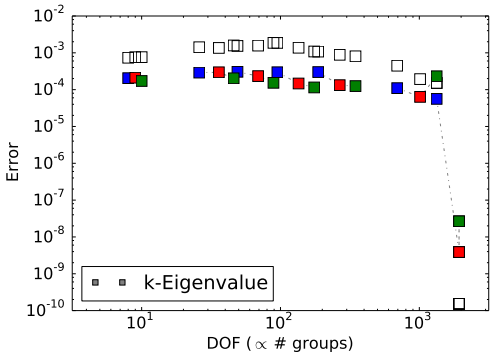

(d) k-Eigenvalue error

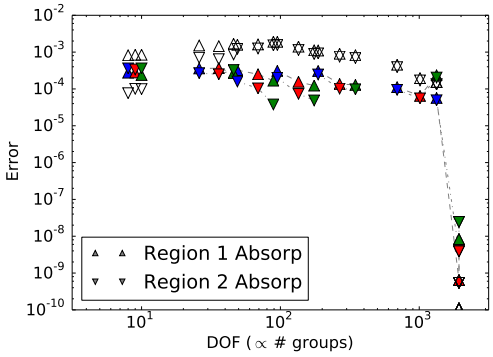

(e) Region absorption rate errors

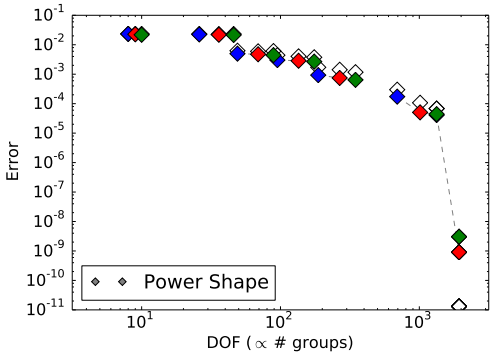

(f) Power shape

Figure 7.22: Error in various QOI relative to the reference solution for several energy group structures and numbers of bands for problem 2 with high-fidelity weight, for the full energy resolution study. The MG solution corresponds to unfilled markers; the PG-FEMG solutions are colored according to how many bands per group were used in the resonance region, be that two, three, four, five, or six. Dotted lines connect multiband cases with the same coarse group structure.

decreased error by another factor of 1.5 to 3. Only for very large DOF count (200+) did increasing the bands per group lower power shape error. This may be due to using a high-fidelity weighting spectrum for the MG. MG has a very accurate absorption QOI in the moderator at low group counts that disappears for higher group counts. The increasing error with increasing bands per groups behavior at very large DOF count ($\sim$ 1000 DOF) for multiple bands per group was caused by the fall-back band boundary calculation which produced poor band boundaries. Future work may include better band boundary algorithms when the DOF are close to the reference DOF.

7.2.2 Generic weighting

In this subsection, a generic, low-fidelity solution was used to weight the cross sections during the condensation process, namely the iwt=5 option in NJOY (mid-life PWR spectrum with O-16 resonances). The same group structure was used as in the reference-weighted results.

7.2.2.1 Base problem

Fig. 7.23 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. 7.23a gives the k-eigenvalues for each of the calculations. Fig. 7.23b gives the error in k-eigenvalue, power shape, and absorption rates. PG-FEMG has a lower error for the all QOI, beating MG by factors of approximately 3.5, 11, 13, and 15 for power shape, k-eigenvalue, absorption in the fuel, and absorption in the moderator, respectively.

For MG, the increases in error commensurate with changing from a high-fidelity to a low-fidelity weighting spectrum are approximately a factor of 2 for power shape and factors of 150-300 for the other QOI. For PG-FEMG, these same factors are approximately 0.7 for power shape, 50 for k-eigenvalue and absorption in the fuel,

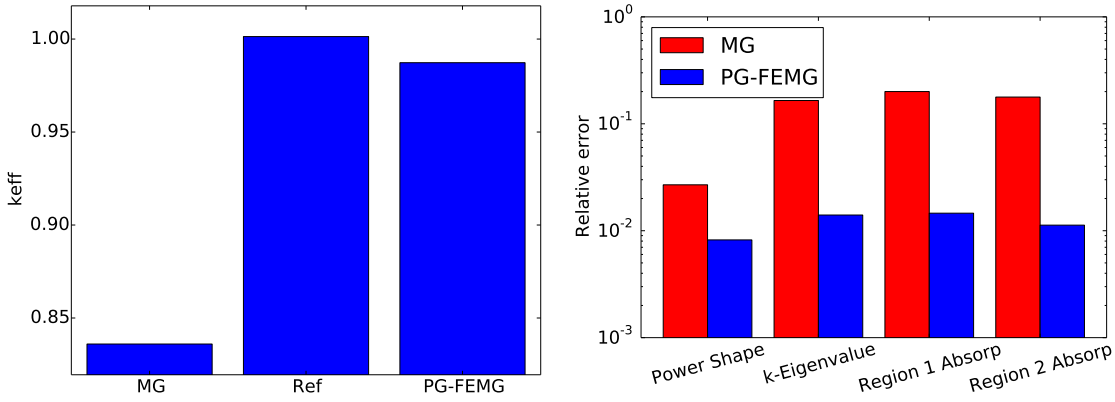

(a) k-Eigenvalues (b) Relative Errors

Figure 7.23: Comparison of eigenvalues and relative errors for problem 2 with low-fidelity weight. Values are given for the reference, MG and PG-FEMG methods.

and 600 for absorption in the moderator. This last value is misleading, because the error in the absorption in the moderator for PG-FEMG with reference weighting was abnormally low (10^{-4}). Differences in the band structures with the two weightings may explain the decrease in error of the power shape for PG-FEMG when going from high-fidelity weighting spectrum to low-fidelity weighting spectrum.

Fig. 7.24a gives the energy dependence of the reference solution source at several equally-spaced radial positions. It shows the source to be affected by the resonances. As the radius is increased, the source becomes smoother.

Fig. 7.24b gives the power profile shapes in space for the reference, MG, and PG-FEMG solutions. Note that the normalization was on this parameter, so they all have the same integral, but may have different spatial shapes. In both cases, the reference solution shows stronger self-shielding effects near the edge of the fuel. The MG and PG-FEMG shapes track the reference shape well until the outside of the fuel, where the reference shows a sharper slope. The error near the edge of the fuel is approximately 2% for PG-FEMG and 6% for MG.

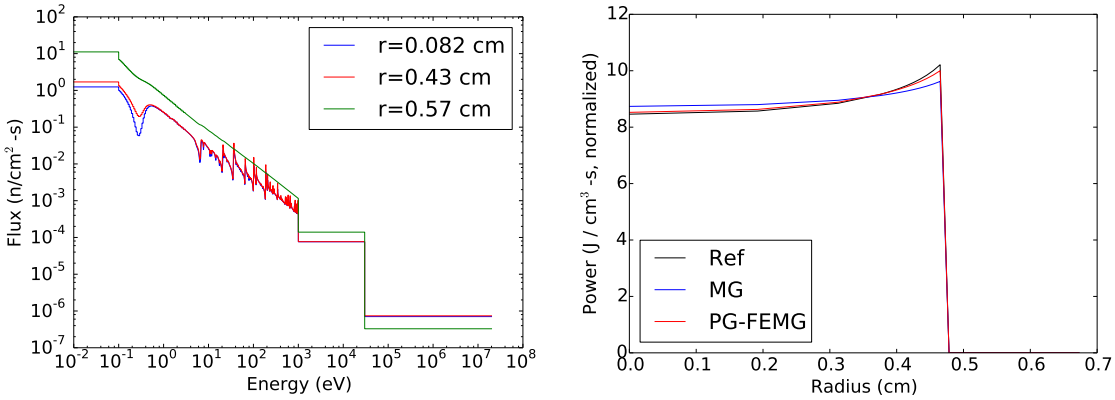

(a) Reference solution neutron source

(b) Power profile

Figure 7.24: Total source and power profile as functions of position for problem 2 with low-fidelity weight. (a) Total source (scattering plus fission) for the reference solution as a function of energy for select radial points. (b) Power profile as a function of position for the reference, MG and PG-FEMG solutions.

Fig. 7.25 (ff.) shows the energy dependence of the reference, MG and PG-FEMG solutions averaged over spatial regions and for the exiting partial current. They show the detailed resonance structure captured by the reference solution and the fidelity of the MG and PG-FEMG reconstructions. Note that the flux dips are space- and direction-dependent. Since a low-fidelity, smooth reference flux was used for both cross section condensation and for flux reexpansion, we expect and find the MG solution to have an inaccurate reconstruction. We find the PG-FEMG solution to have a much more accurate region-averaged reconstruction to the reference solution.

Fig. 7.26 (ff.) shows the energy dependence of the reference, MG and PG-FEMG solutions at selected spatial regions in the fuel. They show the detailed resonance structure captured by the reference solution and the fidelity of the MG and PG-FEMG reconstructions. Since a low-fidelity, smooth reference flux was used for both cross section condensation and for flux reexpansion, we expect and find the MG

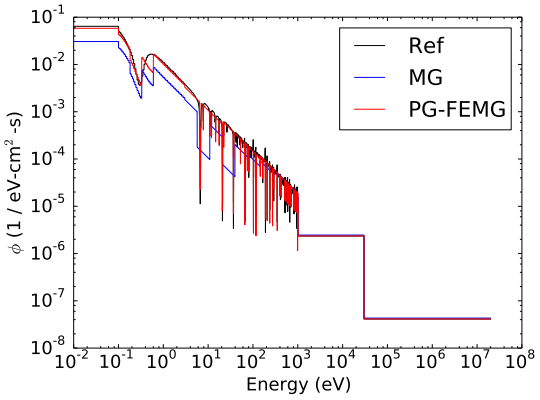

(a) Region 1 average flux

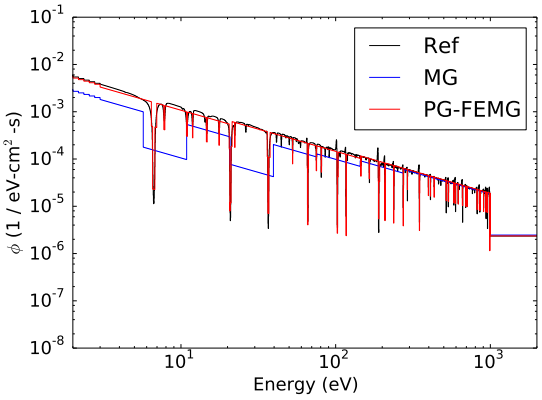

(b) Region 1 average flux (zoom)

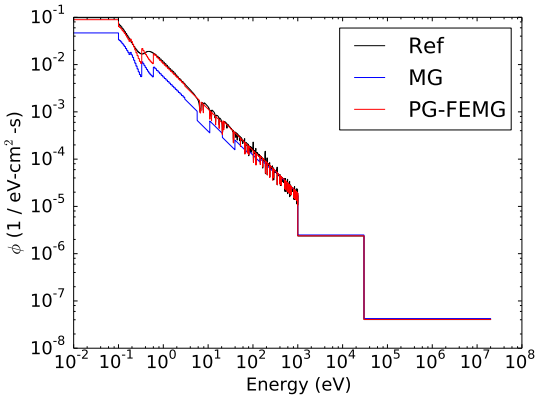

(c) Region 2 average flux

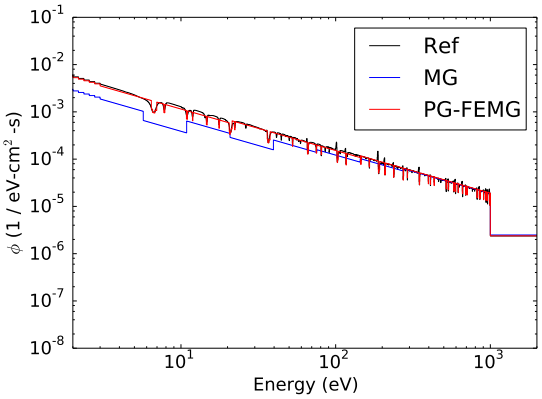

(d) Region 2 average flux (zoom)

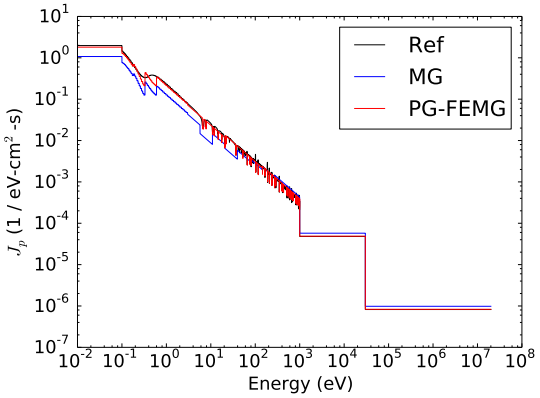

(e) Outgoing partial current

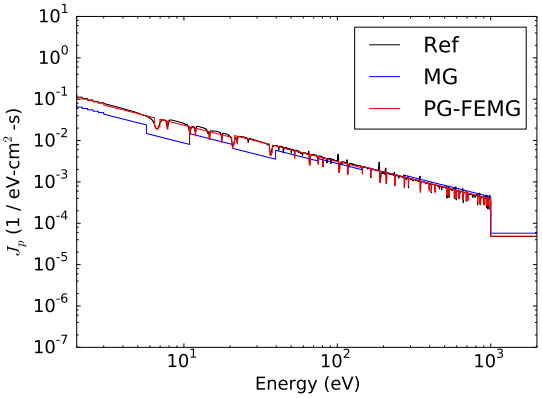

(f) Outgoing partial current (zoom)

Figure 7.25: Energy spectra averaged over spatial regions and at problem boundary for problem 2 with low-fidelity weight.

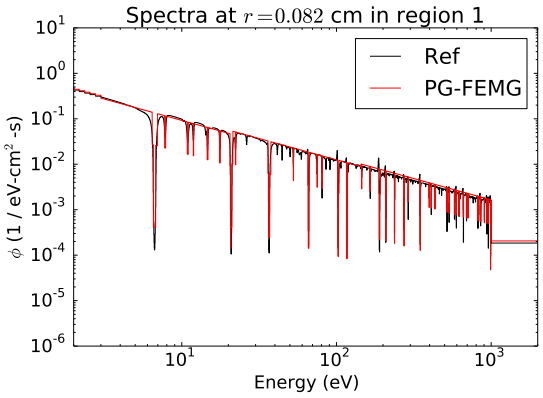

(a) Flux at point 1 for PG-FEMG

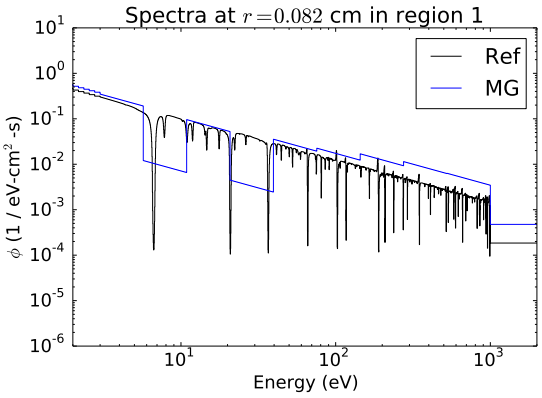

(b) Flux at point 3 for MG

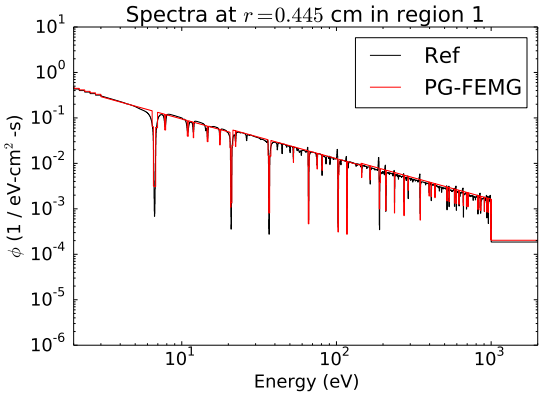

(c) Flux at point 5 for PG-FEMG

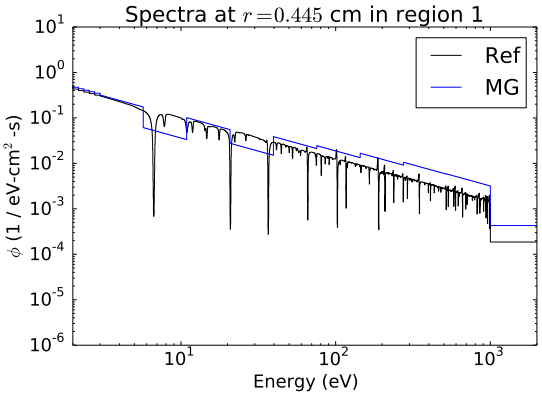

(d) Flux at point 7 for MG

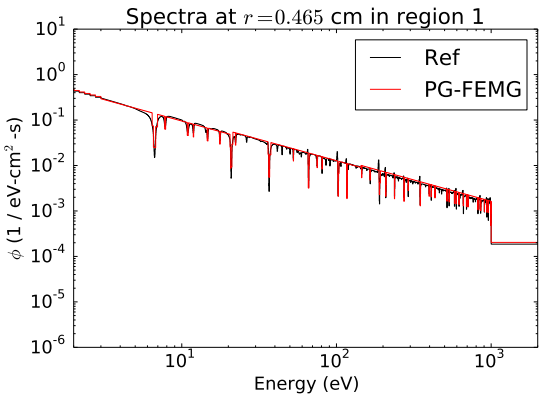

(e) Flux at point 9 for PG-FEMG

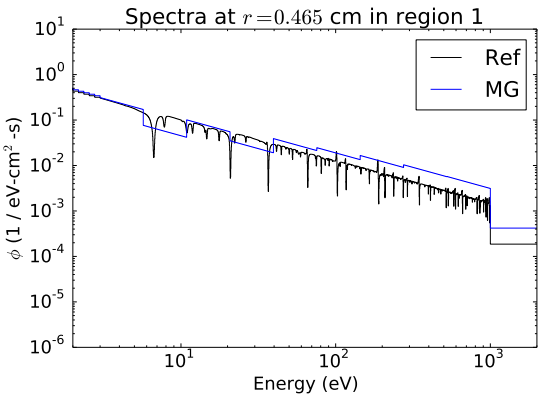

(f) Flux at point 11 for MG

Figure 7.26: Energy spectra at selected spatial points for problem 2 with low-fidelity weight.

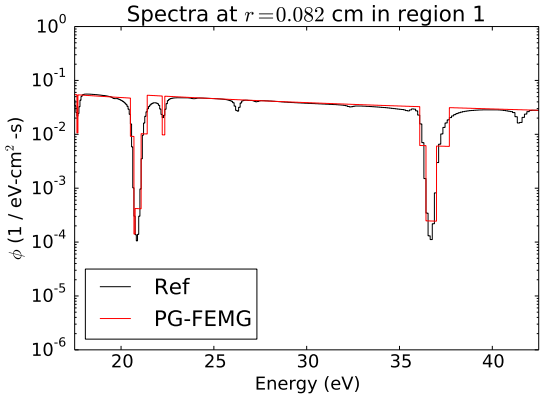

(a) Flux at point 2 for PG-FEMG (zoom)

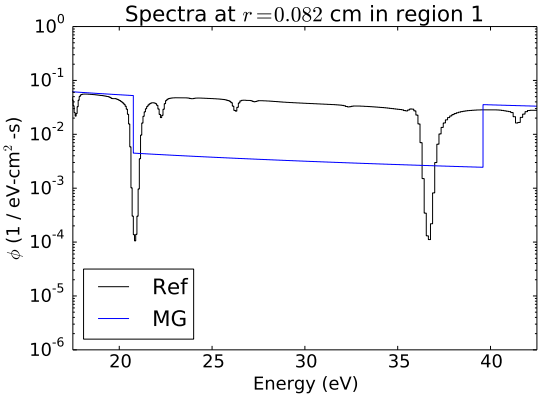

(b) Flux at point 4 for MG (zoom)

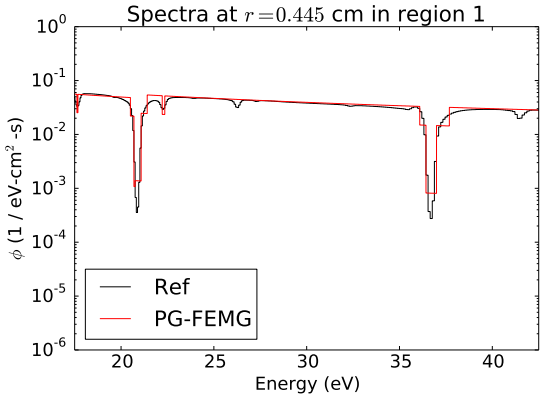

(c) Flux at point 6 for PG-FEMG (zoom)

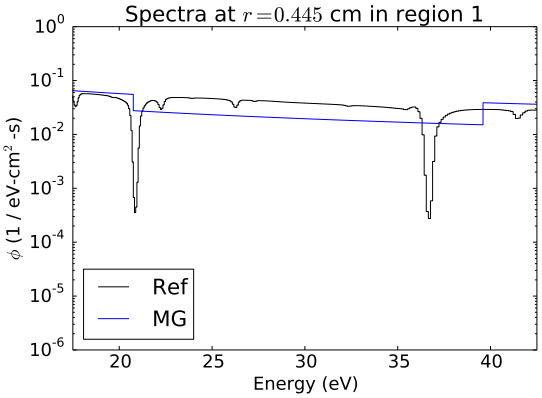

(d) Flux at point 8 for MG (zoom)

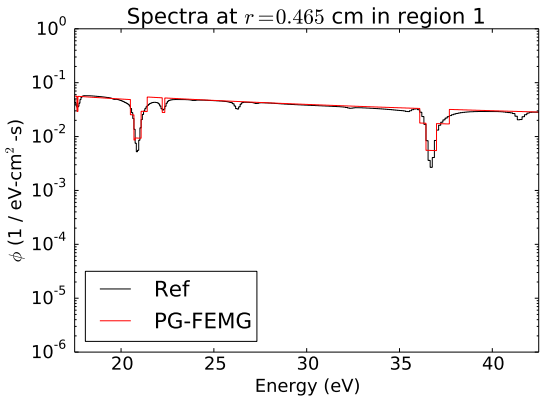

(e) Flux at point 10 for PG-FEMG (zoom)

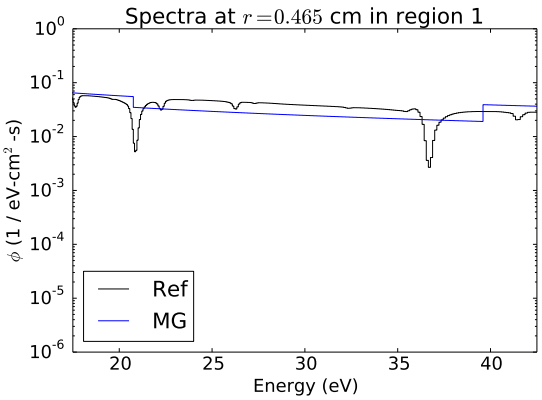

(f) Flux at point 12 for MG (zoom)

Figure 7.27: Energy spectra at selected spatial points for problem 2 with low-fidelity weight (cont.).

solution to have an inaccurate and spatially fixed spectrum reconstruction. The PG-FEMG method has DOF at the resonances. Its spectrum can locally adapt to the correct shape regardless of reference weighting to the extent that the DOF capture the varying portions of the spectrum.

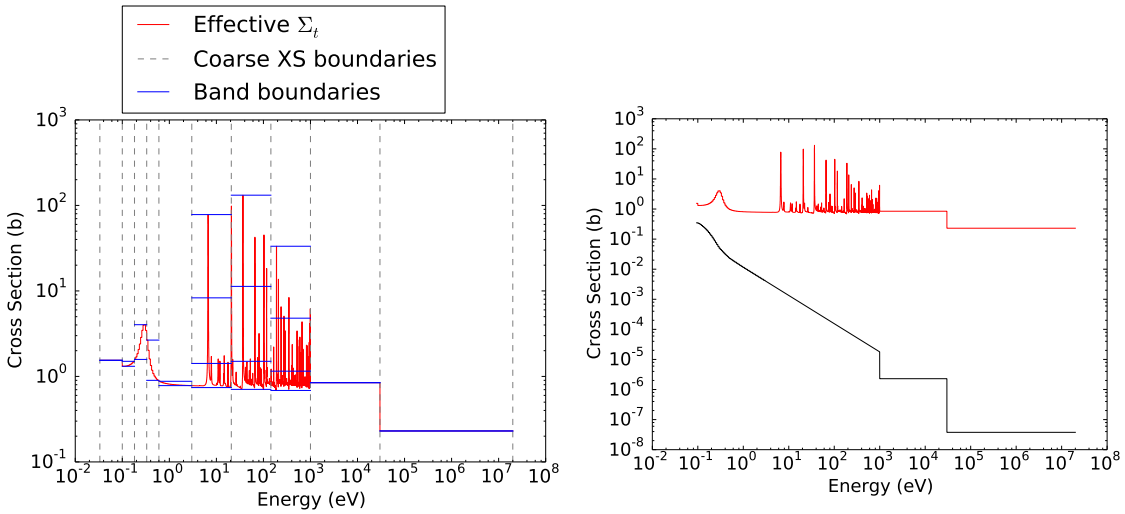

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure 7.28: Effective total cross sections for the PG-FEMG band boundary calculation for problem 2 with low-fidelity weight.

Fig. 7.28a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. 7.28b gives the effective total cross section and the weighting spectrum used. For this calculation, a low-fidelity weighting spectrum was used (note smoothness in b).

Fig. 7.29 ff. give the PG-FEMG (MG, left) and MG (right) cross sections as functions of energy. For the PG-FEMG cases, there may be more than one XS value per coarse group. In an attempt to demonstrate this, lines at each of the band values

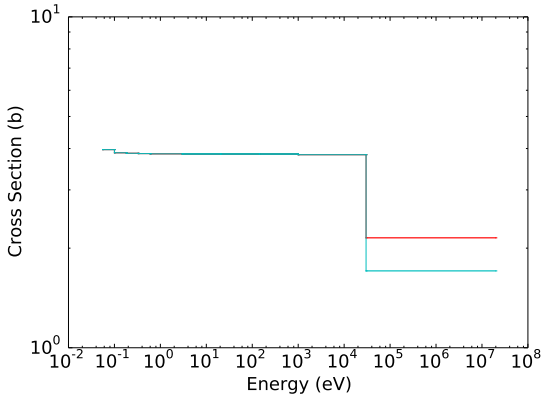

(a) Cross sections for O-16 (PG-FEMG)

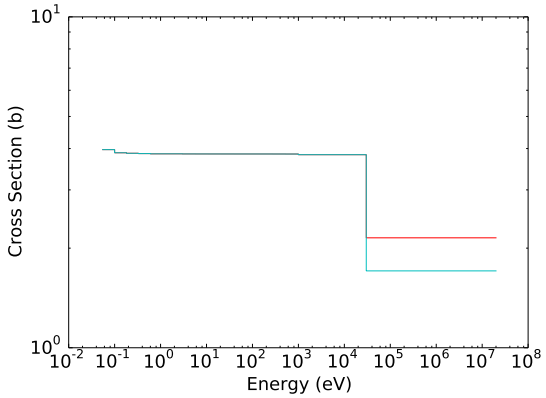

(b) Cross sections for O-16 (MG)

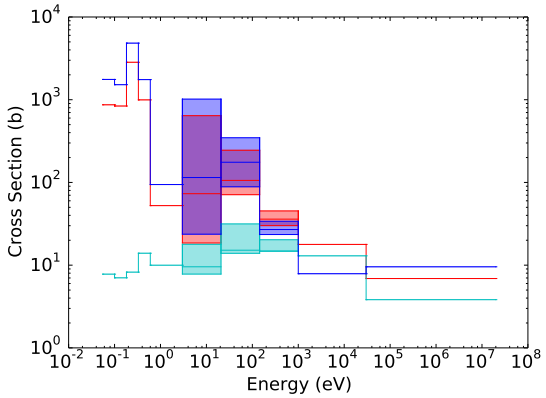

(c) Cross sections for Pu-239 (PG-FEMG)

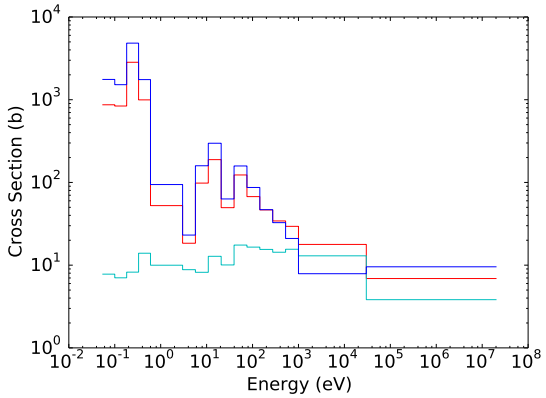

(d) Cross sections for Pu-239 (MG)

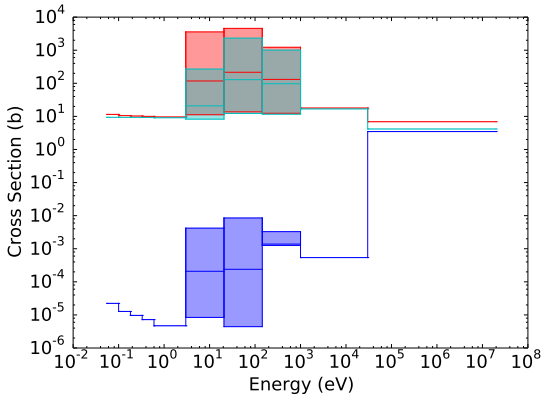

(e) Cross sections for U-238 (PG-FEMG)

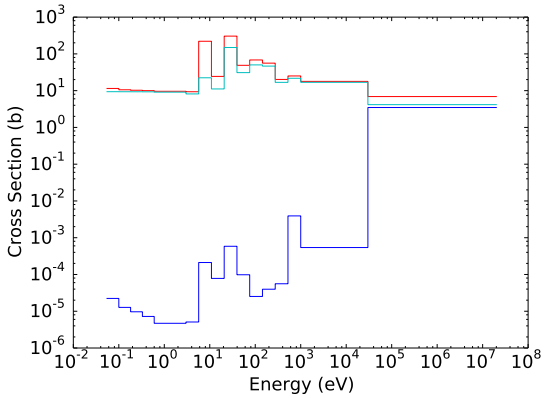

(f) Cross sections for U-238 (MG)

Figure 7.29: Various cross sections as functions of energy for region 1 and problem 2 with low-fidelity weight. Absorption (σ_a), total (σ_t), and fission production ($\nu\sigma_f$) cross sections are given.

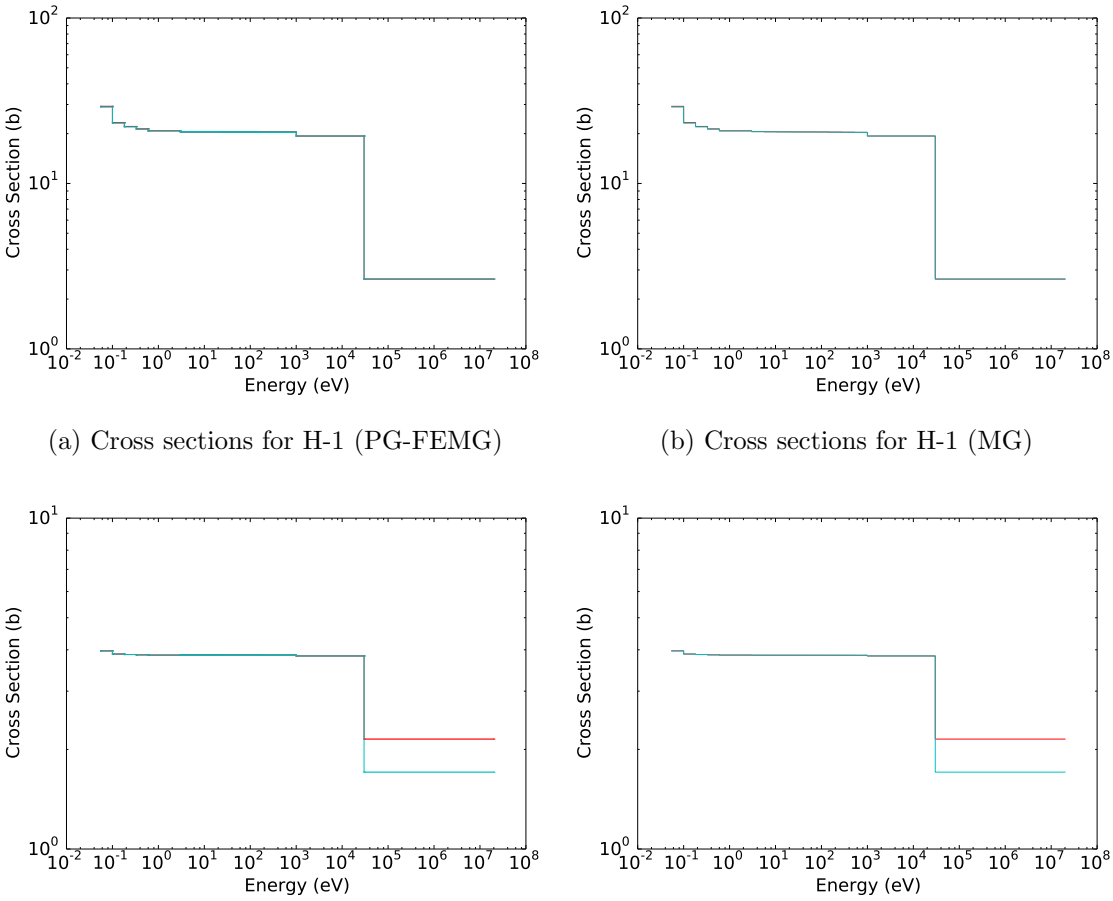

(a) Cross sections for H-1 (PG-FEMG)

(b) Cross sections for H-1 (MG)

(c) Cross sections for O-16 (PG-FEMG)

(d) Cross sections for O-16 (MG)

Figure 7.30: Various cross sections as functions of energy for region 2 and problem 2 with low-fidelity weight. Absorption (σ_a), total (σ_t), and fission production ($\nu\sigma_f$) cross sections are given.

were given with light shading between them to indicate the cross sections range (cf. the picket fence idea of bands).

7.2.2.2 Resonance-only energy resolution study

An energy resolution study was done to determine how various QOIs changed as a function of increasing energy fidelity (DOF) for moderate DOF. The energy group structure was always fully resolved in the non-resonance regions and changed only in the resolved-resonance region. All errors are relative comparisons between the reference solution and the MG or PG-FEMG solutions.

Figure 7.31 gives the results from the energy resolution study. The number of groups and bands per group in the resolved resonance region were varied to determine their effect on QOI error. This resolution study resolved (to the reference level) all non-resonance energies and used varying yet moderate numbers of DOF within the resolved resonance region.

Figure 7.31a-c give work measurements. Figure 7.31a gives efficiency of the k-eigenvalue calculation (1/(work $\times$ error)), Figure 7.31b gives work per DOF, and Figure 7.31c shows total work scaled with the number of groups. Figure 7.31b shows PG-FEMG required approximately the same amount of work per DOF as MG. Both MG and PG-FEMG required substantially less work per DOF than the reference solution (horizontal line at 1.0).

Figure 7.31d-f show how error in the QOI was affected by changing energy resolution. As this resolution study did not use group sizes that resolved individual resonances, errors stayed approximately constant as the number of groups were increased. Errors did decrease as the number of bands per group were increased. For the k-eigenvalue calculation, going from 1 band per group (MG) to 2 bands per group decreased the error by a factor of 8. Increasing from 2 bands to 6 bands decreased

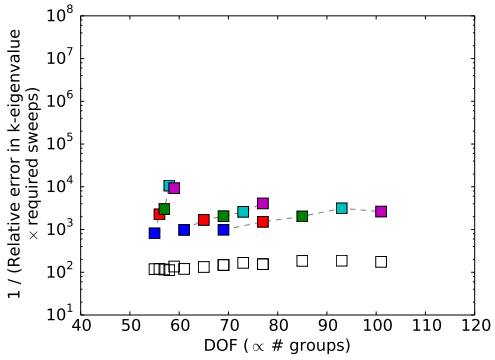

(a) k-Eigenvalue efficiency

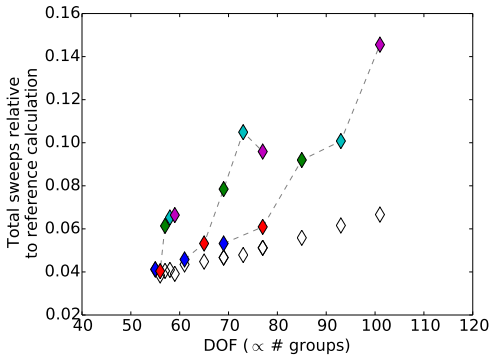

(b) Relative transport sweeps summed over bands / groups (work measurement)

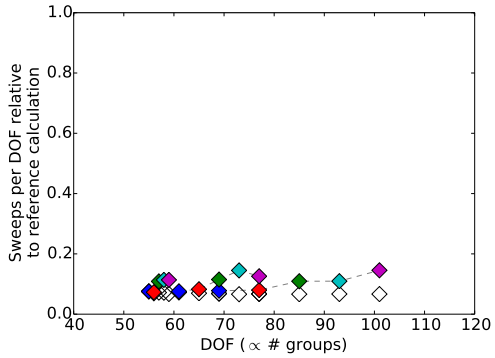

(c) Relative transport sweeps per DOF

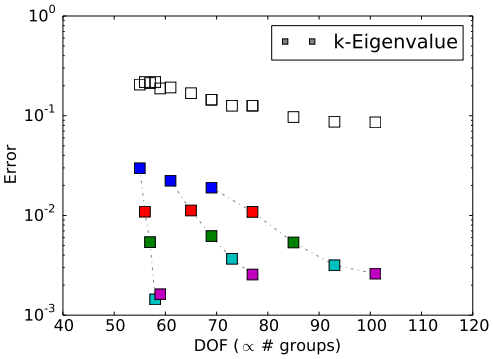

(d) k-Eigenvalue error

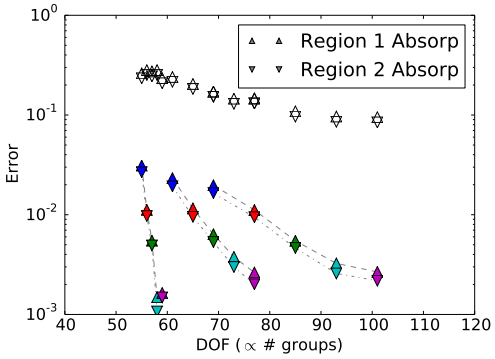

(e) Region absorption rate errors

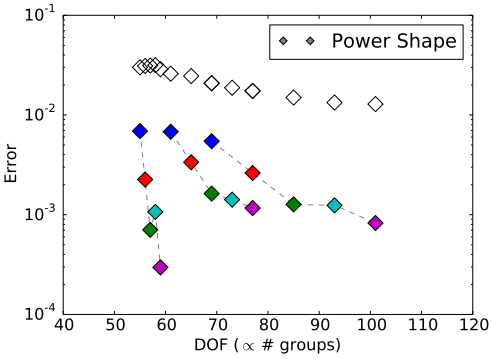

(f) Power shape

Figure 7.31: Error in various QOI relative to the reference solution for several energy group structures and numbers of bands for problem 2 with low-fidelity weight, for the resonance-only energy resolution study. The MG solution corresponds to unfilled markers; the PG-FEMG solutions are colored according to how many bands per group were used in the resonance region, be that two, three, four, five, or six. Dotted lines connect multiband cases with the same coarse group structure.

the same error by another factor of 10. Similar decreases occured for the absorption rates. Increasing from 1 to 2 bands per group decreased the error in the power shape by a factor of almost 4. Increasing from 2 to 6 bands per group decreased the same error by an additional factor of 6-25 depending on group structure. This problem shows error decreasing mostly monotonically as the bands per group were increased. This is expected, as there was only one resonant material region and a poor weighting spectrum was used. As more bands per group were added, the PG-FEMG solutions were better able to adapt their spectra to the true (reference) spectrum.

7.2.2.3 Energy resolution study

An energy resolution study was done to determine how various QOIs changed as a function of increasing energy fidelity (DOF) for large DOF. The energy group structure was resolved in all energy regions for this study. All errors are relative comparisons between the reference solution and the MG or PG-FEMG solutions.

Figure 7.32 gives the results from the energy resolution study. The number of groups and bands per group in the resolved resonance region were varied to determine their effect on QOI error. This resolution study focused on extremely large DOF to investigate asymptotic effects.

Figure 7.32a-c give work measurements. Figure 7.32a gives efficiency of the k-eigenvalue calculation (1/(work $\times$ error)), Figure 7.32b gives work per DOF, and Figure 7.32c shows total work scaled with the number of groups. Figure 7.32b shows PG-FEMG required more work per DOF than MG or the reference solution (horizontal line) in general, but this ratio was approximately constant over all numbers of DOF considered.

Figure 7.32d-f show how error in the QOI was affected by changing energy resolution. They give a sense of the extreme number of groups required to enter the

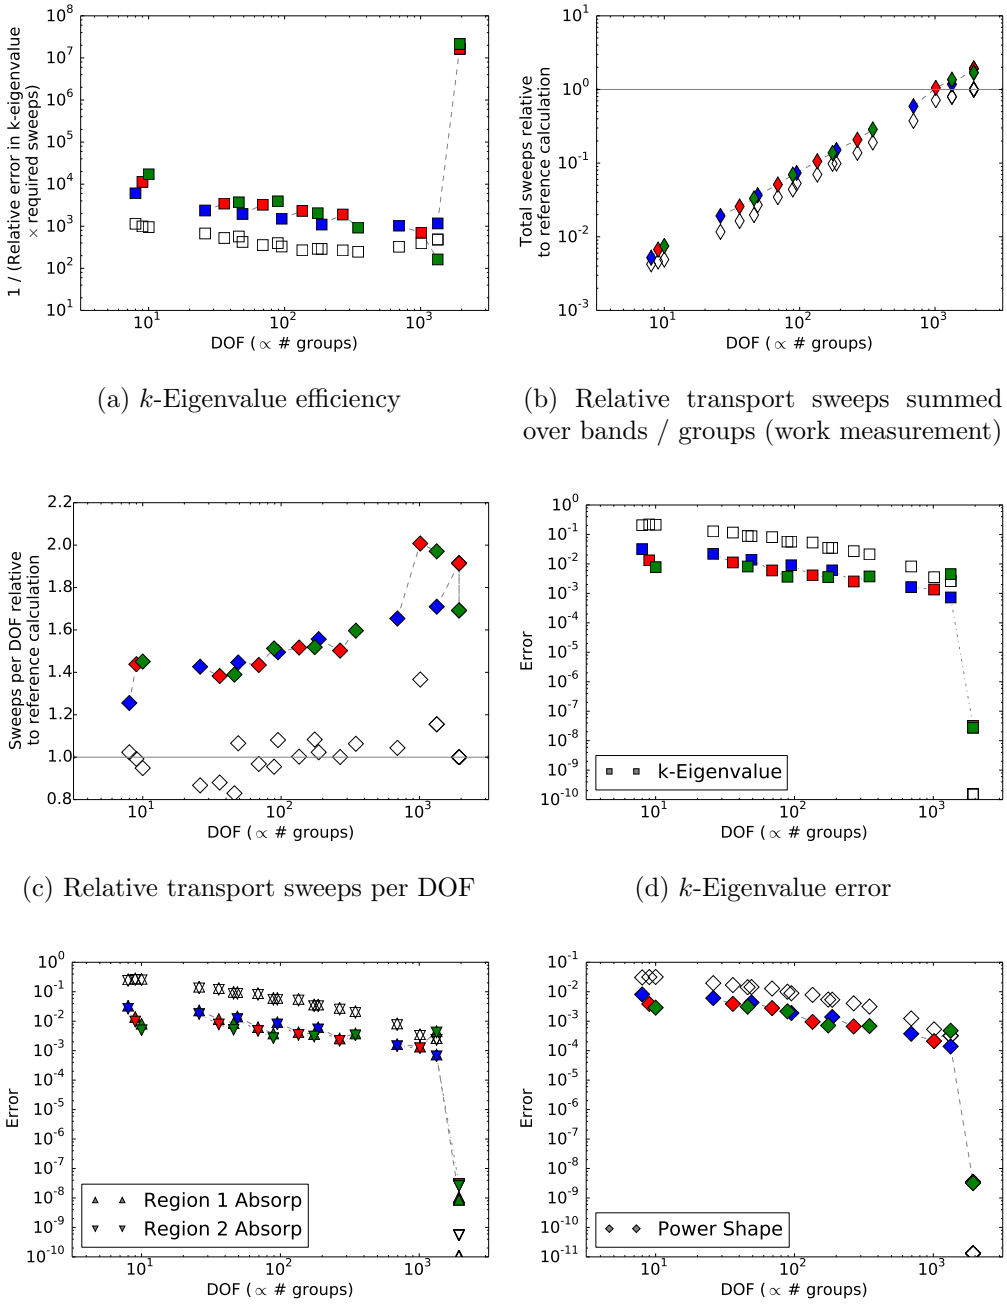

(a) k-Eigenvalue efficiency

(b) Relative transport sweeps summed over bands / groups (work measurement)

(c) Relative transport sweeps per DOF

(d) k-Eigenvalue error

(e) Region absorption rate errors

(f) Power shape

Figure 7.32: Error in various QOI relative to the reference solution for several energy group structures and numbers of bands for problem 2 with low-fidelity weight, for the full energy resolution study. The MG solution corresponds to unfilled markers; the PG-FEMG solutions are colored according to how many bands per group were used in the resonance region, be that two, three, four, five, or six. Dotted lines connect multiband cases with the same coarse group structure.

asymptotic regime. Increasing from 1 band per group (MG) to 2 bands per group decreased the error by a factor of approximately 10 for k-eigenvalue and absorption rates, and about 4 for power shape. Increasing from 2 to 4 bands per group decreased the error by another factor of 4 to 10. The increasing error with increasing bands per group behavior at very large DOF count ($\sim$ 1000 DOF) for PG-FEMG was caused by the fall-back band boundary calculation which produced poor band boundaries.

7.3 Problem 3: Heterogeneous, Multi-T Pin Cells

Problem 3 is a simplified multi-temperature PWR fuel-cell consisting of fuel (U-238, Pu-239, and O-16) and moderator (O-16 and H-1). Four regions are used: three different temperatures in the fuel and one in the moderator. A white boundary was used. Figure 7.33 gives a cartoon of the geometry.

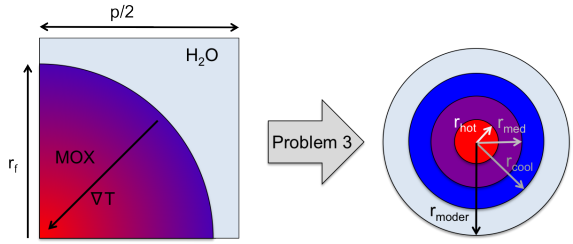

(a) Problem 3

Figure 7.33: Cartoon of the geometry and its simplification to 1D cylindrical geometry for problem 3.

Table 7.3 gives a description of each region for problem 3. Regions are defined as spatially contiguous groups of cells that share a common cross section and are ordered such that larger region numbers correspond to larger radii in the problem. When spatial averaging is done, it is over a region. Unlisted temperatures are 400 K. Tables 6.1 and 6.2 give further geometry and material information, respectively.

Table 7.3: List of spatial regions for problem 3.

Region Number	Description
1	MOX (1000 K)
2	MOX (800 K)
3	MOX (700 K)
4	Light water (550 K)

7.3.1 Reference-calculation weighting

In this subsection, the region-averaged reference solution was used to flux-weight the (ultra-fine-group) cross sections during the condensation processes to produce the MG and PG-FEMG cross sections The same DOF were used for both MG and PG-FEMG runs.

7.3.1.1 Base problem

Fig. 7.34 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. 7.34a gives the k-eigenvalues for each of the calculations. Fig. 7.34b gives the error in k-eigenvalue, power shape, and absorption rates. Both PG-FEMG and MG do approximately equally well with the power shape error. PG-FEMG has a lower error for most other QOI, beating MG by factors of approximately 3, 1.5, 2, and 6 for k-eigenvalue, absorption in the hot fuel, absorption in the medium-temperature fuel, and absorption in the moderator, respectively. MG has lower absorption error in the cool fuel by a factor of approximately 1.6. Since each fuel temperature was treated as a separate region, and cross sections were condensed by region, MG was expected to perform well for this case, as its weighting factors were more local than in other problems where the entire fuel was one region.

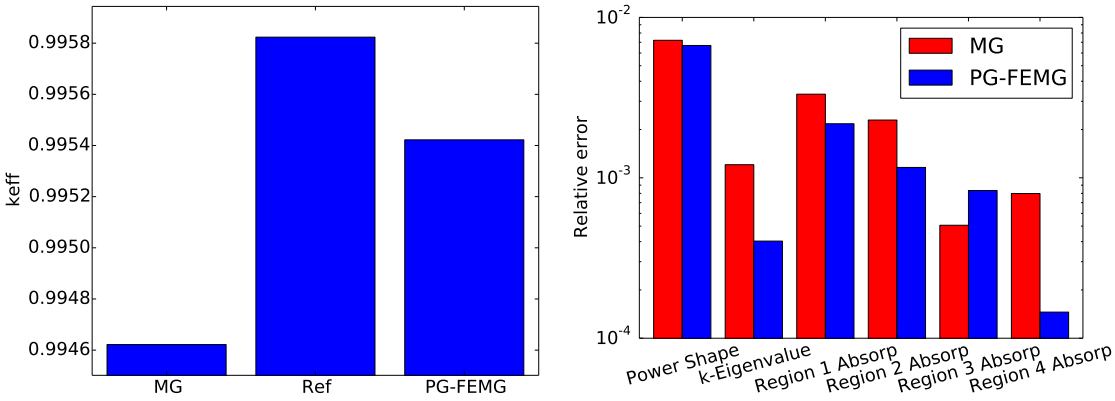

(a) k-Eigenvalues (b) Relative Errors

Figure 7.34: Comparison of eigenvalues and relative errors for problem 3 with high-fidelity weight. Values are given for the reference, MG and PG-FEMG methods.

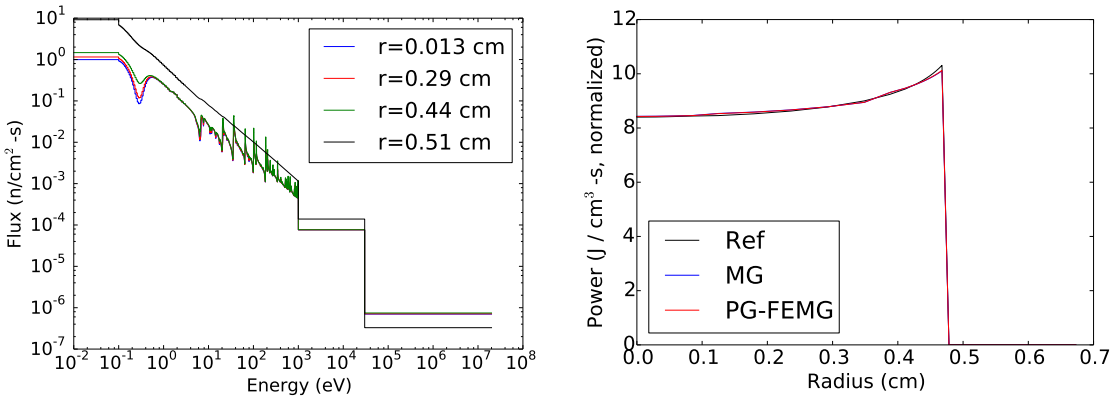

(a) Reference solution neutron source (b) Power profile

Figure 7.35: Total source and power profile as functions of position for problem 3 with high-fidelity weight. (a) Total source (scattering plus fission) for the reference solution as a function of energy for select radial points. (b) Power profile as a function of position for the reference, MG and PG-FEMG solutions.

Fig. 7.35a gives the energy dependence of the reference solution source at several equally-spaced radial positions. It shows the source to be affected by the resonances. As the radius is increased, the source becomes smoother.

Fig. 7.35b gives the power profile shapes in space for the reference, MG, and PG-FEMG solutions. Note that the normalization was on this parameter, so they all have the same integral, but may have different spatial shapes. In both cases, the reference solution shows stronger self-shielding effects near the edge of the fuel. The MG and PG-FEMG shapes track the reference shape well until the outside of the fuel, where the reference shows a sharper slope. The error near the edge of the fuel is approximately 2% for both cases.

Fig. 7.36 (ff.) shows the energy dependence of the reference, MG and PG-FEMG solutions averaged over spatial regions and for the exiting partial current. They show the detailed resonance structure captured by the reference solution and the fidelity of the MG and PG-FEMG reconstructions. Note that the flux dips are space- and direction-dependent. Since a reference solution was used for both cross section condensation and for flux reexpansion, we expect and find both the MG and PG-FEMG solutions to have region-averaged reconstructions accurate to the reference solution.

Fig. 7.38 gives the unprocessed group / band fluxes as functions of position for the reference, MG, and PG-FEMG solutions. A set of representative energies were used, the most interesting of which were the in- and out-of- resonace cases. These show the effect of adding DOF to the resonances with PG-FEMG, where fidelity is preserved, or MG, where features are lost. Note that boundary layers are preserved with PG-FEMG.

Fig. 7.39a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note

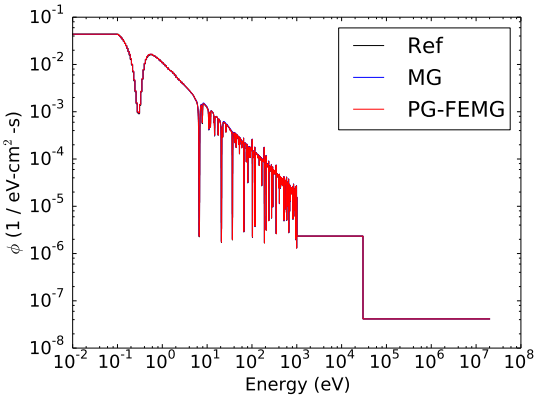

(a) Region 1 average flux

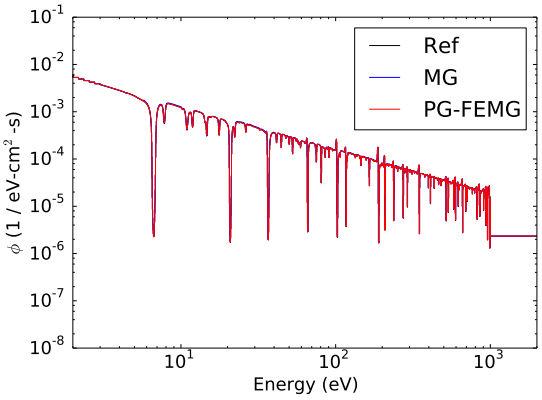

(b) Region 1 average flux (zoom)

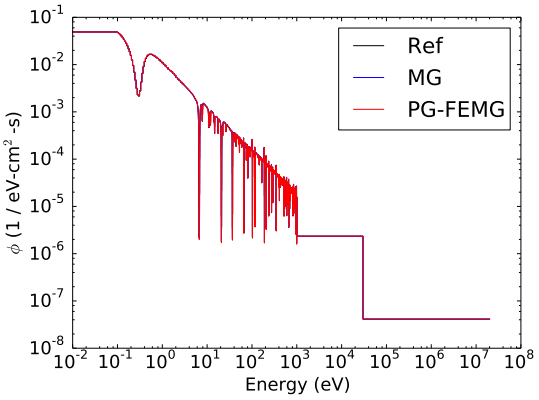

(c) Region 2 average flux

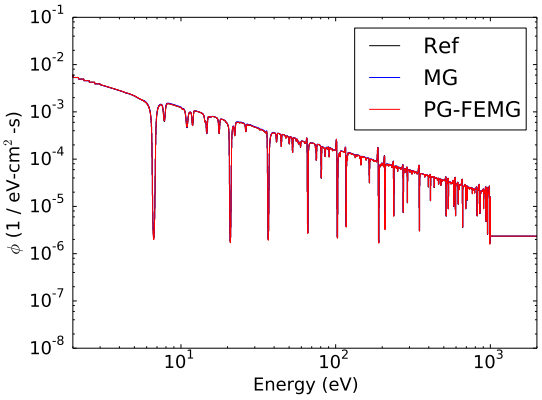

(d) Region 2 average flux (zoom)

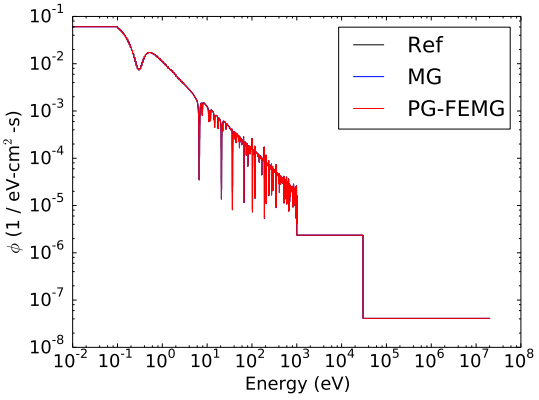

(e) Region 3 average flux

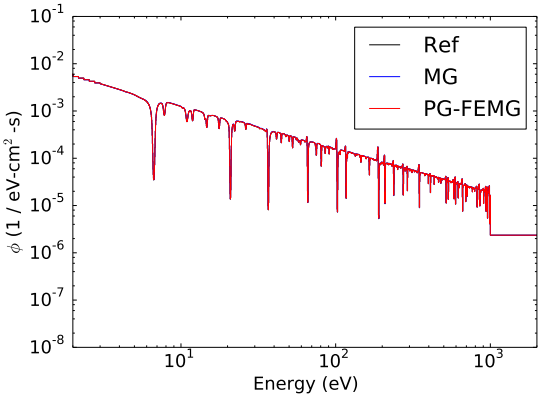

(f) Region 3 average flux (zoom)

Figure 7.36: Energy spectra averaged over spatial regions and at problem boundary for problem 3 with high-fidelity weight.

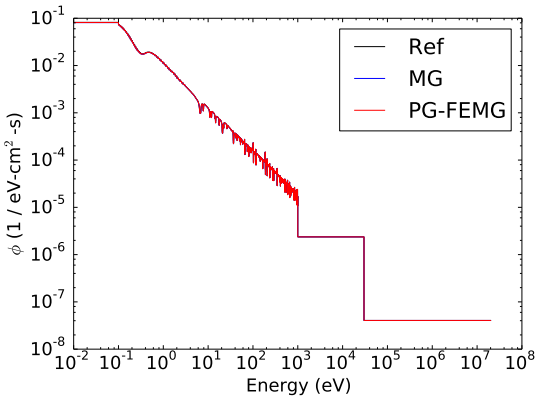

(a) Region 4 average flux

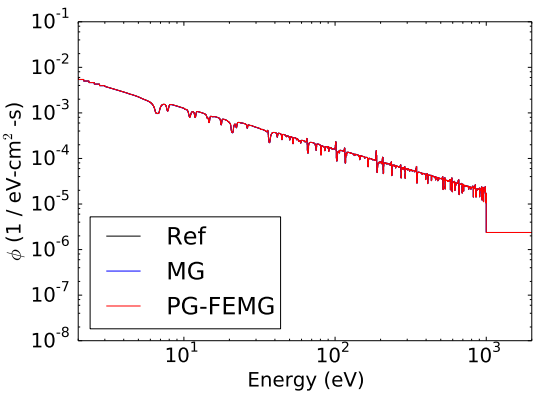

(b) Region 4 average flux (zoom)

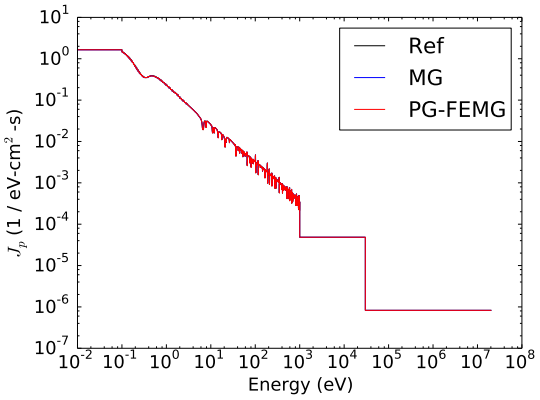

(c) Outgoing partial current

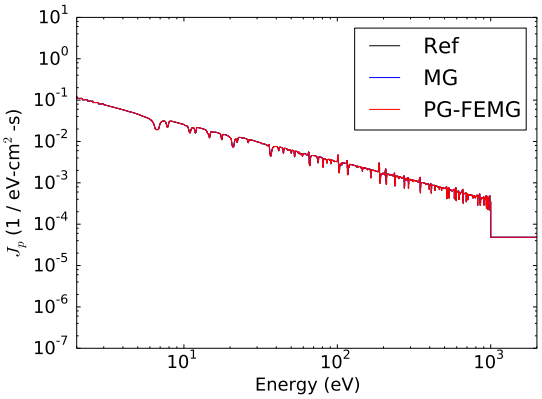

(d) Outgoing partial current (zoom)

Figure 7.37: Energy spectra averaged over spatial regions and at problem boundary for problem 3 with high-fidelity weight (cont.).

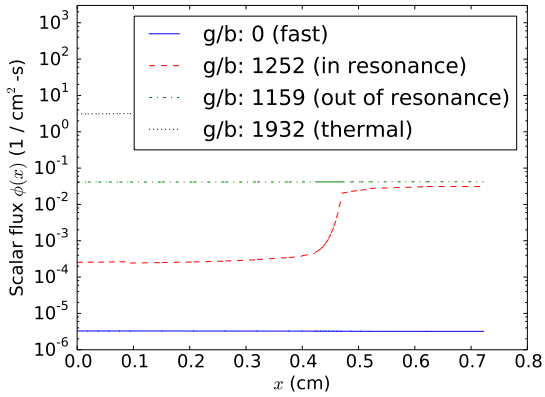

(a) Reference

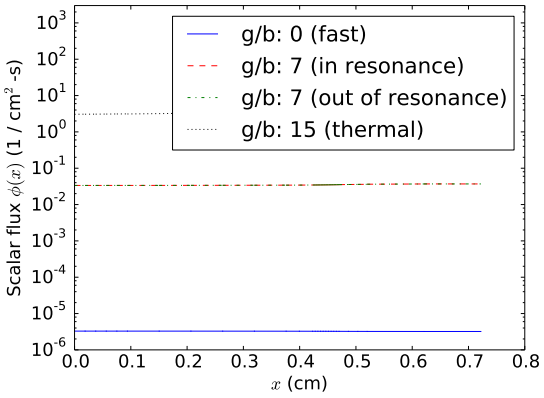

(b) MG

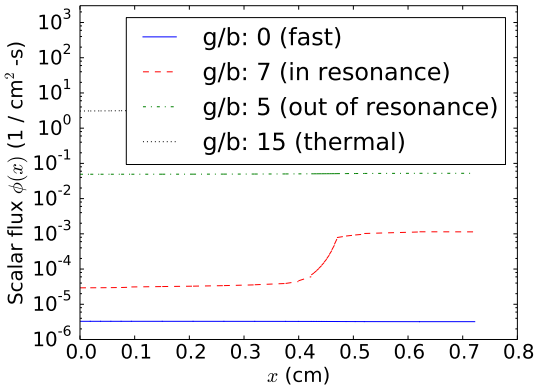

(c) PG-FEMG

Figure 7.38: Spatial dependence for several groups of interest for problem 3 with high-fidelity weight. The group corresponds to the same energy values in all cases. Since no expansion has been done on the multigroup or multiband fluxes, they should only be compared qualitatively to the reference fluxes.

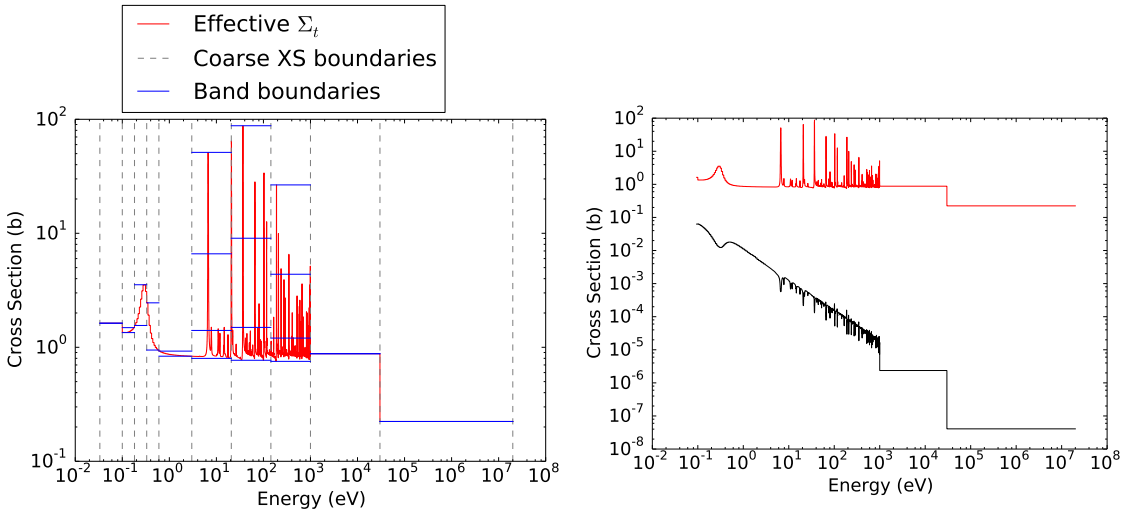

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure 7.39: Effective total cross sections for the PG-FEMG band boundary calculation for problem 3 with high-fidelity weight.

that multiple bands per coarse group were only used in the resonance region. Fig. 7.39b gives the effective total cross section and the weighting spectrum used. For this calculation, a high-fidelity weighting spectrum was used (note resonances in b).

7.3.1.2 Resonance-only energy resolution study

An energy resolution study was done to determine how various QOIs changed as a function of increasing energy fidelity (DOF) for moderate DOF. The energy group structure was always fully resolved in the non-resonance regions and changed only in the resolved-resonance region. All errors are relative comparisons between the reference solution and the MG or PG-FEMG solutions.

Figure 7.40 gives the results from the energy resolution study. The number of groups and bands per group in the resolved resonance region were varied to determine their effect on QOI error. This resolution study resolved (to the reference level) all non-resonance energies and used varying yet moderate numbers of DOF within the

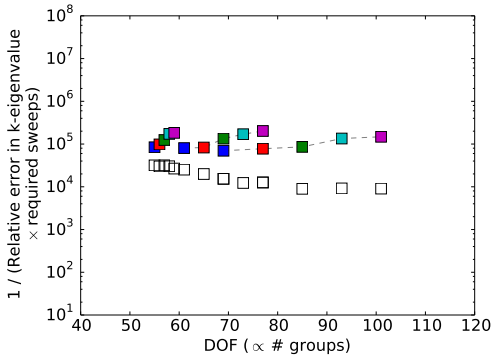

(a) k-Eigenvalue efficiency

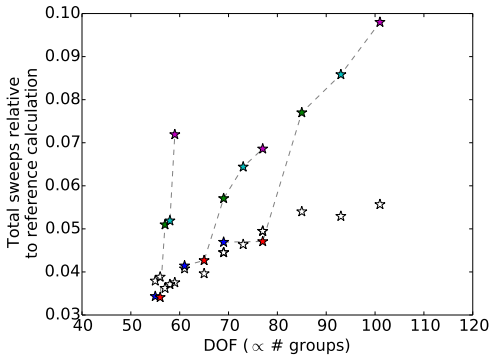

(b) Relative transport sweeps summed over bands / groups (work measurement)

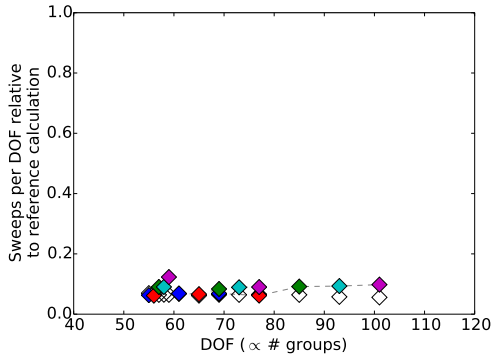

(c) Relative transport sweeps per DOF

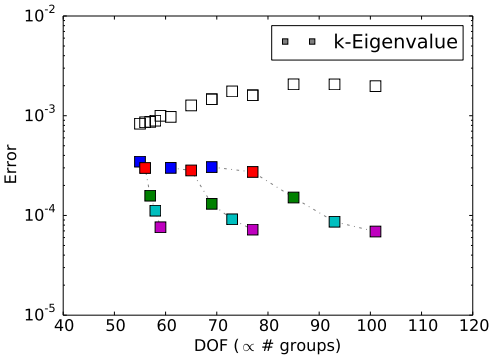

(d) k-Eigenvalue error

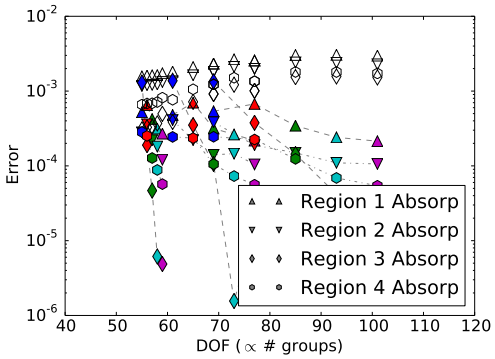

(e) Region absorption rate errors

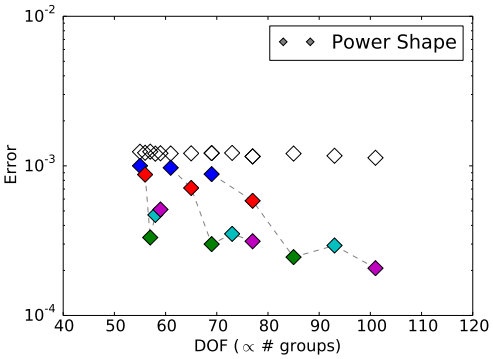

(f) Power shape

Figure 7.40: Error in various QOI relative to the reference solution for several energy group structures and numbers of bands for problem 3 with high-fidelity weight, for the resonance-only energy resolution study. The MG solution corresponds to unfilled markers; the PG-FEMG solutions are colored according to how many bands per group were used in the resonance region, be that two, three, four, five, or six. Dotted lines connect multiband cases with the same coarse group structure.

resolved resonance region.

Figure 7.40a-c give work measurements. Figure 7.40a gives efficiency of the k-eigenvalue calculation (1/(work $\times$ error)), Figure 7.40b gives work per DOF, and Figure 7.40c shows total work scaled with the number of groups. Figure 7.40b shows PG-FEMG required approximately the same amount of work per DOF as MG. Both MG and PG-FEMG required substantially less work per DOF than the reference solution (horizontal line at 1.0).

Figure 7.40d-f show how error in the QOI was affected by changing energy resolution. As this resolution study did not use group sizes that resolved individual resonances, errors stayed approximately constant as the number of groups were increased (for PG-FEMG) or increased as the number of groups increased (for MG). Errors did decrease as the number of bands per group were increased. For the k-eigenvalue calculation, going from 1 band per group (MG) to 2 bands per group decreased the error by a factor of approximately 2-5, depending on group structure. Increasing from 2 bands to 6 bands decreased the same error by another factor of approximately 5. Increasing from 1 to 2 bands per group had a negligible effect on power shape error. Increasing from 2 to 6 bands per group decreased the same error by a factor of approximately 2-5, depending on group structure. Increasing bands per group generally monotonically decreased QOI error except for 5 bands per group and the power shape error, which always increased.

7.3.1.3 Energy resolution study

An energy resolution study was done to determine how various QOIs changed as a function of increasing energy fidelity (DOF) for large DOF. The energy group structure was resolved in all energy regions for this study. All errors are relative comparisons between the reference solution and the MG or PG-FEMG solutions.

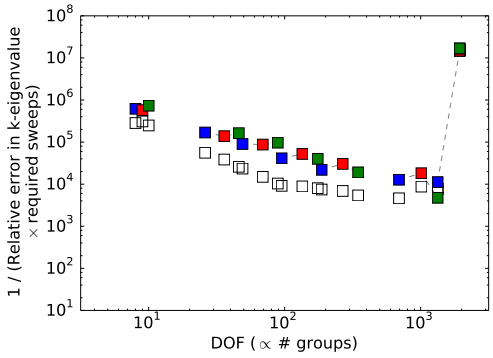

(a) k-Eigenvalue efficiency

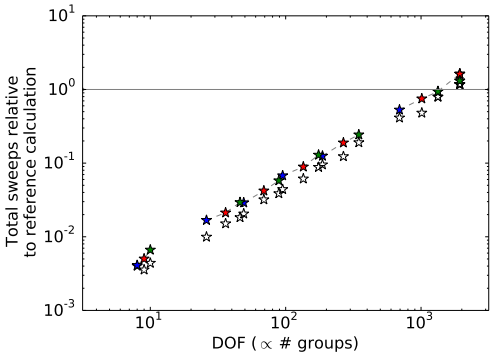

(b) Relative transport sweeps summed over bands / groups (work measurement)

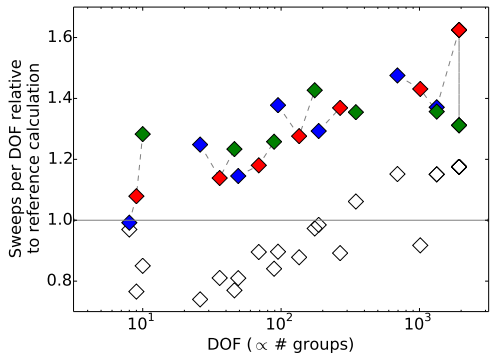

(c) Relative transport sweeps per DOF

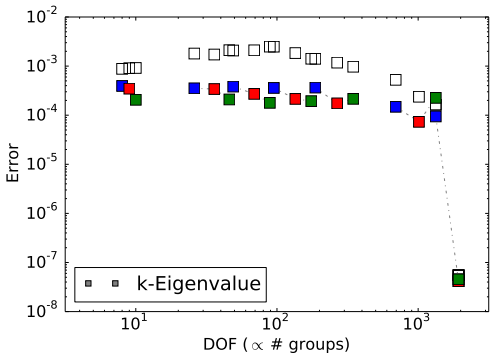

(d) k-Eigenvalue error

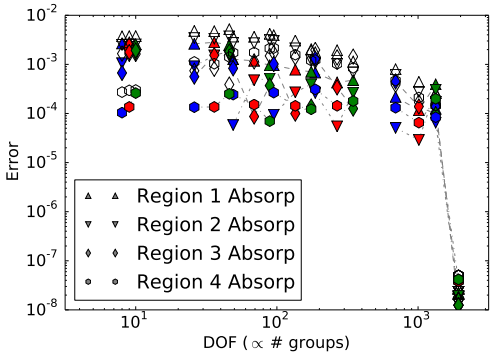

(e) Region absorption rate errors

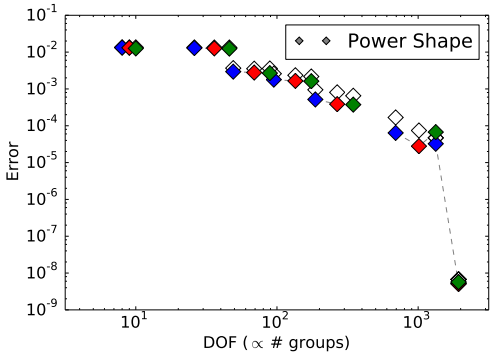

(f) Power shape

Figure 7.41: Error in various QOI relative to the reference solution for several energy group structures and numbers of bands for problem 3 with high-fidelity weight, for the full energy resolution study. The MG solution corresponds to unfilled markers; the PG-FEMG solutions are colored according to how many bands per group were used in the resonance region, be that two, three, four, five, or six. Dotted lines connect multiband cases with the same coarse group structure.

Figure 7.41 gives the results from the energy resolution study. The number of groups and bands per group in the resolved resonance region were varied to determine their effect on QOI error. This resolution study focused on extremely large DOF to investigate asymptotic effects.

Figure 7.41a-c give work measurements. Figure 7.41a gives efficiency of the k-eigenvalue calculation (1/(work $\times$ error)), Figure 7.41b gives work per DOF, and Figure 7.41c shows total work scaled with the number of groups. Figure 7.41b shows PG-FEMG required more work per DOF than MG or the reference solution (horizontal line) in general, but this ratio was approximately constant over all numbers of DOF considered.

Figure 7.41d-f show how error in the QOI was affected by changing energy resolution. They give a sense of the extreme number of groups required to enter the asymptotic regime. Increasing from 1 band per group (MG) to 2 bands per group decreased the error by a factor of approximately 3-7 for k-eigenvalue. Increasing from 2 to 4 bands per group decreased the error by another factor of 1.5-2. Only for very large DOF count (200+) did increasing the bands per group lower power shape error. Absorption rate errors in general decreased with increasing bands per group. The increasing error with increasing bands per group behavior at very large DOF count ($\sim$ 1000 DOF) for PG-FEMG was caused by the fall-back band boundary calculation which produced poor band boundaries.

7.3.2 Generic weighting

In this subsection, a generic, low-fidelity solution was used to weight the cross sections during the condensation process, namely the iwt=5 option in NJOY (mid-life PWR spectrum with O-16 resonances). The same group structure was used as in the reference-weighted results.

7.3.2.1 Base problem

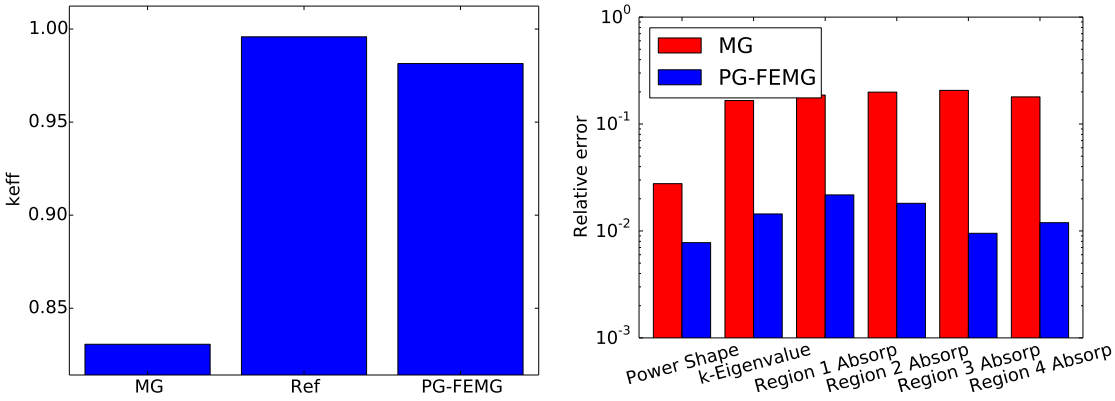

(a) k-Eigenvalues (b) Relative Errors

Figure 7.42: Comparison of eigenvalues and relative errors for problem 3 with low-fidelity weight. Values are given for the reference, MG and PG-FEMG methods.

Fig. 7.42 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. 7.42a gives the k-eigenvalues for each of the calculations. Fig. 7.42b gives the error in k-eigenvalue, power shape, and absorption rates. PG-FEMG has a lower error for the all QOI, beating MG by factors of approximately 3.7, 14, 9, 12, 22, and 15 for power shape, k-eigenvalue, absorption in the hot fuel, absorption in the medium-temperature fuel, absorption in the cool fuel, and absorption in the moderator, respectively.

For MG, the increases in error commensurate with changing from a high-fidelity to a low-fidelity weighting spectrum are approximately a factor of 4 for power shape, a factor of 400 for absorption in the cool fuel, and factors of 50-250 for the other QOI. For PG-FEMG, these same factors are approximately 1.1 for power shape, 33 for k-eigenvalue, 90 for absorption in the moderator, and 10-14 for absorptions in

the fuel. The large increase in error for MG for absorption in the cool fuel between reference weighting and generic weighting shows MG relies on the proper spectrum to determine QOI accurately.

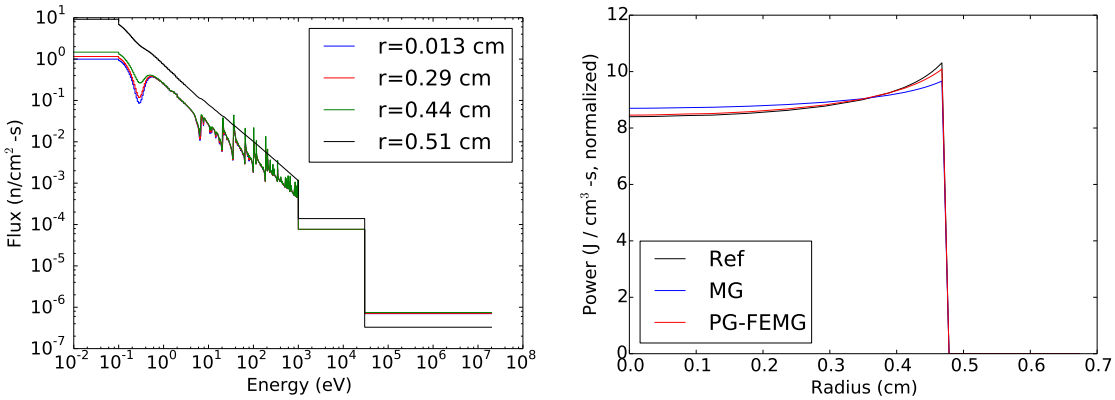

(a) Reference solution neutron source

(b) Power profile

Figure 7.43: Total source and power profile as functions of position for problem 3 with low-fidelity weight. (a) Total source (scattering plus fission) for the reference solution as a function of energy for select radial points. (b) Power profile as a function of position for the reference, MG and PG-FEMG solutions.

Fig. 7.43a gives the energy dependence of the reference solution source at several equally-spaced radial positions. It shows the source to be affected by the resonances. As the radius is increased, the source becomes smoother.

Fig. 7.43b gives the power profile shapes in space for the reference, MG, and PG-FEMG solutions. Note that the normalization was on this parameter, so they all have the same integral, but may have different spatial shapes. In both cases, the reference solution shows stronger self-shielding effects near the edge of the fuel. The MG and PG-FEMG shapes track the reference shape well until the outside of the fuel, where the reference shows a sharper slope. The error near the edge of the fuel

is approximately 2% for PG-FEMG and 6% for MG.

Fig. 7.44 (ff.) shows the energy dependence of the reference, MG and PG-FEMG solutions averaged over spatial regions and for the exiting partial current. They show the detailed resonance structure captured by the reference solution and the fidelity of the MG and PG-FEMG reconstructions. Note that the flux dips are space- and direction-dependent. Since a low-fidelity, smooth reference flux was used for both cross section condensation and for flux reexpansion, we expect and find the MG solution to have an inaccurate reconstruction. We find the PG-FEMG solution to have a much more accurate region-averaged reconstruction to the reference solution.

Fig. 7.46a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. 7.46b gives the effective total cross section and the weighting spectrum used. For this calculation, a low-fidelity weighting spectrum was used (note smoothness in b).

7.3.2.2 Resonance-only energy resolution study

An energy resolution study was done to determine how various QOIs changed as a function of increasing energy fidelity (DOF) for moderate DOF. The energy group structure was always fully resolved in the non-resonance regions and changed only in the resolved-resonance region. All errors are relative comparisons between the reference solution and the MG or PG-FEMG solutions.

Figure 7.47 gives the results from the energy resolution study. The number of groups and bands per group in the resolved resonance region were varied to determine their effect on QOI error. This resolution study resolved (to the reference level) all non-resonance energies and used varying yet moderate numbers of DOF within the resolved resonance region.

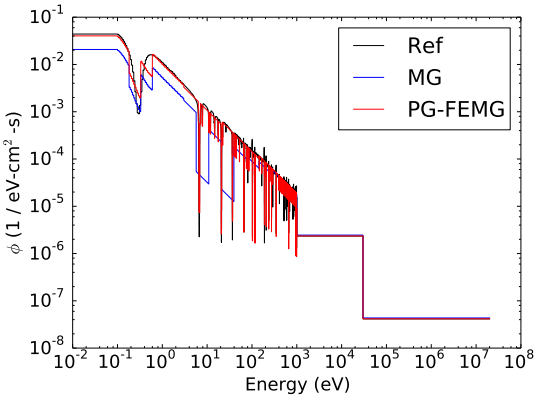

(a) Region 1 average flux

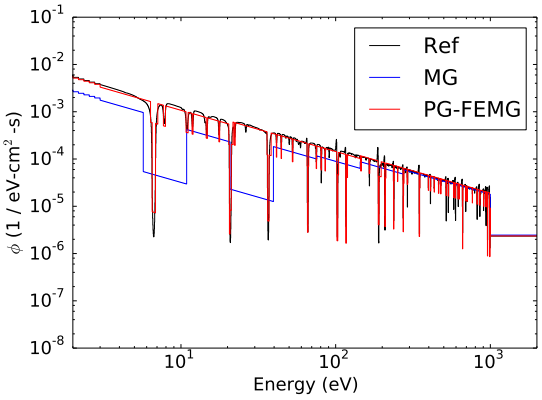

(b) Region 1 average flux (zoom)

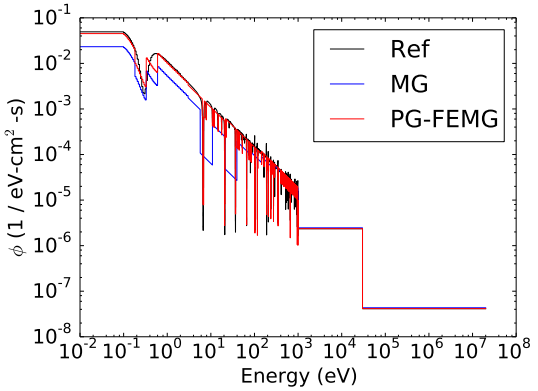

(c) Region 2 average flux

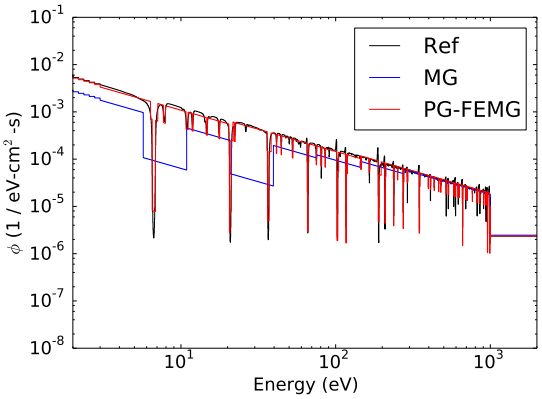

(d) Region 2 average flux (zoom)

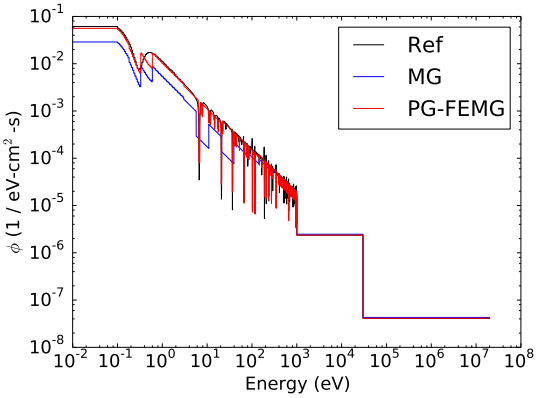

(e) Region 3 average flux

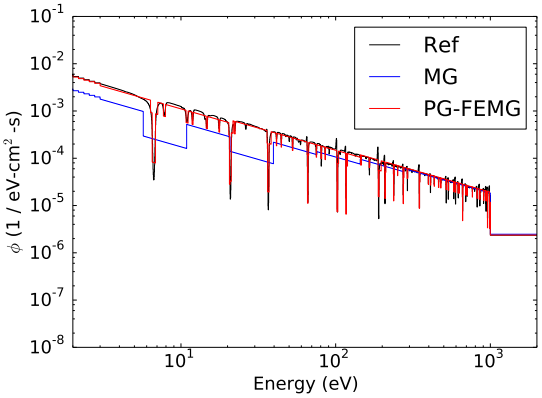

(f) Region 3 average flux (zoom)

Figure 7.44: Energy spectra averaged over spatial regions and at problem boundary for problem 3 with low-fidelity weight.

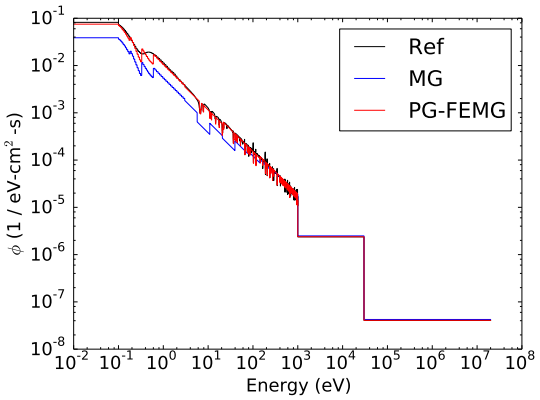

(a) Region 4 average flux

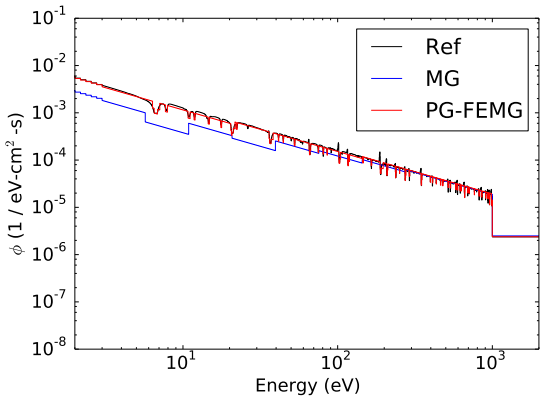

(b) Region 4 average flux (zoom)

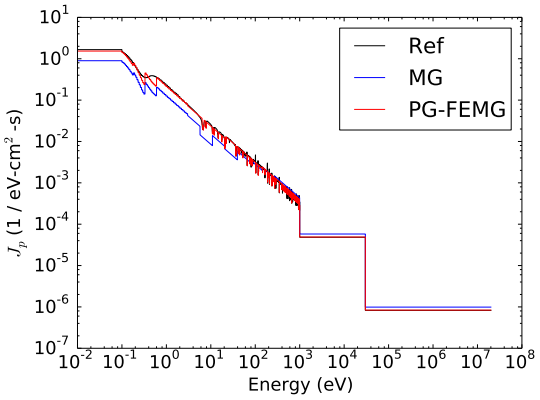

(c) Outgoing partial current

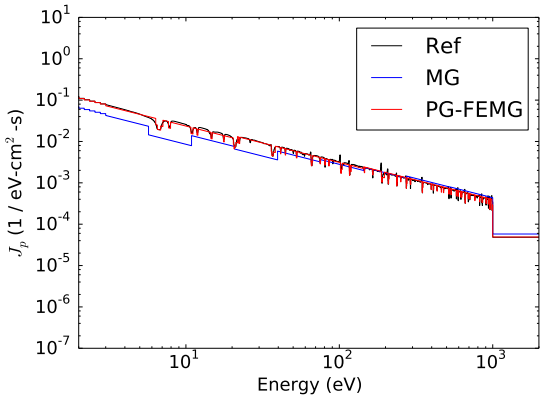

(d) Outgoing partial current (zoom)

Figure 7.45: Energy spectra averaged over spatial regions and at problem boundary for problem 3 with low-fidelity weight (cont.).

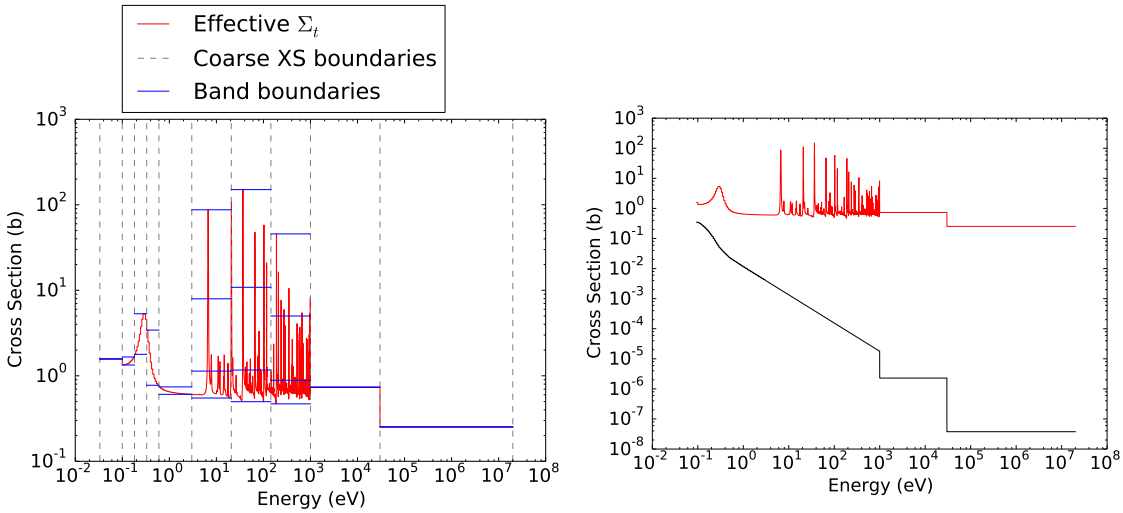

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure 7.46: Effective total cross sections for the PG-FEMG band boundary calculation for problem 3 with low-fidelity weight.

Figure 7.47a-c give work measurements. Figure 7.47a gives efficiency of the k-eigenvalue calculation (1/(work $\times$ error)), Figure 7.47b gives work per DOF, and Figure 7.47c shows total work scaled with the number of groups. Figure 7.47b shows PG-FEMG required approximately the same amount of work per DOF as MG. Both MG and PG-FEMG required substantially less work per DOF than the reference solution (horizontal line at 1.0).

Figure 7.47d-f show how error in the QOI was affected by changing energy resolution. As this resolution study did not use group sizes that resolved individual resonances, errors stayed approximately constant as the number of groups were increased. Errors did decrease as the number of bands per group were increased. For the k-eigenvalue calculation, going from 1 band per group (MG) to 2 bands per group decreased the error by a factor of 8. Increasing from 2 bands to 6 bands decreased the same error by another factor of 15-40, depending on group structure. Similar

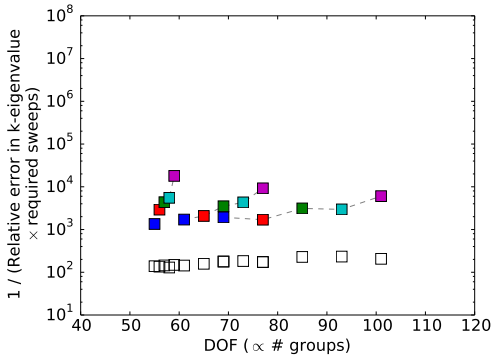

(a) k-Eigenvalue efficiency

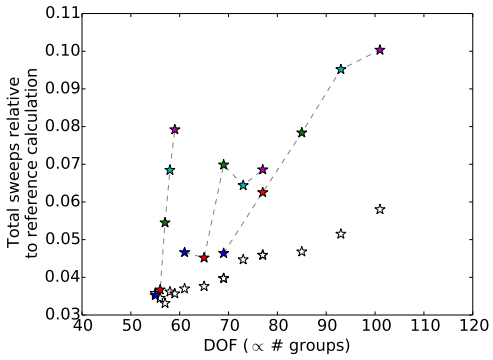

(b) Relative transport sweeps summed over bands / groups (work measurement)

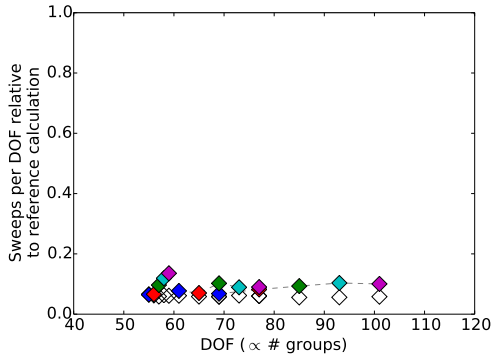

(c) Relative transport sweeps per DOF

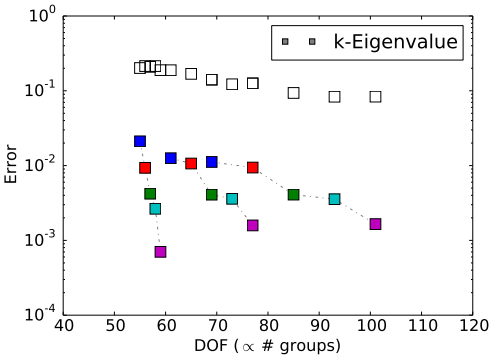

(d) k-Eigenvalue error

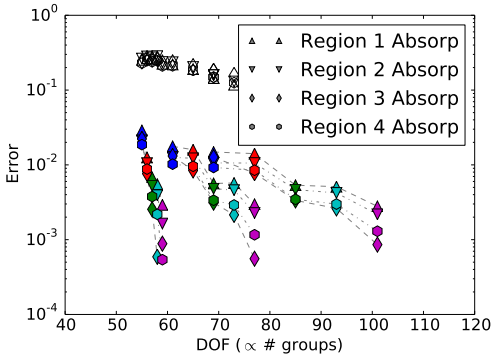

(e) Region absorption rate errors

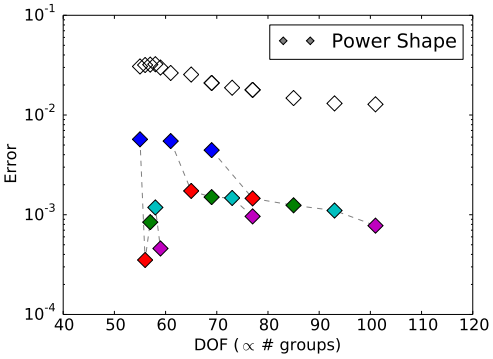

(f) Power shape

Figure 7.47: Error in various QOI relative to the reference solution for several energy group structures and numbers of bands for problem 3 with low-fidelity weight, for the resonance-only energy resolution study. The MG solution corresponds to unfilled markers; the PG-FEMG solutions are colored according to how many bands per group were used in the resonance region, be that two, three, four, five, or six. Dotted lines connect multiband cases with the same coarse group structure.

decreases occured for the absorption rates. Increasing from 1 to 2 bands per group decreased the error in the power shape by a factor of almost 5. Increasing from 2 to 6 bands per group decreased the same error by an additional factor of 6-12 depending on group structure. This problem shows error decreasing mostly monotonically as the bands per group were increased, except for the power shape error at low group counts. This is overall trend was expected, as a poor weighting spectrum was used. As more bands per group were added, the PG-FEMG solutions were better able to adapt their spectra to the true (reference) spectrum.

7.3.2.3 Energy resolution study

An energy resolution study was done to determine how various QOIs changed as a function of increasing energy fidelity (DOF) for large DOF. The energy group structure was resolved in all energy regions for this study. All errors are relative comparisons between the reference solution and the MG or PG-FEMG solutions.

Figure 7.48 gives the results from the energy resolution study. The number of groups and bands per group in the resolved resonance region were varied to determine their effect on QOI error. This resolution study focused on extremely large DOF to investigate asymptotic effects.

Figure 7.48a-c give work measurements. Figure 7.48a gives efficiency of the k-eigenvalue calculation (1/(work $\times$ error)), Figure 7.48b gives work per DOF, and Figure 7.48c shows total work scaled with the number of groups. Figure 7.48b shows PG-FEMG required more work per DOF than MG or the reference solution (horizontal line) in general, but this ratio was approximately constant over all numbers of DOF considered.

Figure 7.48d-f show how error in the QOI was affected by changing energy resolution. They give a sense of the extreme number of groups required to enter the

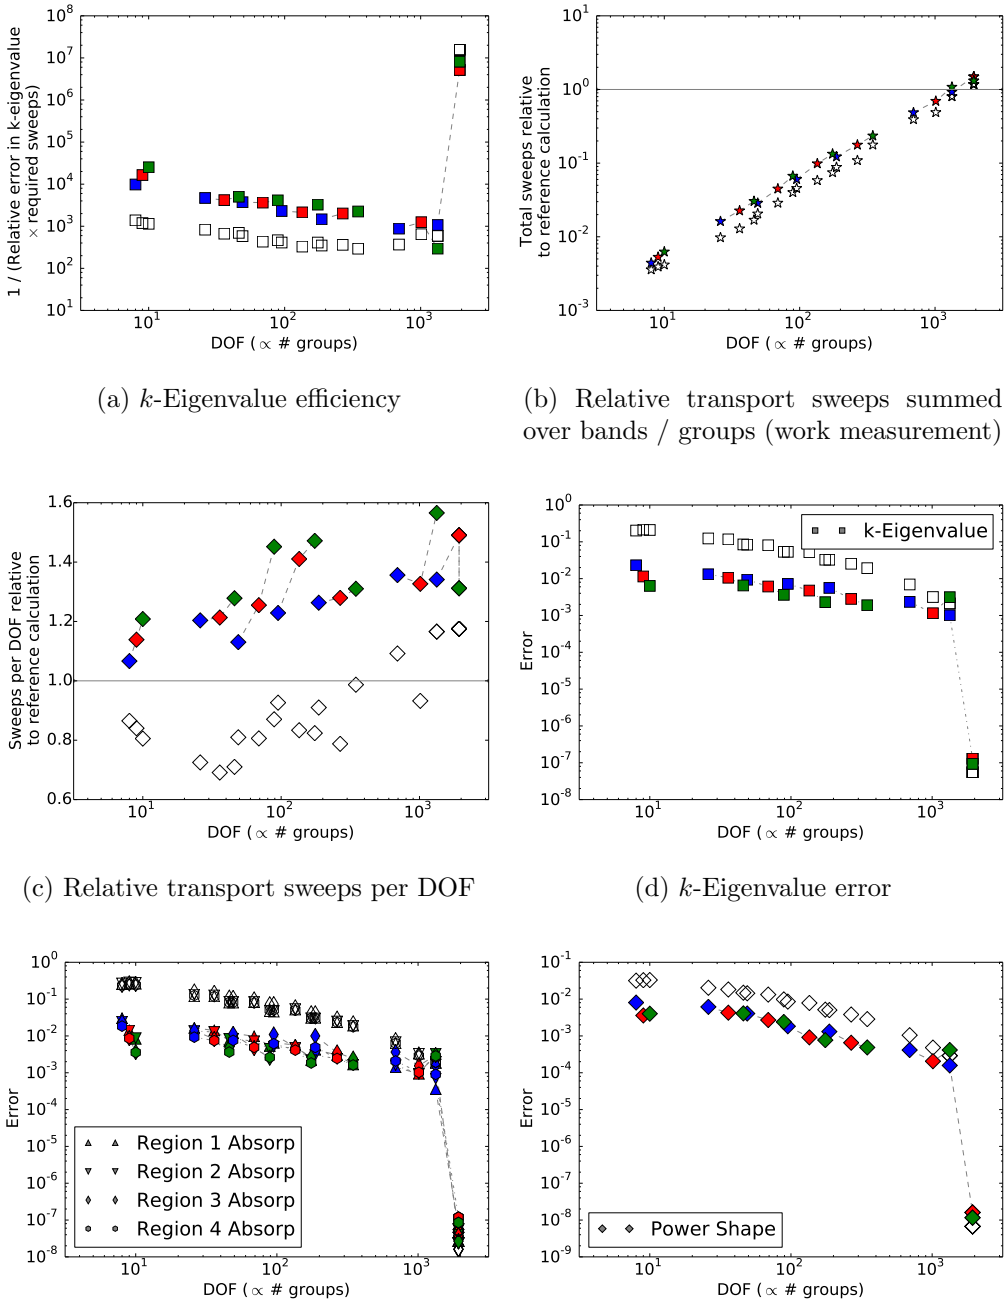

(a) k-Eigenvalue efficiency

(b) Relative transport sweeps summed over bands / groups (work measurement)

(c) Relative transport sweeps per DOF

(d) k-Eigenvalue error

(e) Region absorption rate errors

(f) Power shape

Figure 7.48: Error in various QOI relative to the reference solution for several energy group structures and numbers of bands for problem 3 with low-fidelity weight, for the full energy resolution study. The MG solution corresponds to unfilled markers; the PG-FEMG solutions are colored according to how many bands per group were used in the resonance region, be that two, three, four, five, or six. Dotted lines connect multiband cases with the same coarse group structure.

asymptotic regime. Increasing from 1 band per group (MG) to 2 bands per group decreased the error by a factor of approximately 6-10 for all QOI. Increasing from 2 to 4 bands per group decreased the errors by another factor of 2-5. The increasing error with increasing bands per group behavior at very large DOF count ($\sim$ 1000 DOF) for PG-FEMG was caused by the fall-back band boundary calculation which produced poor band boundaries.

7.4 Probem 4: Heterogeneous, 1-T, Multi-Fuel Pin Cells

Problem 4 is a simplified three-region, one-temperature PWR fuel-cell consisting of ringed fuel surrounded by moderator (O-16 and H-1). Inner MOX fuel (U-238, Pu-239, and O-16) is surrounded by outer LEU fuel (U-238, U-235, and O-16). A white boundary was used. Figure 7.49 gives a cartoon of the geometry.

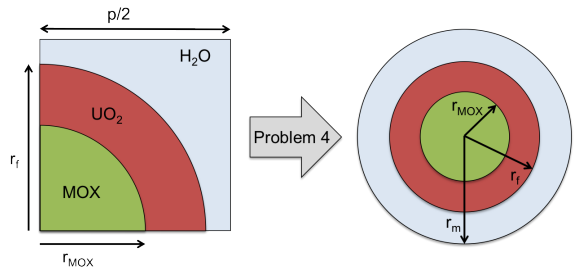

(a) Problem 4

Figure 7.49: Cartoon of the geometry and its simplification to 1D cylindrical geometry for problem 4.

Table 7.4 gives a description of each region for problem 4. Regions are defined as spatially contiguous groups of cells that share a common cross section and are ordered such that larger region numbers correspond to larger radii in the problem. When spatial averaging is done, it is over a region. Unlisted temperatures are 400 K. Tables 6.1 and 6.2 give further geometry and material information, respectively.

Table 7.4: List of spatial regions for problem 4.

Region Number	Description
1	MOX
2	LEU
3	Light water

7.4.1 Reference-calculation weighting

In this subsection, the region-averaged reference solution was used to flux-weight the (ultra-fine-group) cross sections during the condensation processes to produce the MG and PG-FEMG cross sections The same DOF were used for both MG and PG-FEMG runs.

7.4.1.1 Base problem

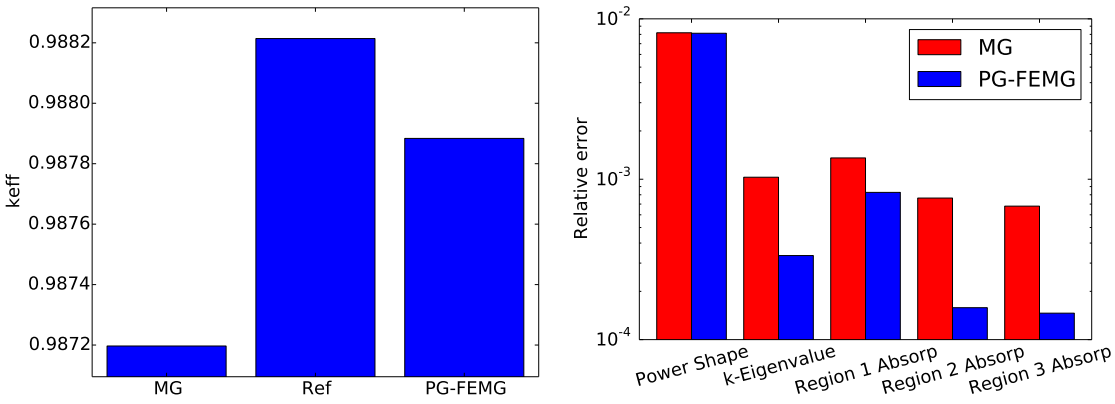

(a) k-Eigenvalues (b) Relative Errors

Figure 7.50: Comparison of eigenvalues and relative errors for problem 4 with high-fidelity weight. Values are given for the reference, MG and PG-FEMG methods.

Fig. 7.50 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. 7.50a gives the k-eigenvalues for each of the calculations. Fig. 7.50b gives the error in k-eigenvalue, power shape, and absorption rates. Both PG-FEMG and MG do approximately equally well with the power shape error. PG-FEMG has a lower error for the other QOI, beating MG by factors of approximately 3, 1.8, 5, and 5 for k-eigenvalue, absorption in the inner fuel, absorption in the outer fuel, and absorption in the moderator, respectively.

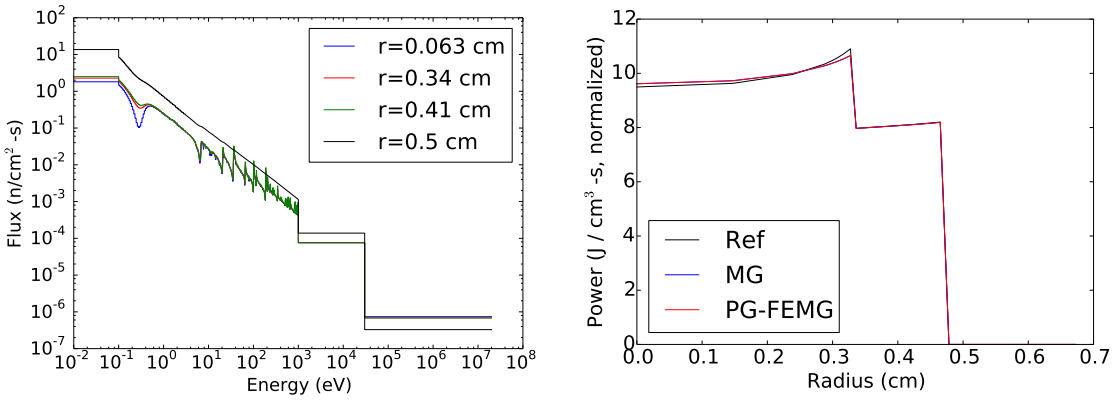

(a) Reference solution neutron source

(b) Power profile

Figure 7.51: Total source and power profile as functions of position for problem 4 with high-fidelity weight. (a) Total source (scattering plus fission) for the reference solution as a function of energy for select radial points. (b) Power profile as a function of position for the reference, MG and PG-FEMG solutions.

Fig. 7.51a gives the energy dependence of the reference solution source at several equally-spaced radial positions. It shows the source to be affected by the resonances. As the radius is increased, the source becomes smoother.

Fig. 7.51b gives the power profile shapes in space for the reference, MG, and PG-FEMG solutions. Note that the normalization was on this parameter, so they

all have the same integral, but may have different spatial shapes. In both cases, the reference solution shows stronger self-shielding effects near the edge of the fuel. The MG and PG-FEMG shapes track the reference shape well until the outside of the inner fuel, where the reference shows a sharper slope. The error near the edge of the inner fuel is approximately 2% for both cases. Both MG and PG-FEMG match the reference shape in the outer fuel to within the resolution of the figure.

Fig. 7.52 (ff.) shows the energy dependence of the reference, MG and PG-FEMG solutions averaged over spatial regions and for the exiting partial current. They show the detailed resonance structure captured by the reference solution and the fidelity of the MG and PG-FEMG reconstructions. Note that the flux dips are space- and direction-dependent. Since a reference solution was used for both cross section condensation and for flux reexpansion, we expect and find both the MG and PG-FEMG solutions to have region-averaged reconstructions accurate to the reference solution.

Fig. 7.54a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. 7.54b gives the effective total cross section and the weighting spectrum used. For this calculation, a high-fidelity weighting spectrum was used (note resonances in b).

7.4.1.2 Resonance-only energy resolution study

An energy resolution study was done to determine how various QOIs changed as a function of increasing energy fidelity (DOF) for moderate DOF. The energy group structure was always fully resolved in the non-resonance regions and changed only in the resolved-resonance region. All errors are relative comparisons between the reference solution and the MG or PG-FEMG solutions.

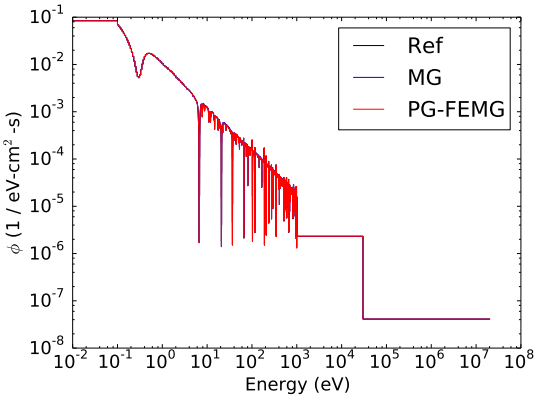

(a) Region 1 average flux

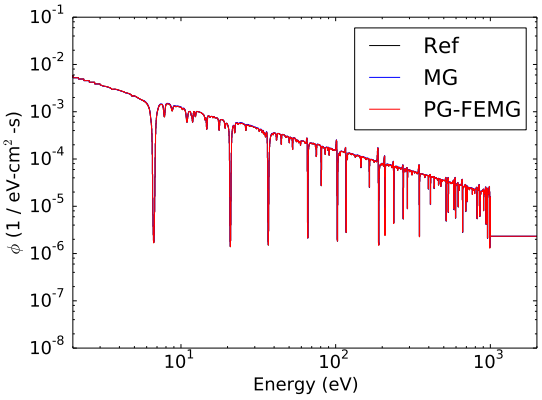

(b) Region 1 average flux (zoom)

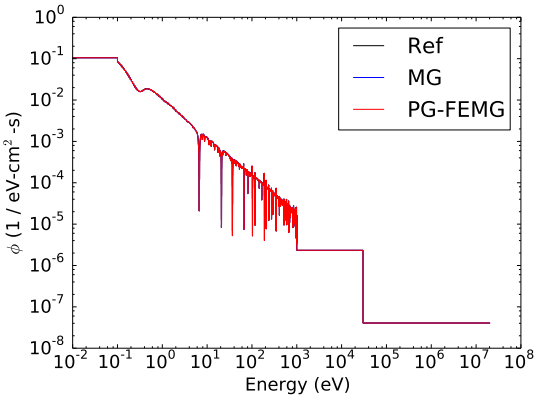

(c) Region 2 average flux

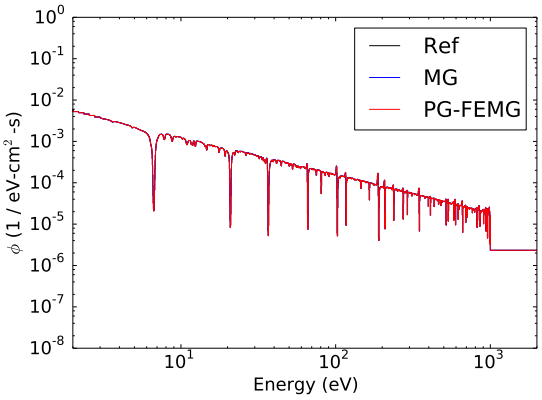

(d) Region 2 average flux (zoom)

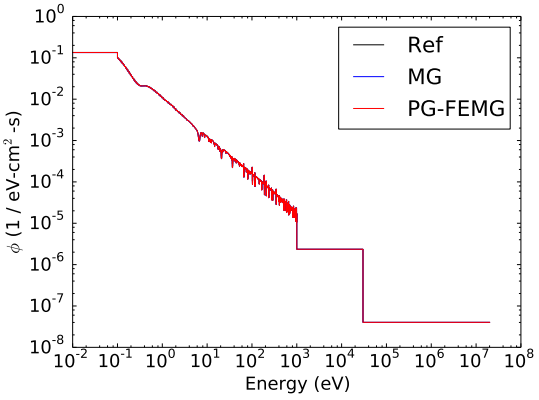

(e) Region 3 average flux

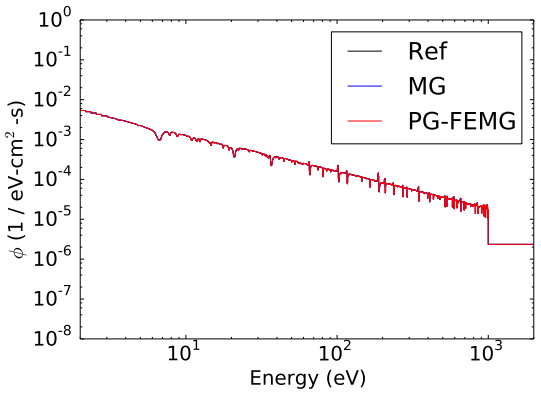

(f) Region 3 average flux (zoom)

Figure 7.52: Energy spectra averaged over spatial regions and at problem boundary for problem 4 with high-fidelity weight.

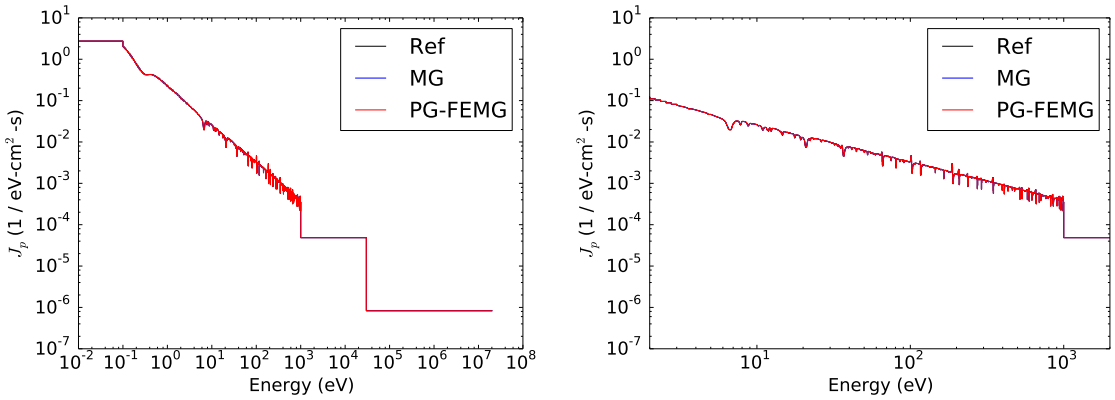

(a) Outgoing partial current

(b) Outgoing partial current (zoom)

Figure 7.53: Energy spectra averaged over spatial regions and at problem boundary for problem 4 with high-fidelity weight (cont.).

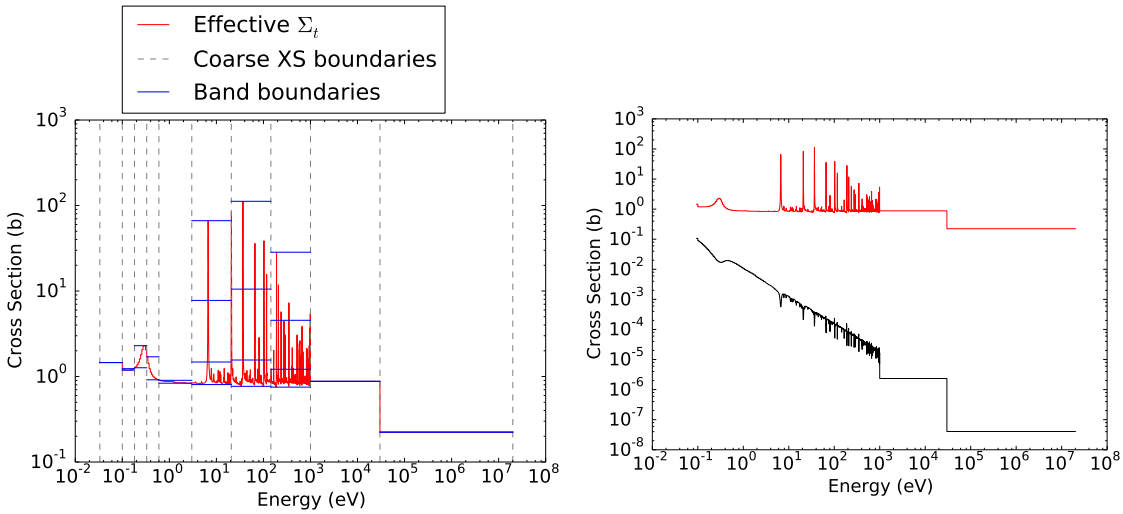

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure 7.54: Effective total cross sections for the PG-FEMG band boundary calculation for problem 4 with high-fidelity weight.

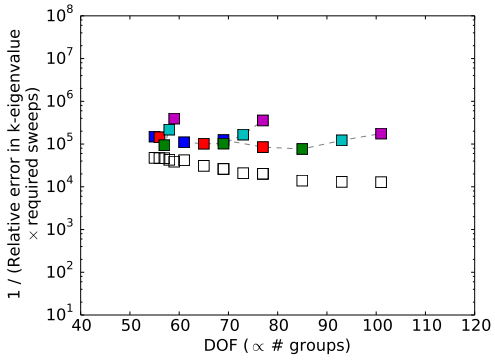

(a) k-Eigenvalue efficiency

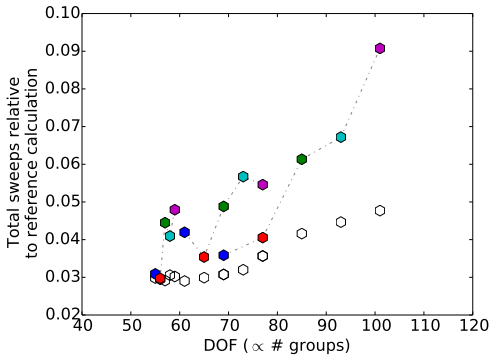

(b) Relative transport sweeps summed over bands / groups (work measurement)

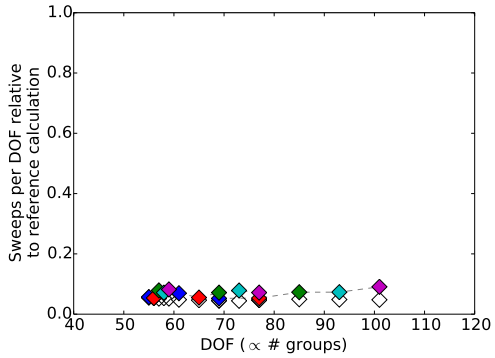

(c) Relative transport sweeps per DOF

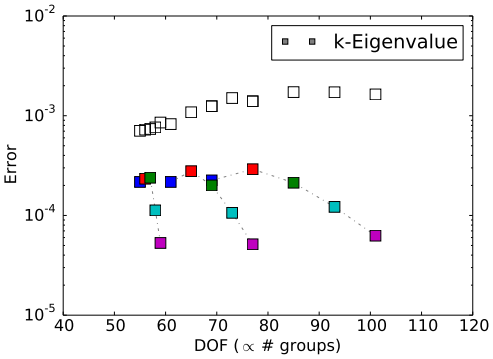

(d) k-Eigenvalue error

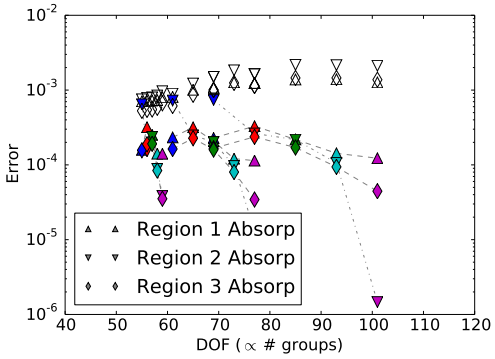

(e) Region absorption rate errors

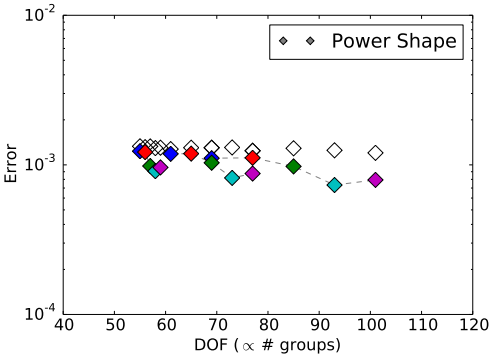

(f) Power shape

Figure 7.55: Error in various QOI relative to the reference solution for several energy group structures and numbers of bands for problem 4 with high-fidelity weight, for the resonance-only energy resolution study. The MG solution corresponds to unfilled markers; the PG-FEMG solutions are colored according to how many bands per group were used in the resonance region, be that two, three, four, five, or six. Dotted lines connect multiband cases with the same coarse group structure.

Figure 7.55 gives the results from the energy resolution study. The number of groups and bands per group in the resolved resonance region were varied to determine their effect on QOI error. This resolution study resolved (to the reference level) all non-resonance energies and used varying yet moderate numbers of DOF within the resolved resonance region.

Figure 7.55a-c give work measurements. Figure 7.55a gives efficiency of the k-eigenvalue calculation (1/(work $\times$ error)), Figure 7.55b gives work per DOF, and Figure 7.55c shows total work scaled with the number of groups. Figure 7.55b shows PG-FEMG required approximately the same amount of work per DOF as MG. Both MG and PG-FEMG required substantially less work per DOF than the reference solution (horizontal line at 1.0).

Figure 7.55d-f show how error in the QOI was affected by changing energy resolution. As this resolution study did not use group sizes that resolved individual resonances, errors stayed approximately constant as the number of groups were increased (for PG-FEMG) or increased as the number of groups increased (for MG). Errors did decrease as the number of bands per group were increased. For the k-eigenvalue calculation, going from 1 band per group (MG) to 2 bands per group decreased the error by a factor of approximately 3-9, depending on group structure. Increasing from 2 bands to 6 bands decreased the same error by another factor of approximately 4.5. The absorption rate errors were always smaller with more bands per group, usually by a factor of 5 or more. Using more than 2 bands per group was important for attaining a small absorption error in the outer fuel. Increasing from 1 to 2 bands per group had a negligible effect on power shape error. Increasing from 2 to 6 bands per group decreased the same error by a factor of approximately 1.3-2, depending on group structure. Increasing bands per group generally monotonically decreased QOI error except for 3 bands per group and the k-eigenvalue error, and 6

bands per group and the power shape error, both of which always increased.

7.4.1.3 Energy resolution study

An energy resolution study was done to determine how various QOIs changed as a function of increasing energy fidelity (DOF) for large DOF. The energy group structure was resolved in all energy regions for this study. All errors are relative comparisons between the reference solution and the MG or PG-FEMG solutions.

Figure 7.56 gives the results from the energy resolution study. The number of groups and bands per group in the resolved resonance region were varied to determine their effect on QOI error. This resolution study focused on extremely large DOF to investigate asymptotic effects.

Figure 7.56a-c give work measurements. Figure 7.56a gives efficiency of the k-eigenvalue calculation (1/(work $\times$ error)), Figure 7.56b gives work per DOF, and Figure 7.56c shows total work scaled with the number of groups. Figure 7.56b shows PG-FEMG required more work per DOF than MG or the reference solution (horizontal line) in general, but this ratio was approximately constant over all numbers of DOF considered.

Figure 7.56d-f show how error in the QOI was affected by changing energy resolution. They give a sense of the extreme number of groups required to enter the asymptotic regime. Increasing from 1 band per group (MG) to 2 bands per group decreased the error by a factor of approximately 2-7 for k-eigenvalue. Increasing from 2 to 4 bands per group did not substantially decrease this error. Only for very large DOF count (200+) did increasing the bands per group lower power shape error. Absorption rate errors were a mixed bag, showing non-monotonic behavior in general while generally being lower when at least 2 bands per group were used. The increasing error with increasing bands per group behavior at very large DOF count

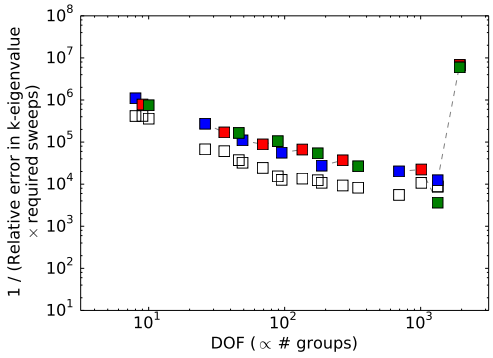

(a) k-Eigenvalue efficiency

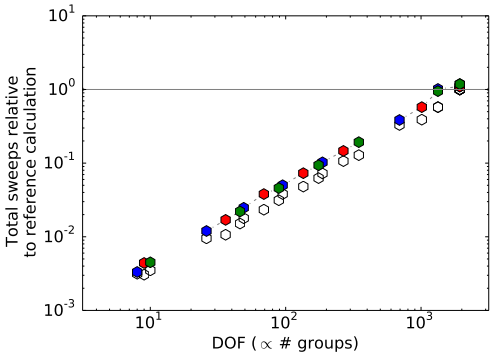

(b) Relative transport sweeps summed over bands / groups (work measurement)

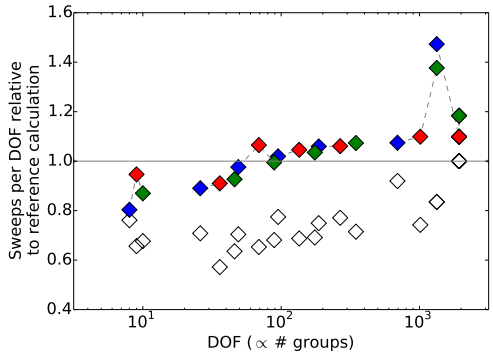

(c) Relative transport sweeps per DOF

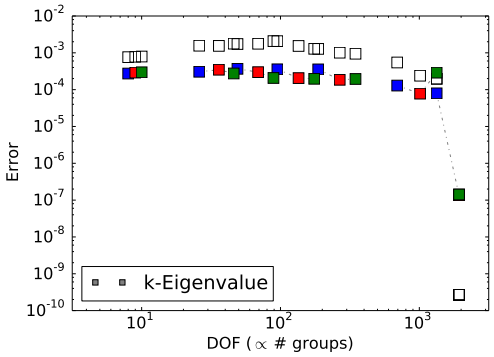

(d) k-Eigenvalue error

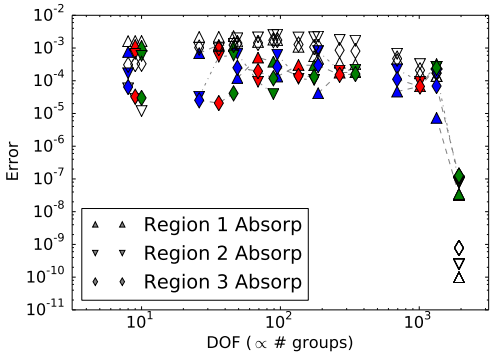

(e) Region absorption rate errors

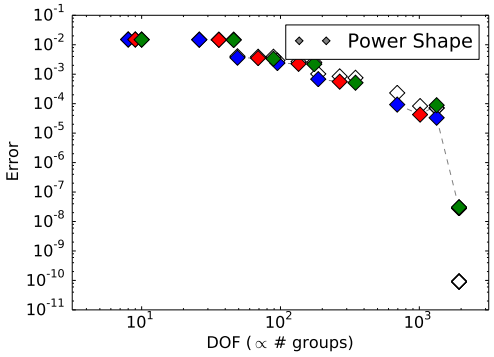

(f) Power shape

Figure 7.56: Error in various QOI relative to the reference solution for several energy group structures and numbers of bands for problem 4 with high-fidelity weight, for the full energy resolution study. The MG solution corresponds to unfilled markers; the PG-FEMG solutions are colored according to how many bands per group were used in the resonance region, be that two, three, four, five, or six. Dotted lines connect multiband cases with the same coarse group structure.

($\sim$ 1000 DOF) for PG-FEMG was caused by the fall-back band boundary calculation which produced poor band boundaries.

7.4.2 Generic weighting

In this subsection, a generic, low-fidelity solution was used to weight the cross sections during the condensation process, namely the iwt=5 option in NJOY (mid-life PWR spectrum with O-16 resonances). The same group structure was used as in the reference-weighted results.

7.4.2.1 Base problem

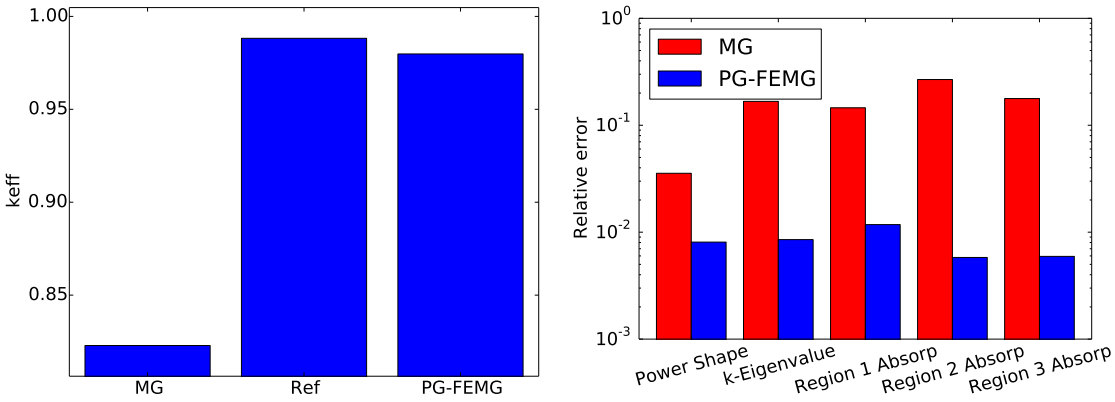

(a) k-Eigenvalues (b) Relative Errors

Figure 7.57: Comparison of eigenvalues and relative errors for problem 4 with low-fidelity weight. Values are given for the reference, MG and PG-FEMG methods.

Fig. 7.57 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. 7.57a gives the k-eigenvalues for each of the calculations. Fig. 7.57b gives the error in k-eigenvalue, power shape, and absorption rates. PG-FEMG has a lower error for the all QOI, beating MG by factors of approximately

4, 22, 12, 42, and 30 for power shape, k-eigenvalue, absorption in the inner fuel, absorption in the outer fuel, and absorption in the moderator, respectively.

For MG, the increases in error commensurate with changing from a high-fidelity to a low-fidelity weighting spectrum are approximately a factor of 4 for power shape and factors of 100-350 for the other QOI. For PG-FEMG, these same factors are approximately 1 for power shape, 25 for k-eigenvalue, 15 for absorption in the inner fuel, and 40-45 for absorptions in the outer fuel and moderator.

A trend is emerging that MG has errors of order 100-300 times greater for k-eigenvalue and absorption rates when given a generic weighting compared to the errors with the reference weighting. PG-FEMG has errors of order 10-50 times greater for the same case. For power shape, MG has errors of 2-5 times larger and PG-FEMG has approximately the same error.

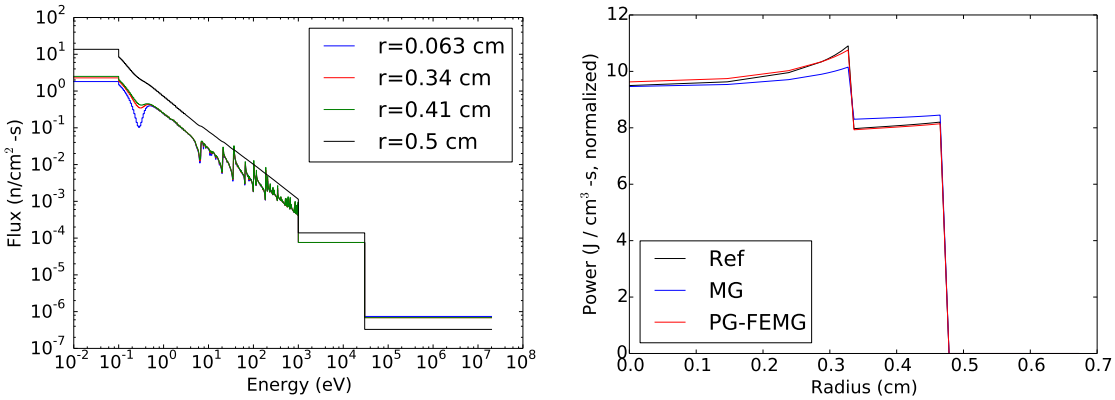

(a) Reference solution neutron source

(b) Power profile

Figure 7.58: Total source and power profile as functions of position for problem 4 with low-fidelity weight. (a) Total source (scattering plus fission) for the reference solution as a function of energy for select radial points. (b) Power profile as a function of position for the reference, MG and PG-FEMG solutions.

Fig. 7.58a gives the energy dependence of the reference solution source at several equally-spaced radial positions. It shows the source to be affected by the resonances. As the radius is increased, the source becomes smoother.

Fig. 7.58b gives the power profile shapes in space for the reference, MG, and PG-FEMG solutions. Note that the normalization was on this parameter, so they all have the same integral, but may have different spatial shapes. In both cases, the reference solution shows stronger self-shielding effects near the edge of the fuel. The MG and PG-FEMG shapes track the reference shape well until the outside of the inner fuel, where the reference shows a sharper slope. The error near the edge of the inner fuel is approximately 1% for PG-FEMG and 7% for MG. The error near the outer edge of the outer fuel is approximately 1% for PG-FEMG and 3% for MG.

Fig. 7.59 (ff.) shows the energy dependence of the reference, MG and PG-FEMG solutions averaged over spatial regions and for the exiting partial current. They show the detailed resonance structure captured by the reference solution and the fidelity of the MG and PG-FEMG reconstructions. Note that the flux dips are space- and direction-dependent. Since a low-fidelity, smooth reference flux was used for both cross section condensation and for flux reexpansion, we expect and find the MG solution to have an inaccurate reconstruction. We find the PG-FEMG solution to have a much more accurate region-averaged reconstruction to the reference solution.

Fig. 7.61a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. 7.61b gives the effective total cross section and the weighting spectrum used. For this calculation, a low-fidelity weighting spectrum was used (note smoothness in b).

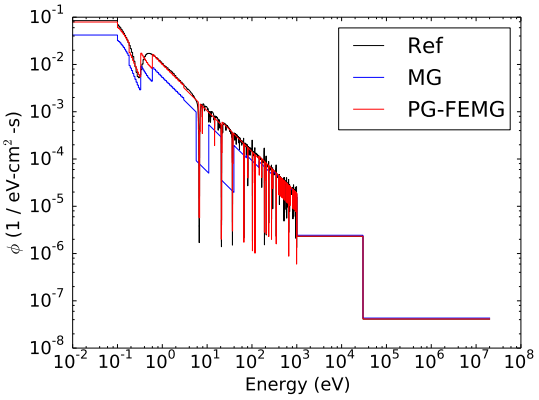

(a) Region 1 average flux

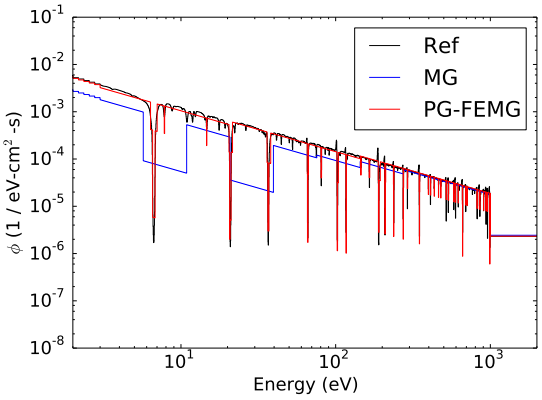

(b) Region 1 average flux (zoom)

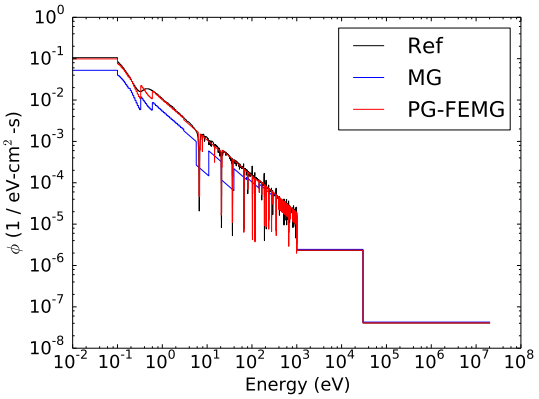

(c) Region 2 average flux

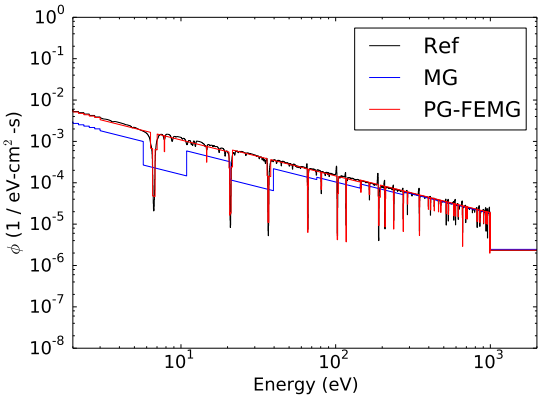

(d) Region 2 average flux (zoom)

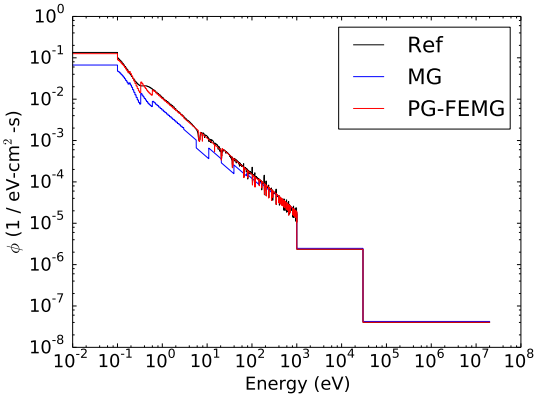

(e) Region 3 average flux

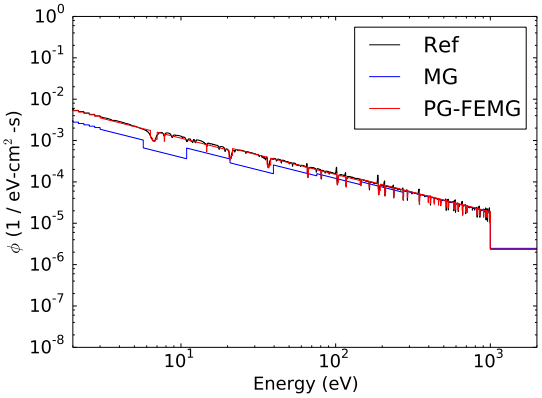

(f) Region 3 average flux (zoom)

Figure 7.59: Energy spectra averaged over spatial regions and at problem boundary for problem 4 with low-fidelity weight.

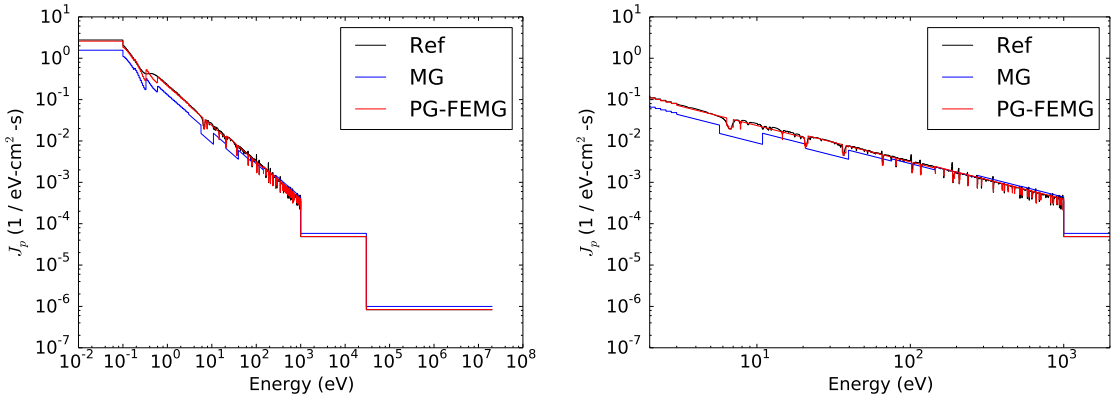

(a) Outgoing partial current

(b) Outgoing partial current (zoom)

Figure 7.60: Energy spectra averaged over spatial regions and at problem boundary for problem 4 with low-fidelity weight (cont.).

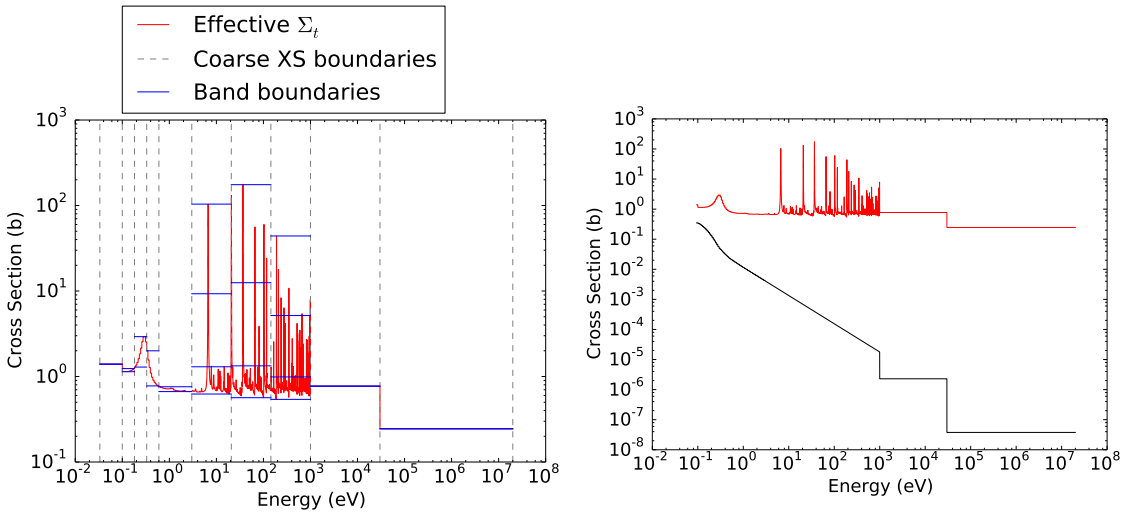

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure 7.61: Effective total cross sections for the PG-FEMG band boundary calculation for problem 4 with low-fidelity weight.

7.4.2.2 Resonance-only energy resolution study

An energy resolution study was done to determine how various QOIs changed as a function of increasing energy fidelity (DOF) for moderate DOF. The energy group structure was always fully resolved in the non-resonance regions and changed only in the resolved-resonance region. All errors are relative comparisons between the reference solution and the MG or PG-FEMG solutions.

Figure 7.62 gives the results from the energy resolution study. The number of groups and bands per group in the resolved resonance region were varied to determine their effect on QOI error. This resolution study resolved (to the reference level) all non-resonance energies and used varying yet moderate numbers of DOF within the resolved resonance region.

Figure 7.62a-c give work measurements. Figure 7.62a gives efficiency of the k-eigenvalue calculation (1/(work $\times$ error)), Figure 7.62b gives work per DOF, and Figure 7.62c shows total work scaled with the number of groups. Figure 7.62b shows PG-FEMG required approximately the same amount of work per DOF as MG. Both MG and PG-FEMG required substantially less work per DOF than the reference solution (horizontal line at 1.0).

Figure 7.62d-f show how error in the QOI was affected by changing energy resolution. As this resolution study did not use group sizes that resolved individual resonances, errors stayed approximately constant as the number of groups were increased. Errors did decrease as the number of bands per group were increased. For the k-eigenvalue calculation, going from 1 band per group (MG) to 2 bands per group decreased the error by a factor of 8. Increasing from 2 bands to 6 bands decreased the same error by another factor of 9-35, depending on group structure. Similar decreases occured for the absorption rates. Increasing from 1 to 2 bands per group

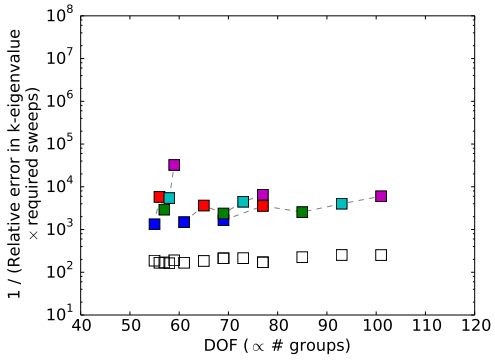

(a) k-Eigenvalue efficiency

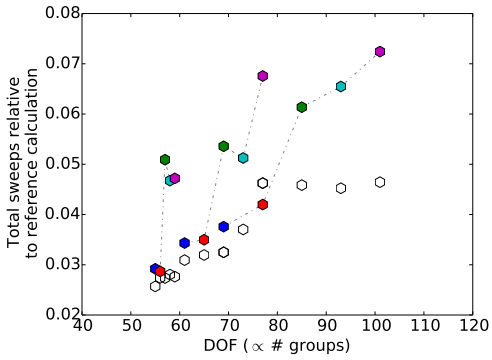

(b) Relative transport sweeps summed over bands / groups (work measurement)

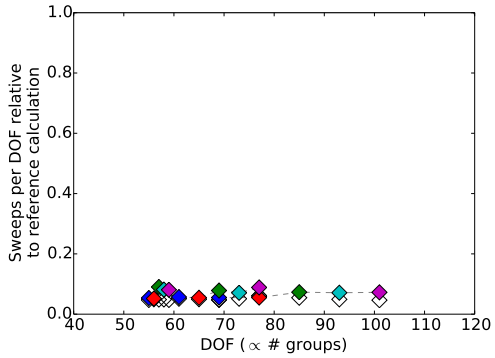

(c) Relative transport sweeps per DOF

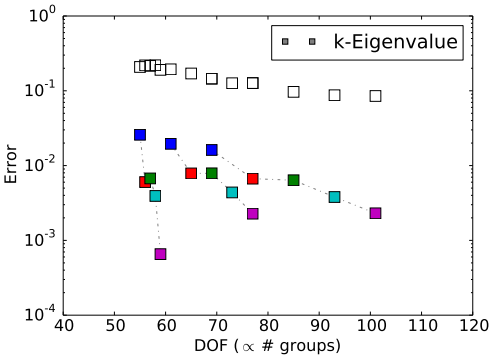

(d) k-Eigenvalue error

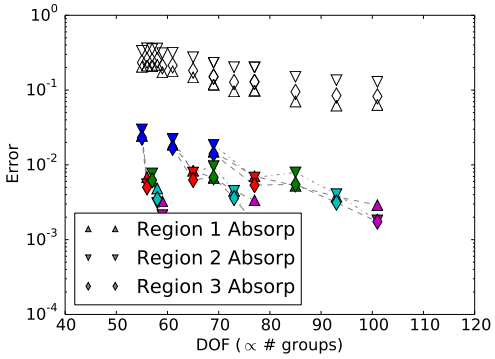

(e) Region absorption rate errors

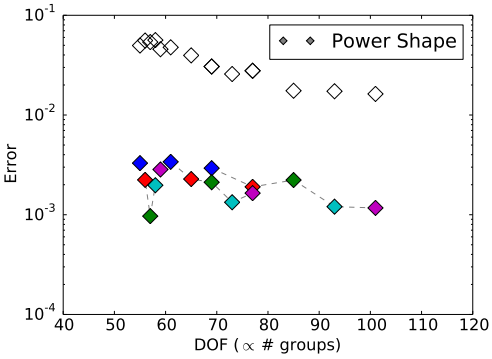

(f) Power shape

Figure 7.62: Error in various QOI relative to the reference solution for several energy group structures and numbers of bands for problem 4 with low-fidelity weight, for the resonance-only energy resolution study. The MG solution corresponds to unfilled markers; the PG-FEMG solutions are colored according to how many bands per group were used in the resonance region, be that two, three, four, five, or six. Dotted lines connect multiband cases with the same coarse group structure.

decreased the error in the power shape by a factor of approximately 7-15, depending on group structure. Increasing from 2 to 6 bands per group decreased the same error by an additional factor of 1-2.5 depending on group structure. This problem shows error decreasing mostly monotonically as the bands per group were increased, except for the power shape error and 6 bands per group and/or low group counts.

7.4.2.3 Energy resolution study

An energy resolution study was done to determine how various QOIs changed as a function of increasing energy fidelity (DOF) for large DOF. The energy group structure was resolved in all energy regions for this study. All errors are relative comparisons between the reference solution and the MG or PG-FEMG solutions.

Figure 7.63 gives the results from the energy resolution study. The number of groups and bands per group in the resolved resonance region were varied to determine their effect on QOI error. This resolution study focused on extremely large DOF to investigate asymptotic effects.

Figure 7.63a-c give work measurements. Figure 7.63a gives efficiency of the k-eigenvalue calculation (1/(work $\times$ error)), Figure 7.63b gives work per DOF, and Figure 7.63c shows total work scaled with the number of groups. Figure 7.63b shows PG-FEMG required more work per DOF than MG or the reference solution (horizontal line) in general, but this ratio was approximately constant over all numbers of DOF considered.

Figure 7.63d-f show how error in the QOI was affected by changing energy resolution. They give a sense of the extreme number of groups required to enter the asymptotic regime. Increasing from 1 band per group (MG) to 2 bands per group decreased the k-eigenvalue error by a factor of approximately 6-10. Increasing from 2 to 4 bands per group decreased this error by a factor of 1.5-3. Absorption rate

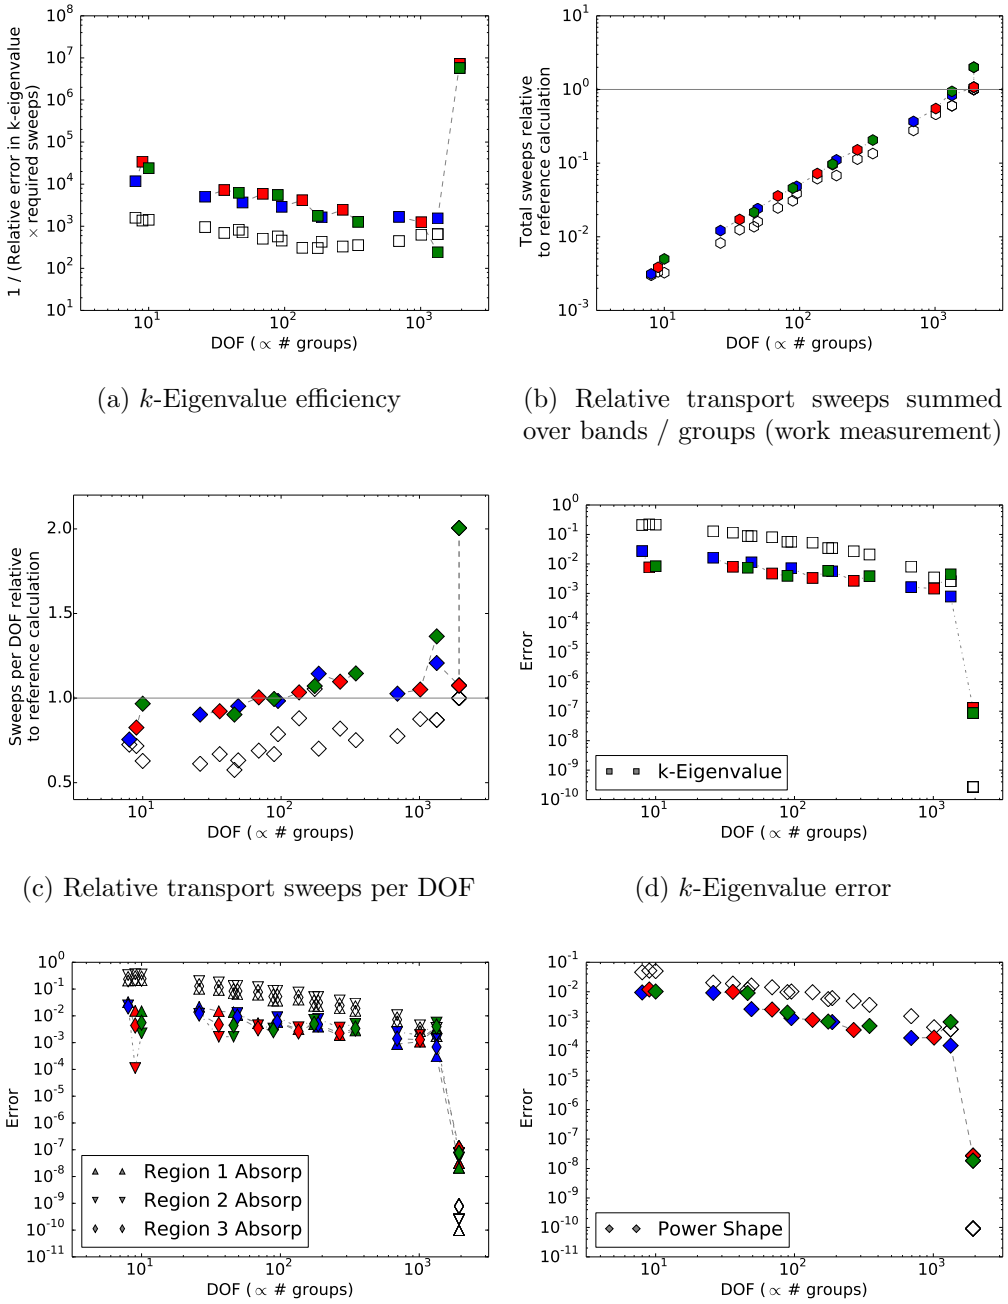

(a) k-Eigenvalue efficiency

(b) Relative transport sweeps summed over bands / groups (work measurement)

(c) Relative transport sweeps per DOF

(d) k-Eigenvalue error

(e) Region absorption rate errors

(f) Power shape

Figure 7.63: Error in various QOI relative to the reference solution for several energy group structures and numbers of bands for problem 4 with low-fidelity weight, for the full energy resolution study. The MG solution corresponds to unfilled markers; the PG-FEMG solutions are colored according to how many bands per group were used in the resonance region, be that two, three, four, five, or six. Dotted lines connect multiband cases with the same coarse group structure.

errors had similar factors, aside from non-monotonic behavior at low group counts for absorption in the outer fuel when using more than 2 bands per group. Power profile shape errors decreased by a factor of 5-6 when increasing from 1 to 2 bands per group. There was negligible decrease in error when moving to 4 bands per group from 2 bands per group. The increasing error with increasing bands per group behavior at very large DOF count ($\sim$ 1000 DOF) for PG-FEMG was caused by the fall-back band boundary calculation which produced poor band boundaries.

7.5 Problem 5: Heterogeneous, 1-T, Multi-Fuel Cylindricized Fuel Lattice

Problem 5 is a simplified three-region lattice of a MOX fuel pin (U-238, Pu-239, and O-16) surrounded by LEU fuel pins (U-238, U-235, O-16). Light water moderator (O-16 and H-1) separates the pins. A white boundary was used. Figure 7.64 gives a cartoon of the geometry.

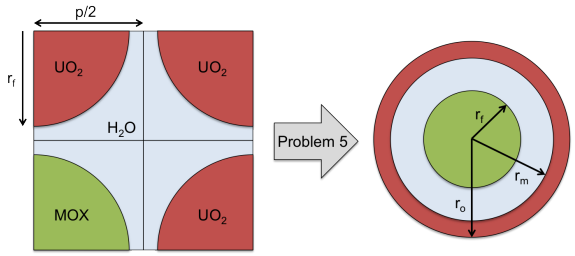

(a) Problem 5

Figure 7.64: Cartoon of the geometry and its simplification to 1D cylindrical geometry for problem 5.

Table 7.5 gives a description of each region for problem 5. Regions are defined as spatially contiguous groups of cells that share a common cross section and are ordered such that larger region numbers correspond to larger radii in the problem. When spatial averaging is done, it is over a region. Unlisted temperatures are 400

K. Tables 6.1 and 6.2 give further geometry and material information, respectively.

Table 7.5: List of spatial regions for problem 5.

Region Number	Description
1	MOX
2	Light water
3	LEU

7.5.1 Reference-calculation weighting

In this subsection, the region-averaged reference solution was used to flux-weight the (ultra-fine-group) cross sections during the condensation processes to produce the MG and PG-FEMG cross sections The same DOF were used for both MG and PG-FEMG runs.

7.5.1.1 Base problem

Fig. 7.65 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. 7.65a gives the k-eigenvalues for each of the calculations. Fig. 7.65b gives the error in k-eigenvalue, power shape, and absorption rates. Both PG-FEMG and MG do approximately equally well with the power shape error. PG-FEMG has a lower error for the other QOI, beating MG by factors of approximately 3, 1.3, 5, and 150 for k-eigenvalue, absorption in the inner fuel, absorption in the moderator, and absorption in the outer fuel, respectively.

Fig. 7.66a gives the energy dependence of the reference solution source at several equally-spaced radial positions. It shows the source to be affected by the resonances.

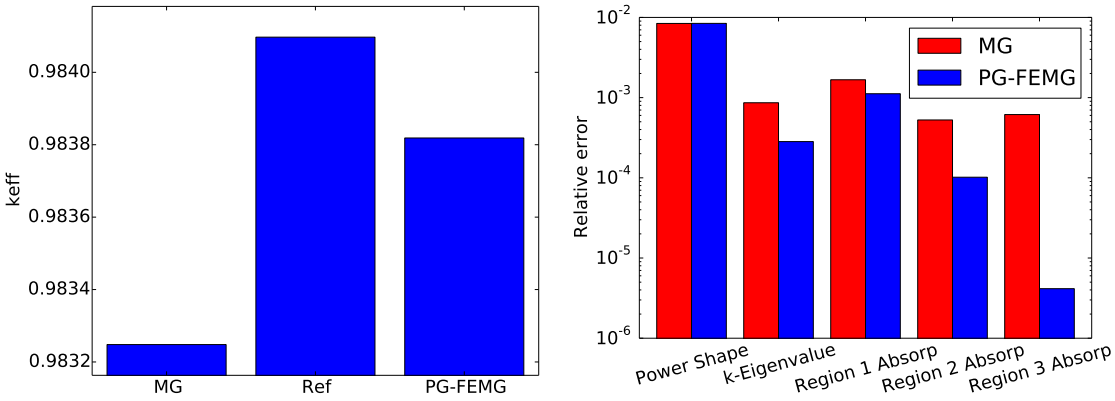

(a) k-Eigenvalues (b) Relative Errors

Figure 7.65: Comparison of eigenvalues and relative errors for problem 5 with high-fidelity weight. Values are given for the reference, MG and PG-FEMG methods.

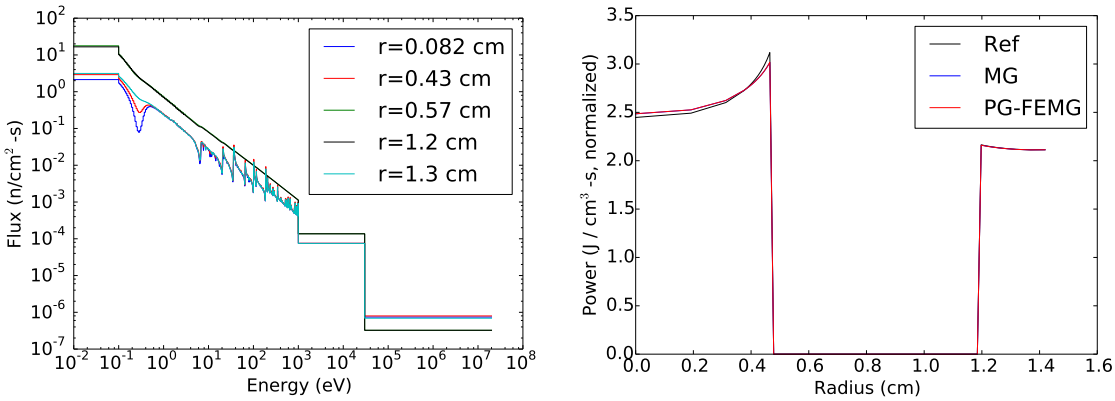

(a) Reference solution neutron source (b) Power profile

Figure 7.66: Total source and power profile as functions of position for problem 5 with high-fidelity weight. (a) Total source (scattering plus fission) for the reference solution as a function of energy for select radial points. (b) Power profile as a function of position for the reference, MG and PG-FEMG solutions.

As the radius is increased in the inner fuel, the source becomes smoother. The source in the outer fuel again exhibits resonances.

Fig. 7.66b gives the power profile shapes in space for the reference, MG, and PG-FEMG solutions. Note that the normalization was on this parameter, so they all have the same integral, but may have different spatial shapes. In both cases, the reference solution shows stronger self-shielding effects near the edge of the fuel. The MG and PG-FEMG shapes track the reference shape well until the outside of the inner fuel, where the reference shows a sharper slope. The error near the edge of the inner fuel is approximately 3% for both cases. Both MG and PG-FEMG match the reference shape in the outer fuel to within the resolution of the figure.

Fig. 7.67 (ff.) shows the energy dependence of the reference, MG and PG-FEMG solutions averaged over spatial regions and for the exiting partial current. They show the detailed resonance structure captured by the reference solution and the fidelity of the MG and PG-FEMG reconstructions. Note that the flux dips are space- and direction-dependent. Since a reference solution was used for both cross section condensation and for flux reexpansion, we expect and find both the MG and PG-FEMG solutions to have region-averaged reconstructions accurate to the reference solution.

Fig. 7.69a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. 7.69b gives the effective total cross section and the weighting spectrum used. For this calculation, a high-fidelity weighting spectrum was used (note resonances in b).

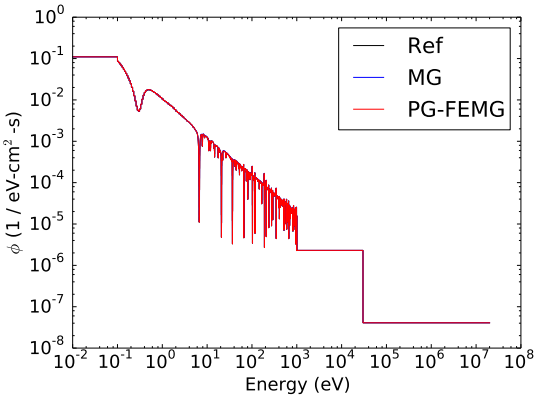

(a) Region 1 average flux

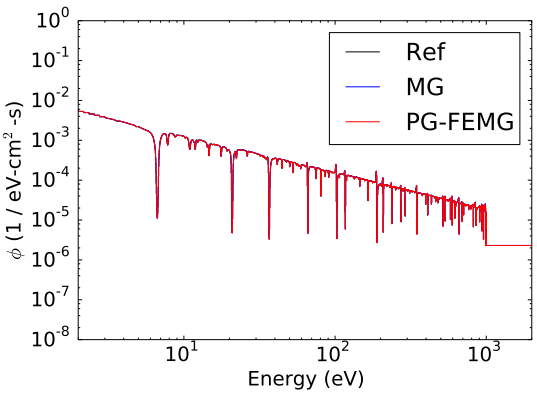

(b) Region 1 average flux (zoom)

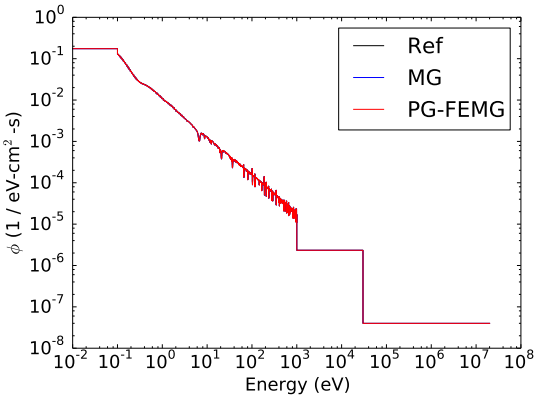

(c) Region 2 average flux

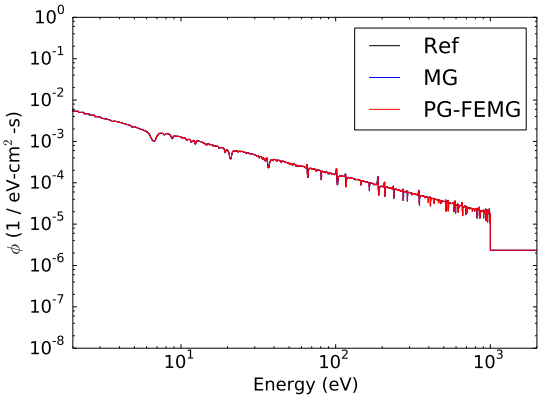

(d) Region 2 average flux (zoom)

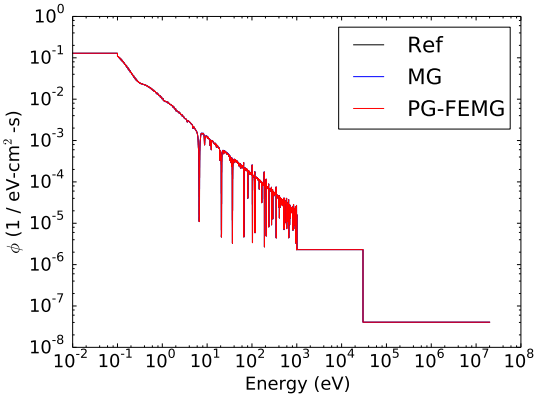

(e) Region 3 average flux

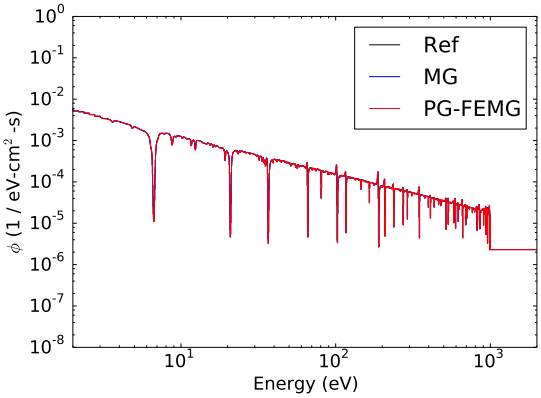

(f) Region 3 average flux (zoom)

Figure 7.67: Energy spectra averaged over spatial regions and at problem boundary for problem 5 with high-fidelity weight.

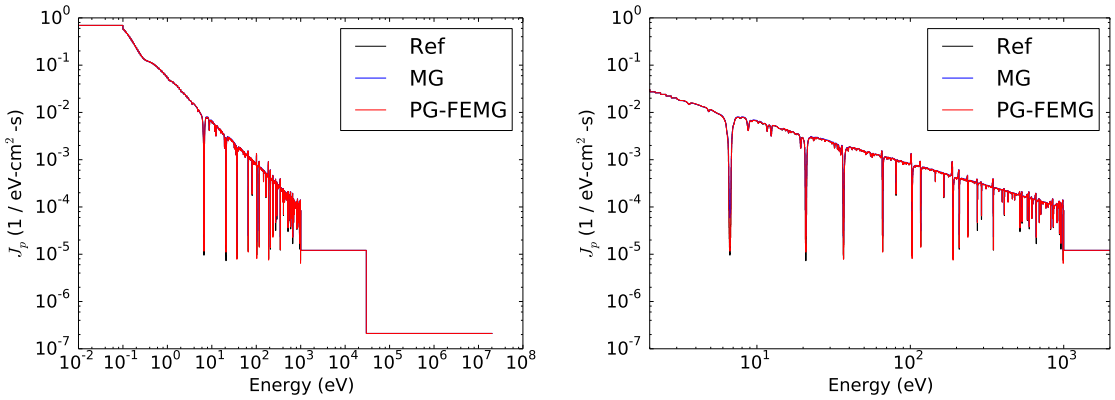

(a) Outgoing partial current

(b) Outgoing partial current (zoom)

Figure 7.68: Energy spectra averaged over spatial regions and at problem boundary for problem 5 with high-fidelity weight (cont.).

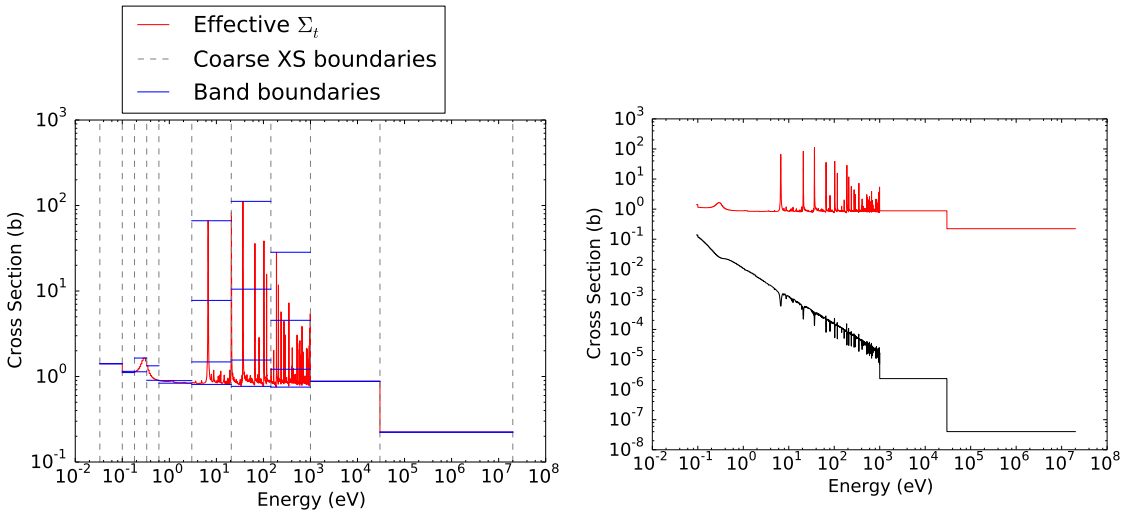

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure 7.69: Effective total cross sections for the PG-FEMG band boundary calculation for problem 5 with high-fidelity weight.

7.5.1.2 Resonance-only energy resolution study

An energy resolution study was done to determine how various QOIs changed as a function of increasing energy fidelity (DOF) for moderate DOF. The energy group structure was always fully resolved in the non-resonance regions and changed only in the resolved-resonance region. All errors are relative comparisons between the reference solution and the MG or PG-FEMG solutions.

Figure 7.70 gives the results from the energy resolution study. The number of groups and bands per group in the resolved resonance region were varied to determine their effect on QOI error. This resolution study resolved (to the reference level) all non-resonance energies and used varying yet moderate numbers of DOF within the resolved resonance region.

Figure 7.70a-c give work measurements. Figure 7.70a gives efficiency of the k-eigenvalue calculation (1/(work $\times$ error)), Figure 7.70b gives work per DOF, and Figure 7.70c shows total work scaled with the number of groups. Figure 7.70b shows PG-FEMG required approximately the same amount of work per DOF as MG. Both MG and PG-FEMG required substantially less work per DOF than the reference solution (horizontal line at 1.0).

Figure 7.70d-f show how error in the QOI was affected by changing energy resolution. As this resolution study did not use group sizes that resolved individual resonances, errors stayed approximately constant as the number of groups were increased (for PG-FEMG) or increased as the number of groups increased (for MG). Errors did decrease as the number of bands per group were increased. For the k-eigenvalue calculation, going from 1 band per group (MG) to 2 bands per group decreased the error by a factor of approximately 5-7, depending on group structure. Increasing from 2 bands to 6 bands decreased the same error by another factor of

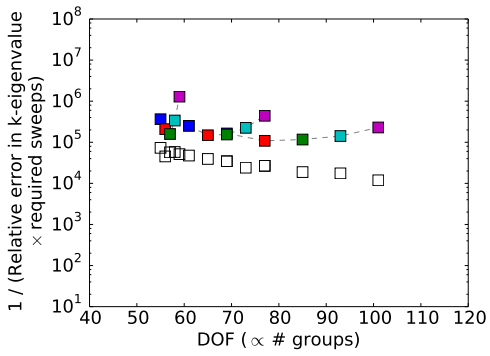

(a) k-Eigenvalue efficiency

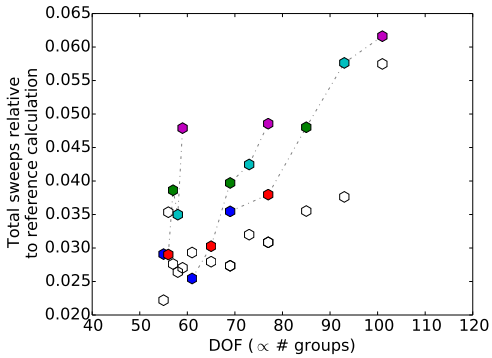

(b) Relative transport sweeps summed over bands / groups (work measurement)

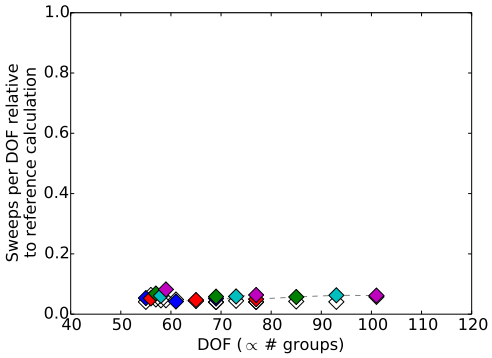

(c) Relative transport sweeps per DOF

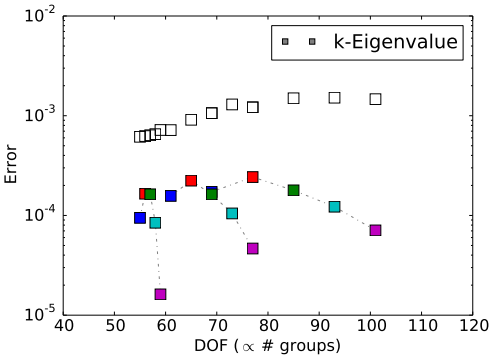

(d) k-Eigenvalue error

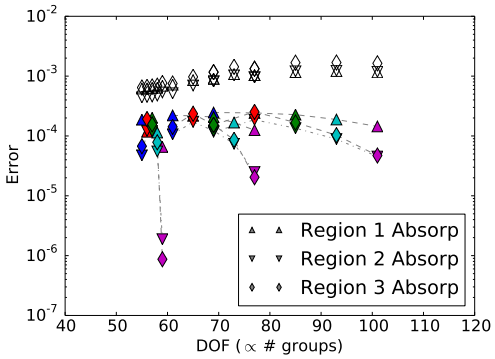

(e) Region absorption rate errors

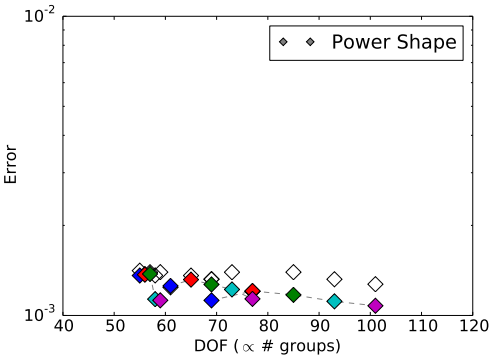

(f) Power shape

Figure 7.70: Error in various QOI relative to the reference solution for several energy group structures and numbers of bands for problem 5 with high-fidelity weight, for the resonance-only energy resolution study. The MG solution corresponds to unfilled markers; the PG-FEMG solutions are colored according to how many bands per group were used in the resonance region, be that two, three, four, five, or six. Dotted lines connect multiband cases with the same coarse group structure.

approximately 3-6. The absorption rate errors had similar factors. Increasing from 1 to 2 bands per group had a negligible effect on power shape error. Increasing from 2 to 6 bands per group decreased the same error by a factor of approximately 1.5. Increasing bands per group generally monotonically decreased QOI error except for 3 bands per group and the k-eigenvalue error; and 3 or 4 bands per group, the absorption rate and power shape errors, and low group counts.

7.5.1.3 Energy resolution study

An energy resolution study was done to determine how various QOIs changed as a function of increasing energy fidelity (DOF) for large DOF. The energy group structure was resolved in all energy regions for this study. All errors are relative comparisons between the reference solution and the MG or PG-FEMG solutions.

Figure 7.71 gives the results from the energy resolution study. The number of groups and bands per group in the resolved resonance region were varied to determine their effect on QOI error. This resolution study focused on extremely large DOF to investigate asymptotic effects.

Figure 7.71a-c give work measurements. Figure 7.71a gives efficiency of the k-eigenvalue calculation (1/(work $\times$ error)), Figure 7.71b gives work per DOF, and Figure 7.71c shows total work scaled with the number of groups. Figure 7.71b shows PG-FEMG required more work per DOF than MG or the reference solution (horizontal line) in general, but this ratio was approximately constant over all numbers of DOF considered.

Figure 7.71d-f show how error in the QOI was affected by changing energy resolution. They give a sense of the extreme number of groups required to enter the asymptotic regime. Increasing from 1 band per group (MG) to 2 bands per group decreased the k-eigenvalue error by a factor of approximately 4-6. Increasing from 2

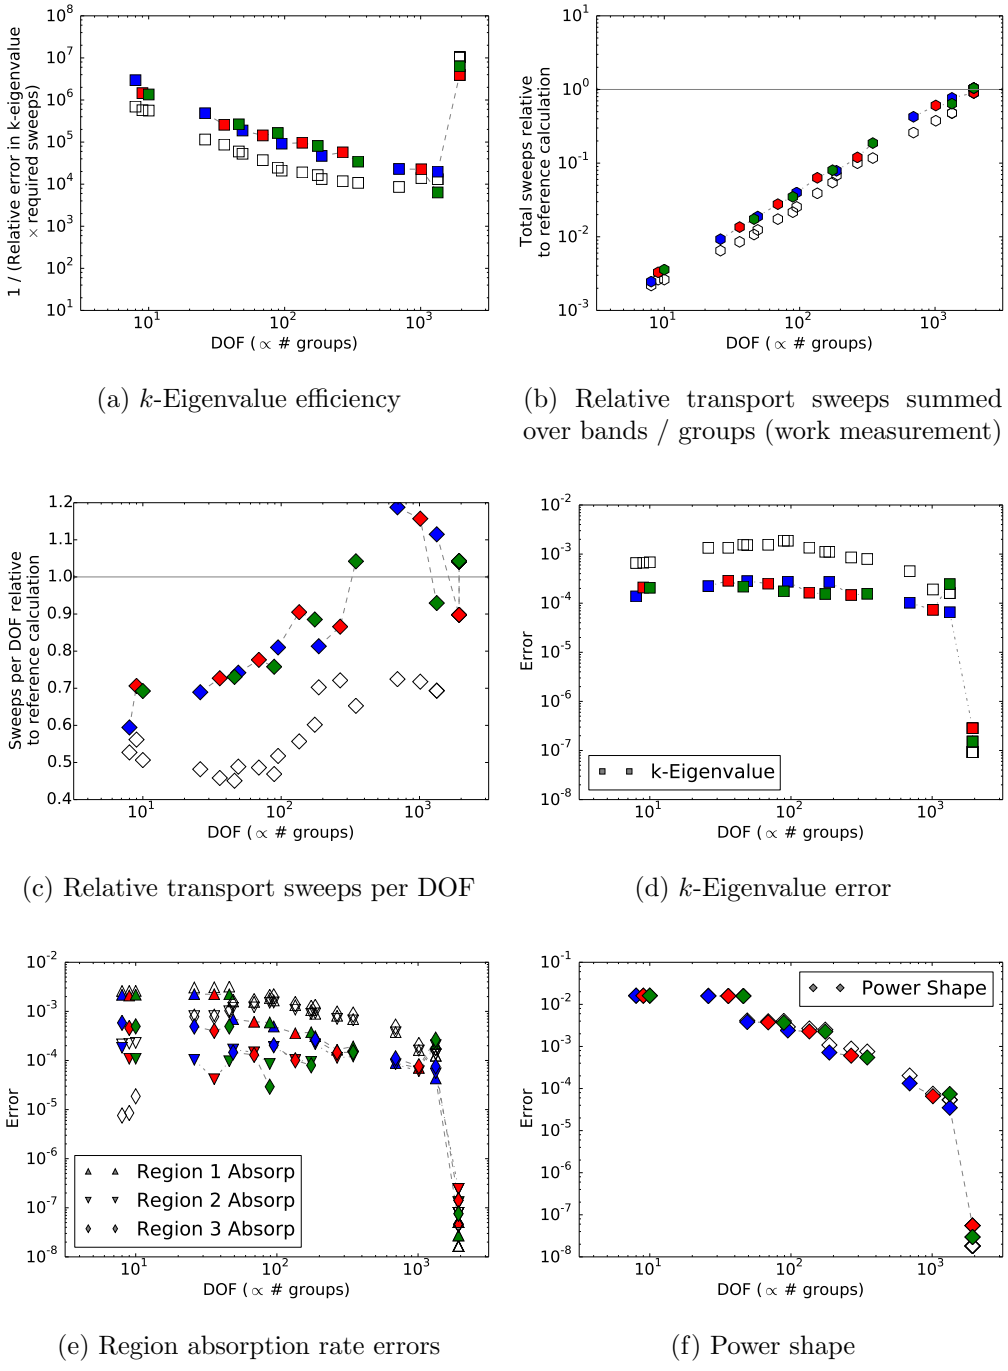

(a) k-Eigenvalue efficiency

(b) Relative transport sweeps summed over bands / groups (work measurement)

(c) Relative transport sweeps per DOF

(d) k-Eigenvalue error

(e) Region absorption rate errors

(f) Power shape

Figure 7.71: Error in various QOI relative to the reference solution for several energy group structures and numbers of bands for problem 5 with high-fidelity weight, for the full energy resolution study. The MG solution corresponds to unfilled markers; the PG-FEMG solutions are colored according to how many bands per group were used in the resonance region, be that two, three, four, five, or six. Dotted lines connect multiband cases with the same coarse group structure.

to 4 bands per group varied from increasing this error to decreasing it by a factor of 2, with the latter occuring at high group counts. Only for very large DOF count (200+) did increasing the bands per group lower power shape error. Absorption rate errors were a mixed bag, showing non-monotonic behavior at some group counts lower than ~ 100. Absorption rate errors with 2 or more bands per group were better than with 1 band per group with the following caveat. The increasing error with increasing bands per group behavior at very large DOF count ($\sim$ 1000 DOF) for PG-FEMG was caused by the fall-back band boundary calculation which produced poor band boundaries.

7.5.2 Generic weighting

In this subsection, a generic, low-fidelity solution was used to weight the cross sections during the condensation process, namely the iwt=5 option in NJOY (mid-life PWR spectrum with O-16 resonances). The same group structure was used as in the reference-weighted results.

7.5.2.1 Base problem

Fig. 7.72 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. 7.72a gives the k-eigenvalues for each of the calculations. Fig. 7.72b gives the error in k-eigenvalue, power shape, and absorption rates. PG-FEMG has a lower error for the all QOI, beating MG by factors of approximately 5, 18, 10, 25, and 30 for power shape, k-eigenvalue, absorption in the inner fuel, absorption in the moderator, and absorption in the outer fuel, respectively.

For MG, the increases in error commensurate with changing from a high-fidelity to a low-fidelity weighting spectrum are approximately a factor of 4.5 for power shape and factors of 85-350 for the other QOI. For PG-FEMG, these same factors are approximately 0.9 for power shape, 35 for k-eigenvalue, 11 for absorption in the

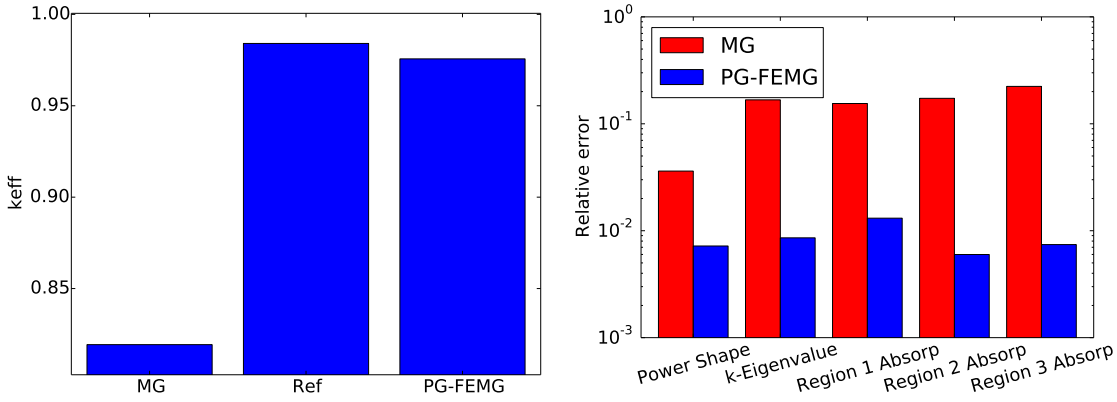

(a) k-Eigenvalues (b) Relative Errors

Figure 7.72: Comparison of eigenvalues and relative errors for problem 5 with low-fidelity weight. Values are given for the reference, MG and PG-FEMG methods.

inner fuel, 60 for absorption in the moderator, and 2000 for absorption in the outer fuel. This last value is misleading, because the error in the absorption in the outer fuel for PG-FEMG with reference weighting was abnormally low (4×10^{-6}).

Fig. 7.73a gives the energy dependence of the reference solution source at several equally-spaced radial positions. It shows the source to be affected by the resonances. As the radius is increased in the inner fuel, the source becomes smoother. The source in the outer fuel again exhibits resonances.

Fig. 7.73b gives the power profile shapes in space for the reference, MG, and PG-FEMG solutions. Note that the normalization was on this parameter, so they all have the same integral, but may have different spatial shapes. In both cases, the reference solution shows stronger self-shielding effects near the edge of the fuel. The MG and PG-FEMG shapes track the reference shape well until the outside of the inner fuel, where the reference shows a sharper slope. The error near the edge of the inner fuel is approximately 1.5% for PG-FEMG and 11% for MG. The error near the outer edge of the outer fuel is approximately <1% for PG-FEMG and 2% for MG.

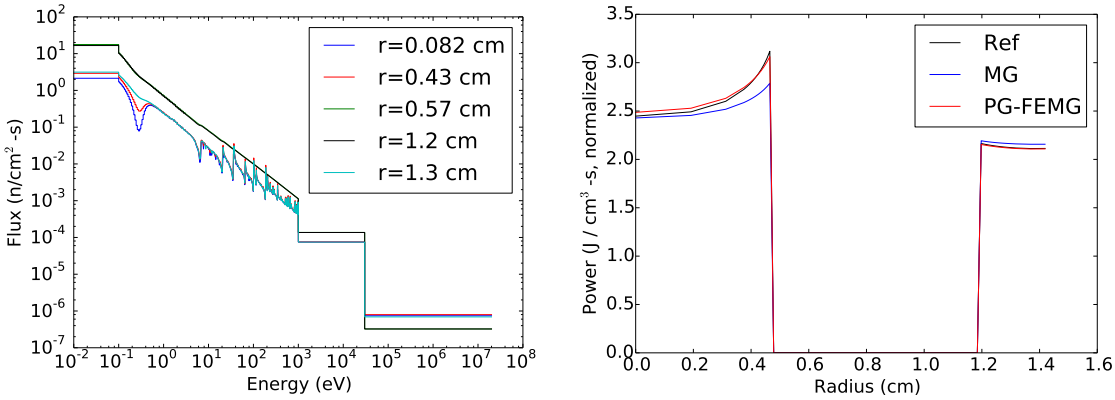

(a) Reference solution neutron source

(b) Power profile

Figure 7.73: Total source and power profile as functions of position for problem 5 with low-fidelity weight. (a) Total source (scattering plus fission) for the reference solution as a function of energy for select radial points. (b) Power profile as a function of position for the reference, MG and PG-FEMG solutions.

Fig. 7.74 (ff.) shows the energy dependence of the reference, MG and PG-FEMG solutions averaged over spatial regions and for the exiting partial current. They show the detailed resonance structure captured by the reference solution and the fidelity of the MG and PG-FEMG reconstructions. Note that the flux dips are space- and direction-dependent. Since a low-fidelity, smooth reference flux was used for both cross section condensation and for flux reexpansion, we expect and find the MG solution to have an inaccurate reconstruction. We find the PG-FEMG solution to have a much more accurate region-averaged reconstruction to the reference solution.

Fig. 7.76a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. 7.76b gives the effective total cross section and the weighting spectrum used. For this calculation, a low-fidelity weighting spectrum was used (note smoothness in b).

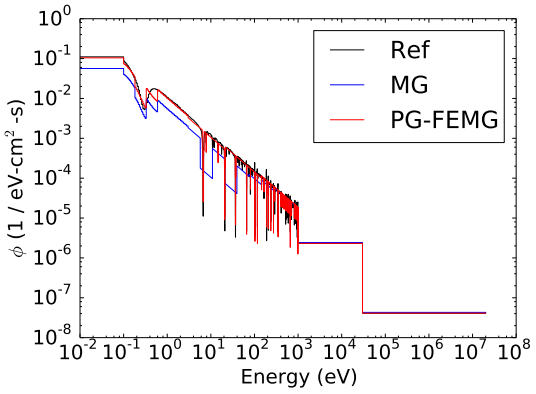

(a) Region 1 average flux

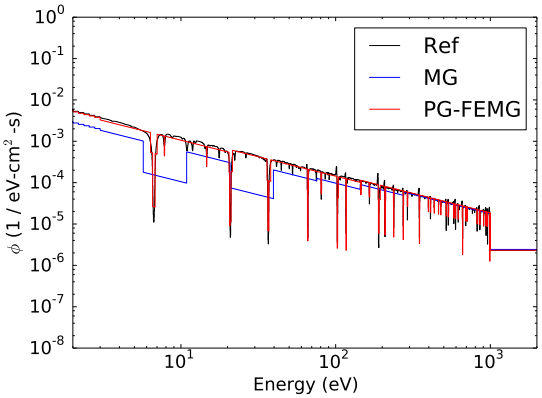

(b) Region 1 average flux (zoom)

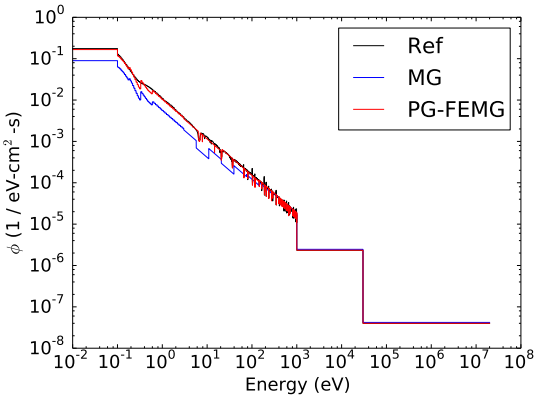

(c) Region 2 average flux

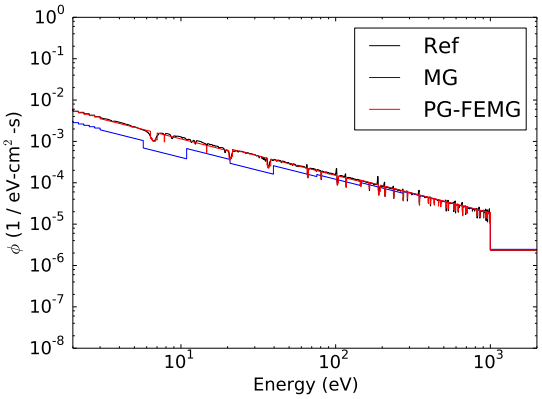

(d) Region 2 average flux (zoom)

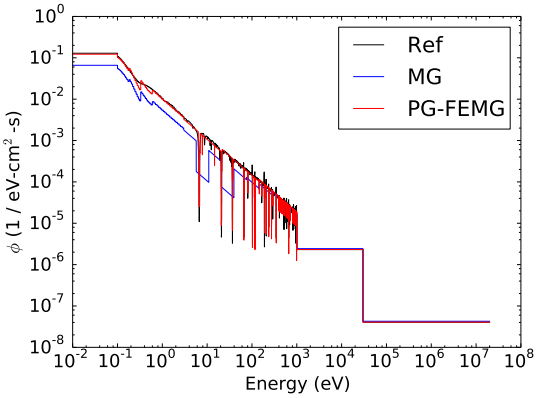

(e) Region 3 average flux

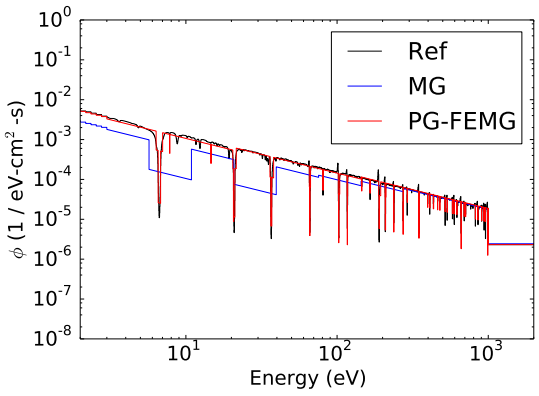

(f) Region 3 average flux (zoom)

Figure 7.74: Energy spectra averaged over spatial regions and at problem boundary for problem 5 with low-fidelity weight.

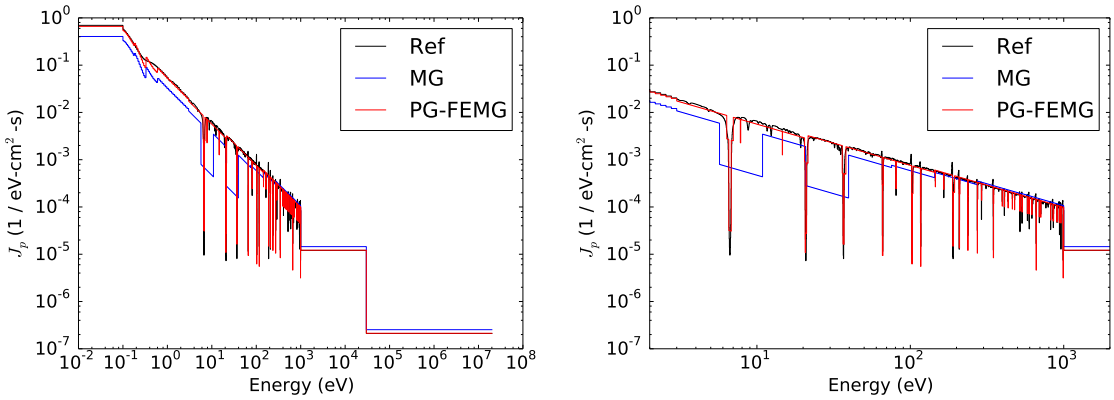

(a) Outgoing partial current

(b) Outgoing partial current (zoom)

Figure 7.75: Energy spectra averaged over spatial regions and at problem boundary for problem 5 with low-fidelity weight (cont.).

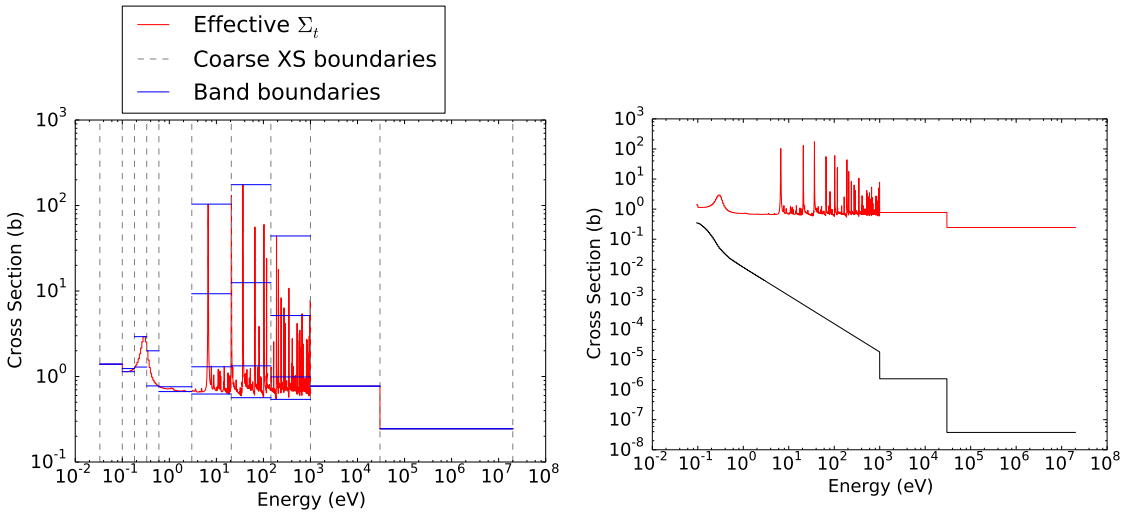

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure 7.76: Effective total cross sections for the PG-FEMG band boundary calculation for problem 5 with low-fidelity weight.

7.5.2.2 Resonance-only energy resolution study

An energy resolution study was done to determine how various QOIs changed as a function of increasing energy fidelity (DOF) for moderate DOF. The energy group structure was always fully resolved in the non-resonance regions and changed only in the resolved-resonance region. All errors are relative comparisons between the reference solution and the MG or PG-FEMG solutions.

Figure 7.77 gives the results from the energy resolution study. The number of groups and bands per group in the resolved resonance region were varied to determine their effect on QOI error. This resolution study resolved (to the reference level) all non-resonance energies and used varying yet moderate numbers of DOF within the resolved resonance region.

Figure 7.77a-c give work measurements. Figure 7.77a gives efficiency of the k-eigenvalue calculation (1/(work $\times$ error)), Figure 7.77b gives work per DOF, and Figure 7.77c shows total work scaled with the number of groups. Figure 7.77b shows PG-FEMG required approximately the same amount of work per DOF as MG. Both MG and PG-FEMG required substantially less work per DOF than the reference solution (horizontal line at 1.0).

Figure 7.77d-f show how error in the QOI was affected by changing energy resolution. As this resolution study did not use group sizes that resolved individual resonances, errors stayed approximately constant as the number of groups were increased. Errors did decrease as the number of bands per group were increased. For the k-eigenvalue calculation, going from 1 band per group (MG) to 2 bands per group decreased the error by a factor of approximately 7. Increasing from 2 bands to 6 bands decreased the same error by another factor of 7-17, depending on group structure. Similar factors held true for the absorption rates. Increasing from 1 to 2 bands

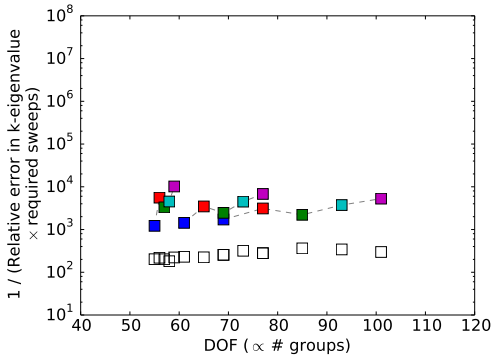

(a) k-Eigenvalue efficiency

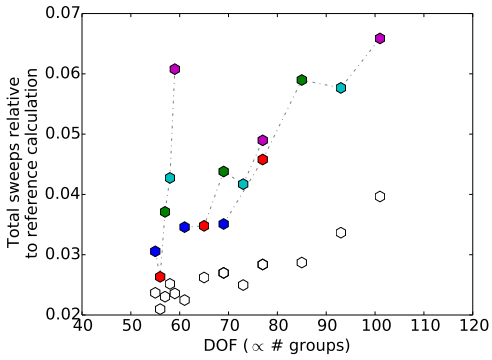

(b) Relative transport sweeps summed over bands / groups (work measurement)

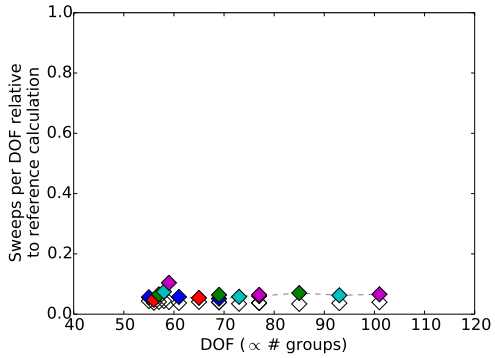

(c) Relative transport sweeps per DOF

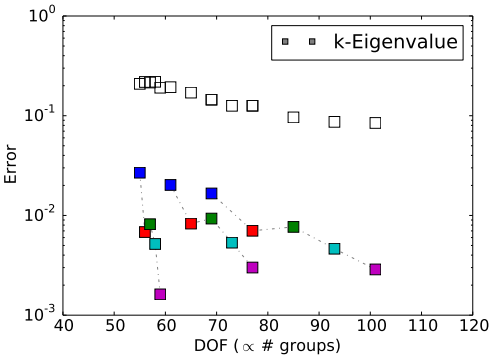

(d) k-Eigenvalue error

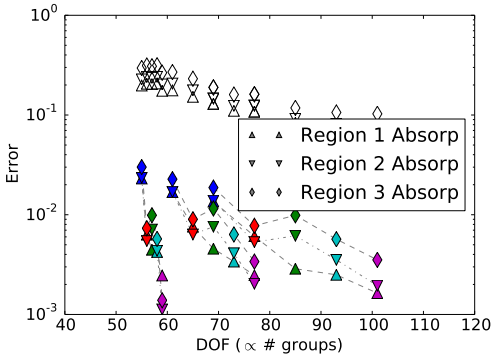

(e) Region absorption rate errors

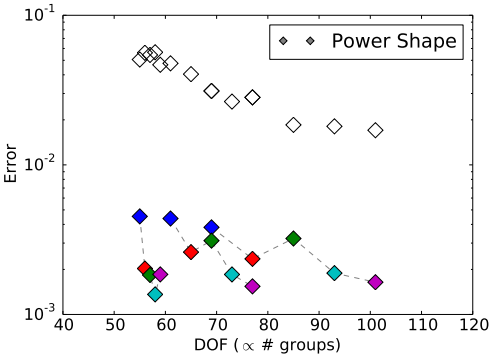

(f) Power shape

Figure 7.77: Error in various QOI relative to the reference solution for several energy group structures and numbers of bands for problem 5 with low-fidelity weight, for the resonance-only energy resolution study. The MG solution corresponds to unfilled markers; the PG-FEMG solutions are colored according to how many bands per group were used in the resonance region, be that two, three, four, five, or six. Dotted lines connect multiband cases with the same coarse group structure.

per group decreased the error in the power shape by a factor of approximately 6-12, depending on group structure. Increasing from 2 to 6 bands per group decreased the same error by an additional factor of 2-3 depending on group structure. This problem shows error decreasing mostly monotonically as the number of bands per group were increased, except when 4 bands per group were used, which might have produced poor band boundaries.

7.5.2.3 Energy resolution study

An energy resolution study was done to determine how various QOIs changed as a function of increasing energy fidelity (DOF) for large DOF. The energy group structure was resolved in all energy regions for this study. All errors are relative comparisons between the reference solution and the MG or PG-FEMG solutions.

Figure 7.78 gives the results from the energy resolution study. The number of groups and bands per group in the resolved resonance region were varied to determine their effect on QOI error. This resolution study focused on extremely large DOF to investigate asymptotic effects.

Figure 7.78a-c give work measurements. Figure 7.78a gives efficiency of the k-eigenvalue calculation (1/(work $\times$ error)), Figure 7.78b gives work per DOF, and Figure 7.78c shows total work scaled with the number of groups. Figure 7.78b shows PG-FEMG required more work per DOF than MG or the reference solution (horizontal line) in general, but this ratio was approximately constant over all numbers of DOF considered.

Figure 7.78d-f show how error in the QOI was affected by changing energy resolution. They give a sense of the extreme number of groups required to enter the asymptotic regime. Increasing from 1 band per group (MG) to 2 bands per group decreased the k-eigenvalue error by a factor of approximately 6-7. Increasing from

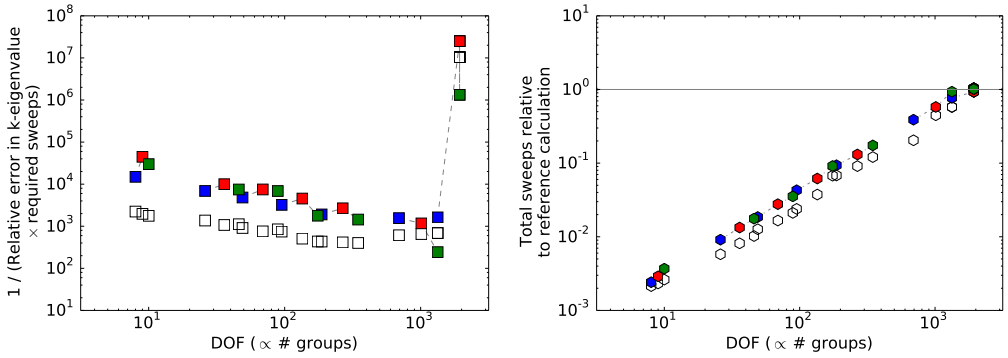

(a) k-Eigenvalue efficiency

(b) Relative transport sweeps summed over bands / groups (work measurement)

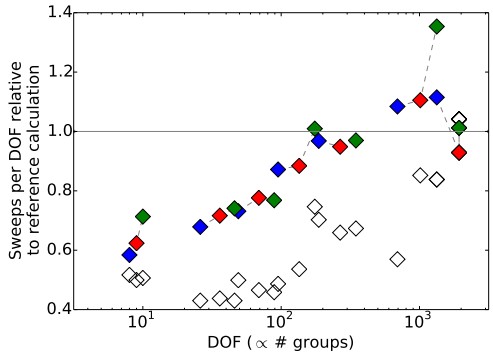

(c) Relative transport sweeps per DOF

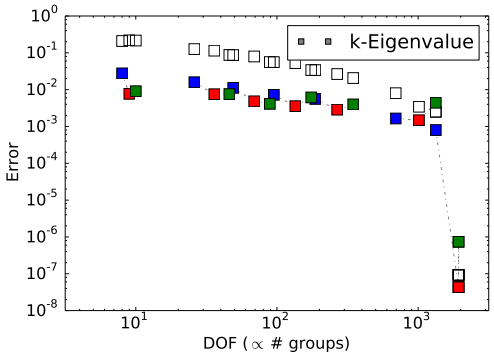

(d) k-Eigenvalue error

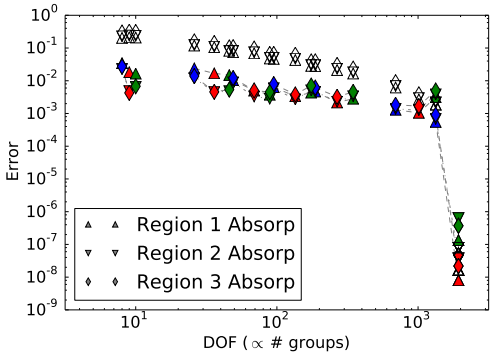

(e) Region absorption rate errors

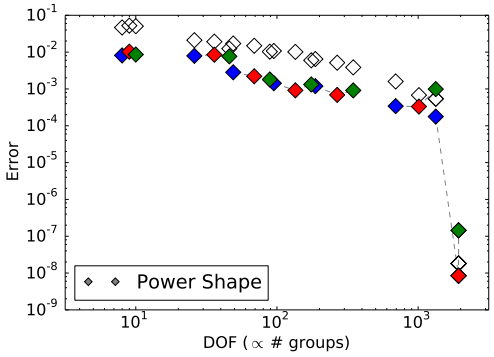

(f) Power shape

Figure 7.78: Error in various QOI relative to the reference solution for several energy group structures and numbers of bands for problem 5 with low-fidelity weight, for the full energy resolution study. The MG solution corresponds to unfilled markers; the PG-FEMG solutions are colored according to how many bands per group were used in the resonance region, be that two, three, four, five, or six. Dotted lines connect multiband cases with the same coarse group structure.

2 to 4 bands per group decreased this error by a factor of 1.5-3. Absorption rate errors had similar factors. Power profile shape errors decreased by a factor of 6-10 when increasing from 1 to 2 bands per group. There was negligible decrease in error when moving to 4 bands per group from 2 bands per group. There was slight non-monotonicity in k-eigenvalue and absorption rate errors for group counts around 100-300. The increasing error with increasing bands per group behavior at very large DOF count ($\sim$ 1000 DOF) for PG-FEMG was caused by the fall-back band boundary calculation which produced poor band boundaries.

8. CONCLUSIONS

8.1 Notable Properties of the PG-FEMG Method

In this section we extract generalizations from the data in the Results section. We note seven points about the PG-FEMG method, some desirable, some undesirable. Our new method

- has better accuracy per DOF than standard MG (factor of 3-10+),
- is robust to condensing spectrum and effective total cross section choice,
- captures space-energy-angle interactions better than standard MG,
- yields better flux reconstructions in energy than standard MG,
- introduces upper diagonal terms in the scattering matrix (looks like within-group upscattering),
- requires more work per DOF than standard MG (factor of <2),
- stresses spatial discretizations by accurately preserving boundary layers.

We illustrate each point more fully below.

8.1.1 Advantageous properties

Our PG-FEMG method has several advantageous properties when compared to standard MG. For the same DOF number, it has lower error in the chosen QOI than MG. When a high-fidelity, reference weighting is used for both methods, this advantage is around a factor of 3 or greater; when a low-fidelity, generic weighting is used, this advantage is around a factor of 10 or greater.

Our method is not sensitive to errors in condensing spectrum. Even when a weighting spectrum far from the local reference solution was used, the PG-FEMG method produced acceptable errors with modest DOF numbers. This is not true for MG. PG-FEMG relative errors of 10^{-4} to 10^{-2} are typical for this case, depending on the number of coarse groups and bands used. For the same cases, MG has relative errors of 10^{-1} or greater.

Our method is robust to band boundary choice in the following sense. Band boundaries in this work were chosen such that they correspond to a global effective total cross section. For a single resonant nuclide, this cross section is simply the total cross section of that nuclide. If the problem contains multiple resonant nuclides or a resonant nuclide at multiple temperatures, the resultant disparate total cross sections must be averaged together in some way. This yields band boundaries that are nowhere ideal and may not accurately represent resonance interference effects. Our test problems, especially problems 3 – 5, showed that our simplistic weighting scheme (volume- and density-weighted averaging) worked well and produced band boundaries that yielded accurate solutions. While this issue merits future study, the results shown here suggest that the method is not overly sensitive to the choice of a global $\Sigma_t(E)$, which is used only to define $\Delta E_{g,b}$ intervals.

Our method captures space-energy-angle interactions better than standard MG. Spatial self-shielding is a space-energy-angle phenomenon: neutrons streaming into the fuel from the moderator are absorbed in or scatter out of the fuel resonances, causing a position-energy-and-angle-dependent flux depression. MG has a fixed spectral shape within a group for all angles and positions that is unable to adapt without using a spatially varying weighting spectrum. PG-FEMG can adapt the spectral shape continuously in space and angle because we have broken up the solution into multiple pieces representing energies in different resonances. This allows PG-FEMG

to naturally self-shield, even when low-fidelity weighting functions are used. MG does not get this right.

Our method yields better flux reconstructions than standard MG. When using finite elements, it is important to note the DOFs are not the solution; the solution is a sum of the unknowns multiplied by shape functions.[1] We call this latter process flux reconstruction. When a high-fidelity reference solution is used, the basis (shape) functions have the resonance information in them and the flux reconstruction has resonance information in it. That is why both MG and PG-FEMG look accurate for reference weighting. When a low-fidelity solution is used, the shape functions have inaccurate resonances. Unless group sizes are used that capture individual resonances, the expanded MG spectrum will be inaccurate; we indeed found this to be the case. For the PG-FEMG case, the DOF themselves adapt to the resonances. Even if a smooth shape function is used, the flux reconstruction has the effects of resonances. Our results showed the PG-FEMG method to reproduce the fluxes accurately in the resonance region even when using low-fidelity weighting schemes.

We note that there is a deep connection between accurate spectra and spatial self-shielding. Self-shielding occurs when the spectrum changes locally due to upstream effects. Having the freedom to locally adapt the energy spectrum in space and angle near resonances gives PG-FEMG the ability to match the true flux shape, which is self-shielded. The MG flux shape for all space and angle is its weighting shape; this rigidity means MG requires an accurate (spatially-varying) weighting spectrum to do self-shielding correctly. We expect PG-FEMG to further outperform MG in multi-dimensional problems where the spectrum carries a stronger angular dependence.

[1] Our DOF do however correspond to a physical quantity: the flux integrated over the discontinuous energy domain of a band.

8.1.2 Disadvantageous properties

The primary disadvantage of the PG-FEMG method is that it is more difficult to set up and solve than standard MG. The difficulty in the setup includes forming the effective total cross section and rewriting portions of the cross-section preparation routines. The difficulty in the solve comes from using a band structure within a coarse group that introduces effective upscattering within the group. True down-scattering in energy space from a resonant band to a non-resonant band may look like upscatter in band-space depending on DOF ordering. This effective upscattering is an unavoidable component of this method.

The presence of the effective upscattering yields systems of equations that are more difficult to solve. This requires more sweeps per group to converge the solution. In practice, we found this increase to be less than a factor of 2 when using 6 or fewer bands per group in the resonance region. Considering that commensurate decreases in error were at least a factor of 3 and usually greater than a factor of 10, we consider this acceptable. Future iterative methods that use block-Jacobi in energy for enhanced parallelism will benefit from the upscattering introduced in PG-FEMG, as one Gauß-Seidel sweep over energy domains will involve several energy-local block-diagonal iterations that can each converge upscattering.

In computational transport, increasing fidelity in one variable often requires commensurate increases in the fidelity of the others. Such is the case with PG-FEMG. The true solution contains spatial boundary layers near interfaces between spatial regions with disparate resonance structures at the energies of those resonances (e.g., the fuel-moderator boundary). In the PG-FEMG method, DOFs represent the resonances and have large cross sections in the resonant materials; in the MG method, the resonances are smeared out over a group, leading to smaller cross sections. The

PG-FEMG method therefore preserves boundary layers while the MG method does not. Capturing their true behavior requires finer spatial discretizations at these boundary layers.

8.1.3 General remarks

We have devised a new method, PG-FEMG, for high-fidelity treatment of the energy variable in neutron and photon transport. PG-FEMG is a finite-element method that can be viewed as a generalization of the standard MG and MB methods.

The new method differs in key aspects from previous MB methods. A major differentiator is the global band structure, which allows PG-FEMG cross sections to be used by any standard MG code. Another is that the sole approximation made in PG-FEMG is an expansion of the angular flux in basis functions. The band-wise basis functions of PG-FEMG generate cross sections averaged over discontinuous energy domains within a coarse group, which differs from standard MB treatment and is reminiscent of the opacity density functional (ODF) method from radiative transfer. However, instead of using a single coarse group and many bands as is often done in radiative applications with ODF,[2] PG-FEMG uses multiple coarse groups in the resonance region and often has multiple bands per group only inside (resolved) resonance regions.

We have shown that PG-FEMG results in reductions in errors for several QOIs for five pin-cell k-eigenvalue problems of increasing complexity. We investigated errors between a reference solution and either a PG-FEMG solution or a MG solution for criticality eigenvalue, power profile, and absorption rates. In all problems and for all QOIs, PG-FEMG produced a substantially lower error for a given total number of unknowns than did MG.

[2]Radiative transfer oftentimes assumes coherent scattering and fission-like absorption-reemission, making this one-group practice more reasonable.

8.2 Summary of Thesis

In this thesis, we have presented a novel method for the discretization of the energy variable in the context of the neutron transport equation applied to k-eigenvalue pin-cell problems. The de-facto standard for energy discretization, multigroup (MG), requires flux weighting from an approximation to the energy shape of the solution. That is, the solution itself is needed to solve for the solution; in reality, things are not this dire, as MG is accurate if given an accurate simulacrum of the solution. MG becomes inaccurate when the flux weight shapes diverge from the true solution shape over the breadth of a group. Standard MG suffers from having a fixed spectral shape within a group, which cannot capture space-energy-angle effects such as the self-shielding that occurs in the resonances of the fuel as particles stream inward.

We have presented a generalized-MG method, which we named the Petrov-Galerkin Finite-Element Multigroup (PG-FEMG) method, to remedy these deficiencies. PG-FEMG, a cousin of the multiband (MB) method, provides multiple degrees of freedom (DOF) within a group. These DOF are concentrated at resonances, where the total cross sections are high. PG-FEMG thus provides fidelity where it is most needed and where it was previously lacking: in the treatment of resonances. Our method utilizes a global band structure, which is its primary differentiator from MB, which historically has used region-dependent band structures. Our global system provides advantages — being able to utilize the method in existing transport solvers — and disadvantages — requiring the construction of, in our current formulation, an accurate, problem-wide effective total cross section on which to build the bands. We believe the advantages in simplicity and usability outweigh the disadvantages in potential non-optimality of the local band structure.

We have presented theory to describe our new method mathematically and have

given practical guidelines on construction of algorithms to implement this theory. In the context of a finite element method, we use weight functions which are Heaviside functions — unity when energy is in the coarse group's energy range and the total effective cross section is in the cross section's band range, and zero elsewhere — and basis functions which are the product of the aforementioned Heaviside functions and normalized local spectral shapes. The basis functions are used to condenses the pointwise or ultra-fine-group cross sections into PG-FEMG cross sections. This framework is a generalization of standard MG, which uses only one band per group. A secondary differentiator between PG-FEMG and typical MB is the definition of the FE space: MB uses as its DOF moments of the cross section weighted with various analytic fluxes; we use integrals over discontinuous portions of the energy domain, reminiscent of Lebesgue integration.

We wrote an S_n code, SCDT, to solve the transport system for k-eigenvalue problems. As this work provides a proof of concept, we restricted ourselves to one-dimensional cylinders so we could study pin cells. We wrote SCDT to incorporate modern discretization and algorithmic practices, including linear-discontinous FE spatial discretization, weighted-diamond difference finite difference angular discretization, starting-direction calculations, S2SA acceleration of the within-group scattering source, augmentation of the state vector to include DOF representing the reflective boundary condition, and non-linear Krylov acceleration (NKA) of the nonlinear eigenvalue system. SCDT and our cross section condensation code were our workhorses, allowing us to test the generalized-multigroup method on several problems.

We thought up a series of five problems of increasing difficulty on which to test the method. The simplest problem was a homogenized pin cell of MOX[3] and hydrogen

[3]We initially used MOX — a low-fissile blend of Pu-239 and U-238 — because of the interesting

with a white boundary.[4] When the high-fidelity spatial-region-averaged reference solution was used as the basis functions (for weighting), this provided a verification of the code, as both the MG and PG-FEMG were expected to get the reference eigenvector and eigenvalue for this problem. Our results showed that the solutions were the same to within solver tolerance. When a low-fidelity weighting scheme was used, the PG-FEMG solution significantly outperformed the MG solution in terms of accuracy of our quantities of interest (QOI) with respect to the reference solution. Our QOI were criticality eigenvalue, reactor power shape (approximated by fission rate), and absorption rates (a proxy for nuclide-specific depletion rates).

The second problem involved an inner cylinder of MOX fuel surrounded by light water in a geometric configuration representative of a typical PWR. When spatial-region-averaged reference fluxes were used as the cross section weighting parameters, we did not expect to see exactness for the MG or PG-FEMG cases due to spatial and angular dependences of the flux smeared out in the region-weighting of the reference solution. Our results bore this out. For this problem, we found PG-FEMG to be more resilient to changes in weighting spectrum than standard MG.

The third problem once again consisted of MOX surrounded by light water; a temperature profile was added. Three representative temperatures were used in the fuel and one in the moderator. Because increasing temperatures broaden the resonances, we expected a degradation in the fidelity of the bands due to the effect of multiple resonance widths on the total effective cross section. We found that the PG-FEMG method still performed well, and much better than the MG method. Once again, we found the PG-FEMG to be less sensitive to changes in weighting spectrum compared to standard MG.

low resonance of Pu-239 and because this nuclide took substantially less time to process in NJOY than U-235. We continued to use it out of habit.

[4] All problems used white boundaries.

Problems four and five involved single-temperature MOX fuel surrounded by LEU fuel and light water. Problem four used the order (with increasing radius): MOX, LEU, water; problem five used the order: MOX, water, LEU, in an attempt to approximate a fuel lattice structure with a one-dimensional cylinder. Both problems involved separated fissile materials with disparate resonances, though in all fuel materials the dominant resonant nuclide was U-238. The results were similar to the ones found previously, with PG-FEMG outperforming MG and generally being less sensitive to weighting spectrum. We did not notice any substantial degradation in band structure when multiple resonant nuclides were added, which gives us hope for problems with many resonant nuclides.

Our results paint a picture in which, for equal DOF in the energy variable, PG-FEMG outperforms MG with both high- and low-fidelity weighting spectra. For the problems tested, PG-FEMG is less sensitive to this weighting spectrum than standard MG is. PG-FEMG outperforms MG both with respect to our QOIs and qualitatively as well. The PG-FEMG framework was used to reconstruct the fine-scale energy shape of the PG-FEMG and MG solutions. When these reconstructions were compared to the reference solution in cases where the reference solution was not used to weight the coarse cross sections, PG-FEMG vastly outperformed MG, capturing resonances and their flux dips. We therefore suggest further investigation into PG-FEMG for flux reconstruction applications.

The one area where MG does outperform PG-FEMG is in required work, which we measure by sweeps[5] required to converge the eigenvalue/eigenvector pair. This is expected, as the PG-FEMG method introduces effective upscatter within a coarse group between bands. We note that this increase was substantially less than the commensurate increase in fidelity, implying that PG-FEMG is more efficient than

[5]Recall, one sweep is required per energy DOF.

MG.

We realize the problems we report on were basic. They used only a handful of resonant nuclides, they were in one dimension, and they were steady-state. Such restrictions were necessary for the initial characterization of the method which is presented here. In future work we plan to apply this method to a realistic reactor geometry and a many-nuclide depletion problem.

In conclusion, we have introduced, characterized, and tested a novel energy discretization scheme and found it to be superior to the currently ubiquitous MG method. Our generalized-MG method is straightforward to implement and has beat MG even when the MG method is given an order of magnitude more DOF in the energy variable. We are excited to see what benefit this framework may offer to the nuclear engineering community as we seek to better model energy dependencies for transport applications.

8.3 Future Work

Future work will include applying PG-FEMG to multi-dimensional problems, comparing PG-FEMG to Monte Carlo calculations using the same cross section data sources, and investigating more advanced band boundary construction algorithms. It will also include investigation of convergence under various refinement scenarios (fixed groups with increasing bands/group, fixed bands/group with increasing groups, and combinations).

REFERENCES

[1] J. J. Duderstadt, L. J. Hamilton, Nuclear Reactor Analysis, John Wiley & Sons, 1976.

[2] E. E. Lewis, W. F. Miller, Computational Methods of Neutron Transport, American Nuclear Society, La Grange Park, Illinois, 1993.

[3] S. Chandrasekhar, The Radiative Equilibrium of the Outer Layers of a Star, with Special Reference to the Blanketing Effect of the Reversing Layer, Mon. Not. R. Astr. Soc. 96 (1935) 21–41.

[4] S. E. Strom, R. L. Kurucz, A Statistical Procedure for Computing Line-Blanketed Model Stellar Atmospheres with Applications to the F5 IV Star Procyon, J. Quant. Spectrosc. Radiat. Transfer 6 (1966) 591–607.

[5] D. F. Carbon, A Comparison of the Straight-Mean, Harmonic-Mean, and Multiple-Picket Approximations for the Line Opacities in Cool Model Atmospheres, The Astrophysical Journal 187 (1974) 135–145.

[6] D. M. Mihalas, B. Mihalas, Foundations of Radiation Hydrodynamics, Dover Publications, Inc., Mineola, New York, 1999.

[7] L. H. Auer, R. B. Lowrie, Dispersion Analysis of Radiation/Thermal Fronts with Fully Resolved Spectral Opacity Variation, Tech. Rep. LA-UR-00-6094, Los Alamos National Laboratory, Los Alamos, New Mexico (Oct. 2000).

[8] J.-F. Ripoll, A. A. Wray, A 3D Multiband Closure for Radiation and Neutron Transfer Moment Models, Journal of Computational Physics 227 (4) (2008) 2212 – 2237.

[9] A. A. Wray, J.-F. Ripoll, D. Prabhu, Investigation of the Opacity Binning Approach for Solving the Shock-Generated Radiation of the Apollo AS-501 Re-

entry, Tech. rep., Center for Turbulence Research (2006).

[10] D. E. Cullen, G. C. Pomraning, The Multiband Method in Radiative Transfer Calculations, Journal of Quantitative Spectroscopy and Radiative Transfer 24 (2) (1980) 97 – 117.

[11] M. N. Nikolaev, V. V. Filippov, Measurement of the Resonance Parameters for Total Cross Sections of Some Elements in the Energy Region 0.3-2.7 MeV, Atom. Energ. 6 (6) (1963) 493, (In Russian).

[12] V. V. Sinitsa, M. N. Nikolaev, Analytical Determination of Subgroup Parameters, Sov. At. Energy 35 (6) (1973) 1129–1131, translated from At. Energ. (USSR).

[13] J. C. Stewart, Non-Grey Radiative Transfer, J. Quant. Spectrosc. Radiat. Transfer 4 (1964) 723–729.

[14] D. E. Cullen, Application of the Probability Table Method to Multigroup Calculations of Neutron Transport, Nucl. Sci. Eng. 55 (4) (1974) 387–400.

[15] D. E. Cullen, E. F. Plechaty, R. J. Doyas, C. R. Weisbin, J. E. White, Cross Section Probability Tables in Multigroup Transport Calculations, Tech. rep., Lawrence Livermore National Laboratory, Livermore, California (Mar. 1978).

[16] T. Takeda, H. Fujimoto, K. Sengoku, S. Shiroya, H. Unesaki, K. Kanda, Application of Multiband Method to KUCA Tight-Pitch Lattice Analysis, Journal of Nuclear Science and Technology 28 (9) (1991) 863–869.

[17] T. Takeda, Y. Kanayama, A Multiband Method with Resonance Interference Effect, Nucl. Sci. Eng. 131 (3) (1999) 401–410.

[18] T. Yamamoto, T. Takeda, Reaction Rate Calculation in Fast Reactor Blanket Using Multiband S_n Theory, Journal of Nuclear Science and Technology 37 (5) (2000) 428–435.

[19] T. Yamamoto, Background-Cross-Section-Dependent Subgroup Parameters,

Journal of Nuclear Science and Technology 40 (6) (2003) 370–375.

[20] T. Yamamoto, T. Takeda, Subgroup Parameters based on Orthogonal Factorization, Journal of Nuclear Science and Technology 44 (1) (2007) 36–42.

[21] S.-E. Huang, K. Wang, D. Yao, An Advanced Approach to Calculation of Neutron Resonance Self-Shielding, Nuclear Engineering and Design 241 (8) (2011) 3051 – 3057.

[22] M. J. Milošević, A New Subgroup Method for Resolved Resonance Treatment, Nucl. Sci. Bull. 2 (1) (1997) 9–13.

[23] A. V. Shilkov, Generalized Multigroup Approximation and Lebesgue Averaging Method in Particle Transport Problems, Transport Theory and Statistical Physics 23 (6) (1994) 781–814.

[24] L. B. Levitt, Probability Table Method for Treating Unresolved Neutron Resonances in Monte Carlo Calculations, Nucl. Sci. Eng. 49 (4) (1972) 540–457.

[25] P. Ribon, J. M. Maillard, Probability Tables and Gauss Quadrature: Application to Neutron Cross-Sections in the Unresolved Energy Range, in: ANS Topical Meeting on Advances in reactor Physics and Safety, ANS, Saratoga Springs, New York, 1986.

[26] A. Hébert, M. Coste, Computing Moment-based Probability Tables for Self-shielding Calculations in Lattice Codes, Nucl. Sci. Eng. 142 (2002) 245–257.

[27] A. Hébert, Development of the Subgroup Projection Method for Resonance Self-Shielding Calculations, Nucl. Sci. Eng. 162 (2009) 56–75.

[28] N. Martin, J. Miss, A. Hébert, Moment-based Probability Tables for Angular Anisotropic Scattering, Annals of Nuclear Energy 38 (5) (2011) 1125 – 1132.

[29] N. Martin, A. Hébert, A Monte Carlo Lattice Code with Probability Tables and Optimized Energy Meshes, Nucl. Sci. Eng. 167 (2011) 177–195.

[30] X-5 Monte Carlo Team, MCNP - A General Monte Carlo N-Particle Trans-

port Code, Version 5 Volume I: Overview and Theory, Los Alamos National Laboratory, lA-UR-03-1987 (February 2008).

[31] T. M. Sutton, F. B. Brown, Implementation of the Probability Table Method in a Continuous-Energy Monte Carlo Code System, Tech. Rep. Conf-981003, KAPL, Schenectady, New York, 12301 (Oct. 1998).

[32] T. M. Sutton, T. J. Donovan, P. S. Dobreff, E. Caro, D. P. Griesheimer, L. J. Tyburski, D. C. Carpenter, H. Joo, The MC21 Monte Carlo Transport Code, in: Joint International Topical Meeting on Mathematics & Computation and Supercomputing in Nuclear Applications, American Nuclear Society, Monterey, California, 2007.

[33] Y. Liu, B. Collins, B. Kochunas, W. Martin, K.-S. Kim, M. Williams, Resonance Self-Shielding Methodology in MPACT, in: International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering, American Nuclear Society, 2013.

[34] Y. Liu, B. Collins, B. Kochunas, W. Martin, K.-S. Kim, M. Williams, Modeling Resonance Interference by 0-D Slowing-Down Solution with Embedded Self-Shielding Method, in: International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering, American Nuclear Society, 2013.

[35] N. A. Gibson, B. Forget, On the Stability of the Discrete Generalized Multigroup Method, Annals of Nuclear Energy 65 (2014) 421–432.

[36] S. Douglass, F. Rahnema, Subgroup Decomposition Method, Annals of Nuclear Energy 48 (0) (2012) 84 – 101.

[37] I. Attieh, R. Pevey, An Adaptive General Multigroup Method, in: Proc. 27th Annual CNS-CAN Student Conference, Toronto, Ontario, Canada, 2002.

[38] I. K. Attieh, Generalized Multigroup Method, Ph.D. thesis, The University of

Tennessee (2004).

[39] B. Forget, F. Rahnema, A Spectral Unfolding Method, Trans. Am. Nucl. Soc. 96 (2007) 669.

[40] F. Rahnema, S. Douglass, B. Forget, Generalized Energy Condensation Theory, Nucl. Sci. Eng. 160 (2008) 41–58.

[41] L. Zhu, B. Forget, A Discrete Generalized Multigroup Energy Expansion Theory, Nucl. Sci. Eng. 166 (2010) 239.

[42] L. Zhu, B. Forget, An Energy Recondensation Method Using the Discrete Generalized Multigroup Energy Expansion Theory, Annals of Nuclear Energy 38 (8) (2011) 1718 – 1727.

[43] C. N. McGraw, M. L. Adams, W. D. Hawkins, M. P. Adams, T. Smith, Accuracy of the Linear Discontinuous Galerkin Method for Reactor Analysis with Resolved Fuel Pins, in: Proc. International Conference on the Physics of Reactors, American Nuclear Society, Kyoto, Japan, 2014, submitted.

[44] D. E. Cullen, R. N. Blomquist, M. G. Greene, E. Lent, R. MacFarlane, S. McKinley, E. F. Plechaty, J. C. Sublet, How Accurately Can We Calculate Neutrons Slowing Down in Water?, Tech. rep., Lawrence Livermore National Laboratory, Livermore, California (Apr. 2006).

[45] R. E. MacFarlane, R. M. Roicourt, NJOY: A Neutron and Photon Processing System, Trans. Am. Nucl. Soc. 22 (1975) 720.

[46] R. E. MacFarlane, D. W. Muir, The NJOY Nuclear Data Processing System, Version 91, Los Alamos National Laboratory, Los Alamos, New Mexico, report LA-12750-M (1994).

[47] J. S. Warsa, T. A. Wareing, J. E. Morel, Krylov Iterative Methods and the Degraded Effectiveness of Diffusion Synthetic Acceleration for Multidimensional S_n Calculations in Problems with Material Discontinuities, Nucl. Sci. Eng. 147 (3)

(2004) 218–248.

[48] N. N. Carlson, K. Miller, Design and Application of a Gradient-Weighted Moving Finite Element Code I: in One Dimension, SIAM J. Sci. Comput. 19 (3) (1998) 728–765.

[49] M. T. Calef, E. D. Fichtl, J. S. Warsa, M. Berndt, N. N. Carlson, Nonlinear Krylov Acceleration Applied to a Discrete Ordinates Formulation of the k-eigenvalue Problem, Journal of Computational Physics 238 (2013) 188–209.

[50] N. N. Carlson, K. Miller, Design and Application of a Gradient-Weighted Moving Finite Element Code II: in Two Dimensions, SIAM J. Sci. Comput. 19 (3) (1998) 766–798.

[51] D. G. Anderson, Iterative Procedures for Nonlinear Integral Equations, J. ACM 12 (4) (1965) 547–560.

[52] D. A. Knoll, H. Park, C. Newman, Acceleration of k-Eigenvalue/Criticality Calculations Using the Jacobian-Free Newton-Krylov Method, Nucl. Sci. Eng. 167 (2) (2011) 133–140.

[53] Y. Saad, M. H. Schultz, GMRES: A Generalized Minimal Residual Algorithm for Solving Non-Symmetric Linear Systems, SIAM J. Sci. Stat. Comput. 7 (1986) 856.

[54] D. F. Gill, Y. Y. Azmy, Newton's Method for solving k-Eigenvalue/ Problems in Neutron Diffusion Theory, Nucl. Sci. Eng. 167 (2) (2011) 141–153.

[55] D. F. Gill, Y. Y. Azmy, J. S. Warsa, J. D. Densmore, Newton's Method for the Computation of k-Eigenvalues in S_n Transport Applications, Nucl. Sci. Eng. 168 (1) (2011) 37–58.

[56] M. L. Adams, N. M. Amato, P. Nelson, L. Rauchwerger, Efficient Massively-Parallel Implementations of Modern Deterministic Transport Calculations, Tech. rep., Texas A&M University, report to the Department of Energy (2002).

APPENDIX A

SPACE AND ANGLE DISCRETIZATIONS OF THE TRANSPORT EQUATION

Many standard references, e.g., [1, 2], are of sufficient age not to cover modern discretization practices, esp. the use of weighted-diamond finite-difference angular derivatives, starting-direction sweeps, Galerkin quadrature, and linear-discontinuous finite element spatial discretizations. While discussion and derivation of such topics is profitable, they remain beyond the scope of the thesis proper due to their length and because they do not represent novel work. For completeness of discussion, they are included here.

A.1 Restriction to Cylindrically-Symmetric Geometries (cont.)

For reference, we copy Eq. (2.1) below:

$$\frac{1}{v(E)}\frac{\partial}{\partial t}\psi(\mathbf{r},E,\boldsymbol{\Omega},t)+\boldsymbol{\Omega}\cdot\boldsymbol{\nabla}\psi(\mathbf{r},E,\boldsymbol{\Omega},t)+\Sigma_t(\mathbf{r},E,t)\psi(\mathbf{r},E,\boldsymbol{\Omega},t)$$

$$=S\psi(\mathbf{r},E,\boldsymbol{\Omega},t)+F\psi(\mathbf{r},E,\boldsymbol{\Omega},t)+q_{\text{ext}}, \quad \text{(A.1a)}$$

$$S\psi(\mathbf{r},E,\boldsymbol{\Omega},t)=\int_0^\infty \mathrm{d}E' \int_{4\pi} \mathrm{d}\boldsymbol{\Omega}' \;\; \Sigma_s(\mathbf{r},E'\to E,\boldsymbol{\Omega}\cdot\boldsymbol{\Omega}',t)\;\psi(\mathbf{r},E',\boldsymbol{\Omega}',t), \quad \text{(A.1b)}$$

$$F\psi(\mathbf{r},E,\boldsymbol{\Omega},t)=\frac{\chi(\mathbf{r},E,t)}{4\pi}\int_0^\infty \mathrm{d}E' \;\; \nu\Sigma_f(\mathbf{r},E',t)\int_{4\pi} \mathrm{d}\boldsymbol{\Omega}' \;\; \psi(\mathbf{r},E',\boldsymbol{\Omega}',t), \quad \text{(A.1c)}$$

$$\psi(\mathbf{r},E,\boldsymbol{\Omega},t)=v(E)\,n(\mathbf{r},E,\boldsymbol{\Omega},t), \quad \text{(A.1d)}$$

$$\psi(\mathbf{r},E,\boldsymbol{\Omega},t)=f_{\text{in}}(\mathbf{r},E,\boldsymbol{\Omega},t) \qquad \boldsymbol{\Omega}\cdot\hat{\mathbf{n}}<0, \quad \text{(A.1e)}$$

$$\psi(\mathbf{r},E,\boldsymbol{\Omega},0)=f_0(\mathbf{r},E,\boldsymbol{\Omega}). \quad \text{(A.1f)}$$

In cylindrical geometries, the streaming operator is given by

$$\boldsymbol{\Omega} \cdot \boldsymbol{\nabla} \psi = \frac{\mu}{r} \frac{\partial}{\partial r} \big(r \, \psi(r, \boldsymbol{\Omega}, E) \big) - \frac{1}{r} \frac{\partial}{\partial \omega} \big(\eta \, \psi(r, \boldsymbol{\Omega}, E) \big), \tag{A.2a}$$

$$\boldsymbol{\Omega} \cdot \boldsymbol{\nabla} \psi = \mu \frac{\partial \psi}{\partial r}(r, \boldsymbol{\Omega}, E) - \frac{\eta}{r} \frac{\partial \psi}{\partial \omega}(r, \boldsymbol{\Omega}, E), \tag{A.2b}$$

in conservative and non-conservative forms, respectively.

The angular variable is defined by axis-cosines:

$$\mu = \hat{\mathbf{x}} \cdot \boldsymbol{\Omega}, \tag{A.3a}$$

$$\eta = \hat{\mathbf{y}} \cdot \boldsymbol{\Omega}, \tag{A.3b}$$

$$\xi = \hat{\mathbf{z}} \cdot \boldsymbol{\Omega}, \tag{A.3c}$$

or equivalently,

$$\mu = \sqrt{(1 - \xi^2)} \cos(\omega), \tag{A.3d}$$

$$\eta = \sqrt{(1 - \xi^2)} \sin(\omega), \tag{A.3e}$$

where ω is the azimuthal angle and ξ is the polar cosine. By cylindrical symmetries, it is known that ξ and η are symmetric about 0, but μ is not. These symmetries are exploited and the transport equation, Eq. (A.1), is solved in the two unique octants defined by $\xi > 0$, $\eta > 0$.

The range of our angular cosines are

$$\mu \in [-1, 1], \tag{A.4a}$$

$$\eta \in [0, 1], \tag{A.4b}$$

$$\xi \in [0, 1], \tag{A.4c}$$

or equivalently,

$$\omega \in [0, \pi], \tag{A.4d}$$

$$\theta \in [0, \pi/2], \tag{A.4e}$$

$$(\xi = \cos\theta). \tag{A.4f}$$

This comprises two of the eight octants.

A.2 Discretization in Angle (cont.)

We take the standard S_n approach of collocating the transport equation at discrete angular directions. This makes the transport operator, $\nabla \cdot \Omega$, diagonal but the scattering operator in general full.

We use the standard Legendre-Chebyshev quadrature with levels of equal ξ which have a variable number of azimuths, ω. Let k_l be the index of the quadrature points for level l, such that $l \in \{1, \ldots, n/2\}$, $k_l \in \{1, \ldots, 2l\}$. The cardinality of the set $\{k_l\}\ \forall\, l, k$ is $(n/2)(n/2+1)$. We call the collection of these points and their associated weights the S_n quadrature. We assign weights to the quadrature to preserve scattering kernels, which requires preserving angular moments.

This quadrature is called Legendre-Chebyshev because polar (level) spacing is given by the Gauß-Legendre quadrature and azimuthal spacing is given by Gauß-

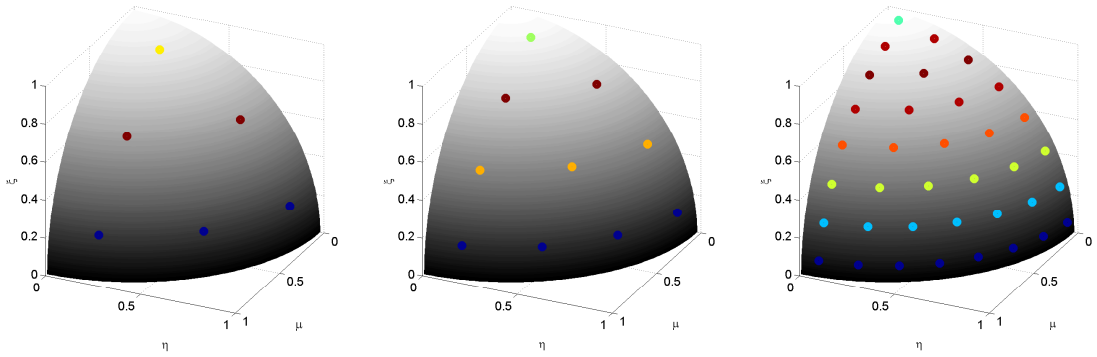

(a) S_6 (b) S_8 (c) S_{16}

Figure A.1: Examples of the triangular Legendre-Chebyshev quadratures used. The other octant used, with $\mu < 0$ and $\eta > 0$, is not shown. Colors correspond to the weights of each direction, with redder colors signifying larger weights.

Chebyshev quadrature, i.e., equal spacing. Figure A.1 provides an example of Legendre-Chebyshev quadratures. We note this quadrature is a hierarchical quadrature of Gauß-Legendre on the levels and Gauß-Chebyshev within each level. Specifically, the top polar level (with ξ closest to unity) has 2 azimuthal directions (one per octant), the next 4, etc., with the azimuthal points evenly spread out in the $\eta > 0$, $\xi > 0$ octants.

The points' weights were the Legendre weights for the level divided by the number of directions for that level, such that the sum over all the weights was unity. This gave rise to the so-called uniradians unit. A useful consequence of this normalization is that, for isotropic fluxes, the magnitude of the angular and scalar fluxes would be the same, not off by a factor of 4π.

Galerkin quadrature was used to enable movement between the discrete directions and the Legendre moments of the angular fluxes. Namely, a moment-to-discrete $\mathcal{M}$ and discrete-to-moment $\mathcal{D}$ matrix were used to map one to the other. It was desired

that these matrices be inverses of one another. To that end, the following definitions were used:

$$\mathcal{M}_{i,j} = (2l(i)+1)\, Y_{l(i)}^{m(i)}(\mathbf{\Omega}_j), \tag{A.5a}$$

$$\mathcal{D} = \mathcal{M}^{-1}, \tag{A.5b}$$

$$Y_l^m(\mathbf{\Omega}) = \sqrt{\left[(2-\delta_{m,0})\,\frac{(l-|m|)!}{(l+|m|)!}\right]}\, P_l^m(\xi)\; cos(m\,\omega), \tag{A.5c}$$

where ξ is the polar cosine, ω is the azimuthal angle, the $\mathbf{\Omega}_j$ are the discrete quadrature points, $\delta_{m,0}$ is the Kronecker delta, the P_m^l are the associated Legendre polynomials, and l and m are chosen such that $l+m$ is even and $0 \leq m \leq l$. Note that Eq. (A.5c) only holds for $m \geq 0$. As $\mathcal{M}^{-1}$ was subject to numerical noise, the first row of $\mathcal{D}$, which gave the scalar flux when multiplied by ψ, was replaced with the quadrature weights. (**Verify** this.)

The angular derivative term in Eq. (A.2a) was approximated using finite difference (FD):

$$\left.\frac{\partial}{\partial\omega}(\eta\psi)\right|_k \simeq \frac{\alpha_{k+1/2}\,\psi_{k+1/2} - \alpha_{k-1/2}\,\psi_{k-1/2}}{w_k}, \tag{A.6a}$$

$$\alpha_{k+1/2} = \alpha_{k-1/2} - \mu_k w_k, \tag{A.6b}$$

where $k = k_l$ is the collocation index and the w_k are the quadrature weights. The starting condition was $\alpha_{k_l+1/2} = 0$ per level. As k increases, μ_k increases from -1 to +1; for a given level, ξ is fixed. Using the skew-symmetry of the μ_k and the symmetry of the w_k, this implies that the $\alpha_{k+1/2}$ are symmetric and that $\alpha_{K_l+1/2} = 0$ for $K_l - k_l + 1$ quadrature points per level. We note that Eq. (A.6a) is a discrete version of $\int_0^{\omega_k} \mathrm{d}\hat{\omega}\; \mu(\hat{\omega})$, which is even (symmetric in k).

A weighted-diamond difference (WDD) approach was used to define ψ as

$$\psi_k = \alpha_k \psi_{k+1/2} + (1 - \alpha_k)\, \psi_{k-1/2}, \tag{A.7a}$$

$$\alpha_k = \frac{\cos\omega_k - \cos\omega_{k+1/2}}{\cos\omega_{k-1/2} - \cos\omega_{k+1/2}}, \tag{A.7b}$$

where the ω_k and $\omega_{k\pm1/2}$ are evenly spaced from the Gauß-Chebyshev quadrature. (Note full- and half-integer subscripts are different.) For a given level, l, as k_l increases, ξ stays the same, but η and μ change. One way to motivate the cosines in Eq. (A.7b) is by requiring discrete exactness for $Y_l^0(\Omega) \sim f(\xi)$ and $Y_l^1(\Omega) \sim f(\xi)\cos(\omega)$.

A.3 Discretization in Space (cont.)

Now-standard linear-discontinuous (LD) finite element (FE) methodologies were used for the spatial discretization of the transport equation, Eq. (A.1). The solution was cast as a linear combination of basis functions, in this case the standard

$$B_L(r) = \frac{r_1 - r}{r_1 - r_0}, \tag{A.8a}$$

$$B_R(r) = \frac{r - r_0}{r_1 - r_0}. \tag{A.8b}$$

The angularly-discretized system was multiplied by a test function and integrated over a cell using cylindrical ($\mathrm{d}r\ r$) weighting. As the FEM was Galerkin, the test / weight functions were equal to the basis functions. The integral itself was carried out exactly using Gauß-Legendre quadrature[1]. As the largest-order polynomials that needed to be integrated were in the mass terms and involved the product of two basis functions and the r-weighting (3^{rd}-order), S_2 quadrature was used, which is exact for

[1]Quadrature is used for integrals in both the angular and spatial domains. The quadratures used are in general different.

0^{th}- through 3^{rd}-order polynomials. The integrand and quadrature were mapped to the reference element, $[0,1]$, where the basis functions were simply

$$B_L(u) = 1 - u, \tag{A.9a}$$

$$B_R(u) = u. \tag{A.9b}$$

An integration by parts was used on the spatial derivative term, $\nabla \cdot \Omega\, \psi$. Aside from the boundary terms created by the integration by parts, all unknowns in the integrals relate to within-cell DOF. The correct physical definition of the boundary terms uses up-winded quantities, which couple the current cell with the one in the $-\mathbf{\Omega}$ direction. This allows a sweep in the direction of $\mathbf{\Omega}$ through the grid: start at a spatial boundary of the domain and traverse the cells in the $+\mathbf{\Omega}$ direction to the opposite boundary. Such traversal allows each cell to be solved in sequence with no lagging of neighbors' information.

The FE representation of ψ and ϕ in space is given by the inner product

$$\psi_k(r)\big|_{\text{cell s}} = \begin{pmatrix} B_{s,L}(r) & B_{s,R}(r) \end{pmatrix} \begin{pmatrix} \psi_{k,s,L} \\ \\ \psi_{k,s,R} \end{pmatrix}, \tag{A.10a}$$

$$\psi_k(r)\big|_{\text{cell s}} = \mathbf{B}_s(r)^{\text{T}} \boldsymbol{\psi}_{k,s}, \tag{A.10b}$$

for direction or moment k.

For simplicity, we present a separate treatment of each term in the transport equation when applying the LD FE discretization. The first set of terms are mass-like terms: terms that have no spatial derivatives or factors of r and are proportional to ϕ or ψ; these terms show up in the reaction terms. The second set of terms are also mass-like, but have factors of $1/r$; these terms show up in the angular derivative.

The third set of terms are the streaming terms: they have spatial derivatives and represent advection along $\mathbf{\Omega}$.

$$\text{Term 1} = \sigma_x(r)\psi_k(r), \tag{A.11a}$$

$$\text{Term 2} = \frac{1}{r}\frac{\alpha_{k+1/2}\,\psi_{k+1/2}(r) - \alpha_{k-1/2}\,\psi_{k-1/2}(r)}{w_k} \sim \frac{1}{r}\psi(r), \tag{A.11b}$$

$$\text{Term 3} = \frac{\mu_k}{r}\frac{\partial}{\partial r}\Big(r\psi_k(r)\Big). \tag{A.11c}$$

The mass matrix for cell s is defined as

$$\left(\mathcal{N}_s^p\right)_{i,j} = \int_{r_{s-1/2}}^{r_{s+1/2}} \mathrm{d}r\ r^p\ B_i(r)B_j(r), \tag{A.12}$$

for $i = L, R$ and $j = L, R$.[2]

The standard FE approach of multiplying by a weight function and integrating over all space is performed. The weight functions have local support over the cell, s, that simplifies the bounds of integration. The fact that the cross sections are constant over a cell allows them to be pulled out of the integral.

[2]Since $\mathcal{M}$ was already taken by the moment-to-discrete matrix in angle, $\mathcal{N}$ is used for the mass matrix in space.

The Term 1 system becomes:

$$\int_0^R \mathrm{d}r \;\; r \begin{pmatrix} B_{s,L}(r) \\ \\ B_{s,R}(r) \end{pmatrix} \text{(Term 1)}$$

$$= \int_{r_{s-1/2}}^{r_{s+1/2}} \mathrm{d}r \;\; r \begin{pmatrix} B_{s,L}(r) \\ \\ B_{s,R}(r) \end{pmatrix} \sigma_x(r)\psi_k(r) \qquad \text{(A.13a)}$$

$$= \sigma_x \int_{r_{s-1/2}}^{r_{s+1/2}} \mathrm{d}r \;\; r \begin{pmatrix} B_{s,L}(r) \\ \\ B_{s,R}(r) \end{pmatrix} \begin{pmatrix} B_{s,L}(r) & B_{s,R}(r) \end{pmatrix} \begin{pmatrix} \psi_{k,s,L} \\ \\ \psi_{k,s,R} \end{pmatrix} \qquad \text{(A.13b)}$$

$$= \sigma_x \left\{ \int_{r_{s-1/2}}^{r_{s+1/2}} \mathrm{d}r \;\; r \begin{pmatrix} B_{s,L}(r)B_{s,L}(r) & B_{s,L}(r)B_{s,R}(r) \\ \\ B_{s,R}(r)B_{s,L}(r) & B_{s,R}(r)B_{s,R}(r) \end{pmatrix} \right\} \begin{pmatrix} \psi_{k,s,L} \\ \\ \psi_{k,s,R} \end{pmatrix} \qquad \text{(A.13c)}$$

$$= \sigma_x \begin{pmatrix} (\mathcal{N}_s^1)_{L,L} & (\mathcal{N}_s^1)_{L,R} \\ \\ (\mathcal{N}_s^1)_{R,L} & (\mathcal{N}_s^1)_{R,R} \end{pmatrix} \begin{pmatrix} \psi_{k,s,L} \\ \\ \psi_{k,s,R} \end{pmatrix} \qquad \text{(A.13d)}$$

$$= \sigma_x \, \mathcal{N}_s^1 \, \boldsymbol{\psi}_{k,s}, \qquad \text{(A.13e)}$$

where we define

$$\boldsymbol{\psi}_{k,s} \equiv \begin{pmatrix} \psi_{k,s,L} \\ \\ \psi_{k,s,R} \end{pmatrix}. \qquad \text{(A.13f)}$$

By similar arguments and using the WDD relation Eq. (A.7a), the Term 2 system is:

$$\int_0^R \mathrm{d}r\ r \begin{pmatrix} B_{s,L}(r) \\ \\ B_{s,R}(r) \end{pmatrix} \text{(Term 2)}$$

$$= \int_{r_{s-1/2}}^{r_{s+1/2}} \mathrm{d}r\ r \begin{pmatrix} B_{s,L}(r) \\ \\ B_{s,R}(r) \end{pmatrix} \frac{1}{r} \frac{\alpha_{k+1/2}\,\psi_{k+1/2}(r) - \alpha_{k-1/2}\,\psi_{k-1/2}(r)}{w_k} \quad \text{(A.14a)}$$

$$= \frac{\alpha_{k+1/2}}{w_k}\,\mathcal{N}_s^0\,\boldsymbol{\psi}_{k+1/2,s} - \frac{\alpha_{k-1/2}}{w_k}\,\mathcal{N}_s^0\,\boldsymbol{\psi}_{k-1/2,s} \quad \text{(A.14b)}$$

$$= \left(\frac{\alpha_{k+1/2}}{w_k}\frac{1}{\alpha_k}\right)\,\mathcal{N}_s^0\,\boldsymbol{\psi}_{k,s} - \left(\frac{\alpha_{k-1/2}}{w_k} + \frac{\alpha_{k+1/2}}{w_k}\left(\frac{1-\alpha_k}{\alpha_k}\right)\right)\,\mathcal{N}_s^0\,\boldsymbol{\psi}_{k-1/2,s}. \quad \text{(A.14c)}$$

The outflow matrix represents the loss term for the current cell due to outflow into the downwind cell. It is defined as

$$\mathcal{B}_{s,k}^p = \begin{cases} \begin{pmatrix} -\mu_k\, r_{s-1/2}^p & 0 \\ 0 & 0 \end{pmatrix} & \mu_k < 0, \\ \\ \begin{pmatrix} 0 & 0 \\ 0 & \mu_k\, r_{s+1/2}^p \end{pmatrix} & \mu_k > 0. \end{cases} \quad \text{(A.15)}$$

The within-cell streaming matrix represents the streaming within a cell between internal DOF. It is defined as

$$\left(\mathcal{C}_s^p\right)_{i,j} = \int_{r_{s-1/2}}^{r_{s+1/2}} \mathrm{d}r\ r^p\ \frac{\partial B_i}{\partial r}(r) B_j(r). \quad \text{(A.16)}$$

Note that neither $\mathcal{B}$ nor $\mathcal{C}$ is symmetric.

Finally, the upwind source represents the instreaming from the upwind cell. It is

defined as

$$\mathbf{b}^p_{s-1,k} = \begin{cases} \begin{pmatrix} 0 \\ -\mu_k \, r^p_{s+1/2} \, \psi_{s+1,k,L} \end{pmatrix} & \mu_k < 0, \\ \\ \begin{pmatrix} \mu_k \, r^p_{s-1/2} \, \psi_{s-1,k,R} \\ 0 \end{pmatrix} & \mu_k > 0. \end{cases} \tag{A.17}$$

Using integration by parts and neglecting for now the boundary terms, the Term 3 system is:

$$\int_0^R \mathrm{d}r \;\; r \begin{pmatrix} B_{s,L}(r) \\ B_{s,R}(r) \end{pmatrix} \text{(Term 3)}$$

$$= \int_{r_{s-1/2}}^{r_{s+1/2}} \mathrm{d}r \;\; r \begin{pmatrix} B_{s,L}(r) \\ B_{s,R}(r) \end{pmatrix} \frac{\mu_k}{r} \frac{\partial}{\partial r}\Big(r\psi_k(r)\Big) \tag{A.18a}$$

$$= -\mu_k \int_{r_{s-1/2}}^{r_{s+1/2}} \mathrm{d}r \;\; r \frac{\partial}{\partial r} \begin{pmatrix} B_{s,L}(r) \\ B_{s,R}(r) \end{pmatrix} \psi_k(r) \tag{A.18b}$$

$$= -\mu_k \int_{r_{s-1/2}}^{r_{s+1/2}} \mathrm{d}r \;\; r \frac{\partial}{\partial r} \begin{pmatrix} B_{s,L}(r) \\ B_{s,R}(r) \end{pmatrix} \begin{pmatrix} B_{s,L}(r) & B_{s,R}(r) \end{pmatrix} \begin{pmatrix} \psi_{k,s,L} \\ \psi_{k,s,R} \end{pmatrix} \tag{A.18c}$$

$$= -\mu_k \int_{r_{s-1/2}}^{r_{s+1/2}} \mathrm{d}r \;\; r \begin{pmatrix} \frac{\partial B_{s,L}}{\partial r}(r) B_{s,L}(r) & \frac{\partial B_{s,L}}{\partial r}(r) B_{s,R}(r) \\ \frac{\partial B_{s,R}}{\partial r}(r) B_{s,L}(r) & \frac{\partial B_{s,R}}{\partial r}(r) B_{s,R}(r) \end{pmatrix} \begin{pmatrix} \psi_{k,s,L} \\ \psi_{k,s,R} \end{pmatrix} \tag{A.18d}$$

$$= -\mu_k \begin{pmatrix} (\mathcal{C}^1_s)_{L,L} & (\mathcal{C}^1_s)_{L,R} \\ (\mathcal{C}^1_s)_{R,L} & (\mathcal{C}^1_s)_{R,R} \end{pmatrix} \begin{pmatrix} \psi_{k,s,L} \\ \psi_{k,s,R} \end{pmatrix} \tag{A.18e}$$

$$= -\mu_k \, \mathcal{C}^1_s \, \boldsymbol{\psi}_{k,s}. \tag{A.18f}$$

The boundary contribution to Term 3 is upwinded and is shown below for $\mu_k > 0$:

$$\int_0^R \mathrm{d}r \; r \begin{pmatrix} B_{s,L}(r) \\ \\ B_{s,R}(r) \end{pmatrix} (\text{Term 3}) \Big|_{\text{bdr}}$$

$$= \mu_k \left[\begin{pmatrix} B_{s,L}(r) \\ \\ B_{s,R}(r) \end{pmatrix} r\psi_k(r) \right] \Bigg|_{r_{s-1/2}}^{r_{s+1/2}} \tag{A.19a}$$

$$= \mu_k \left[\begin{pmatrix} B_{s,L}(r_{s+1/2}) \\ \\ B_{s,R}(r_{s+1/2}) \end{pmatrix} r_{s+1/2}\, \psi_{k,s,R} - \begin{pmatrix} B_{s,L}(r_{s-1/2}) \\ \\ B_{s,R}(r_{s-1/2}) \end{pmatrix} r_{s-1/2}\, \psi_{k,s-1,R} \right] \tag{A.19b}$$

$$= \begin{pmatrix} 0 \\ \\ 1 \end{pmatrix} \mu_k\, r_{s+1/2}\, \psi_{k,s,R} - \begin{pmatrix} 1 \\ \\ 0 \end{pmatrix} \mu_k\, r_{s-1/2}\, \psi_{k,s-1,R} \tag{A.19c}$$

$$= \boldsymbol{\mathcal{B}}^1_{s,k}\, \boldsymbol{\psi}_{k,s} - \mathbf{b}^1_{s-1,k}. \tag{A.19d}$$

A similar derivation yields the same result for $\mu_k < 0$.

To find the equation that must be solved per cell, Eq. (A.1) in the cylindrical case (Eq. (A.2a)) with the angular FD approximation (Eq. (A.6a)) is used:

$$\int_0^R \mathrm{d}r \; r \begin{pmatrix} B_{s,L}(r) \\ \\ B_{s,R}(r) \end{pmatrix} \left\{ \frac{\mu_k}{r} \frac{\partial}{\partial r}\big(r\, \psi_{g,k}(r)\big) + \frac{1}{r} \frac{\alpha_{k+1/2}\, \psi_{g,k+1/2}(r) - \alpha_{k-1/2}\, \psi_{g,k-1/2}(r)}{w_k} + \right.$$

$$\left. \Sigma_{t,g}(r)\, \psi_{g,k}(r) = \sum_{g'=1}^{G} \sum_{k'} \mathcal{M}_{k,k'}\, \Sigma_{\text{scat},g'\to g,k'}(r)\; \phi_{g',k'}(r) + q_{g,k}(r) \right\}, \tag{A.20a}$$

$$\int_0^R \mathrm{d}r \; r \begin{pmatrix} B_{s,L}(r) \\ \\ B_{s,R}(r) \end{pmatrix} \Big\{ \text{Term 3} + \text{Term 2} +$$

$$\text{Term 1} = \text{Like Term 1} \Big\}, \tag{A.20b}$$

$$\mathcal{B}^1_{s,k}\,\boldsymbol{\psi}_{s,g,k} - \mathbf{b}^1_{s-1,g,k} - \mu_k\,\mathcal{C}^1_s\,\boldsymbol{\psi}_{s,g,k} + \left(\frac{\alpha_{k+1/2}}{w_k}\frac{1}{\alpha_k}\right)\,\mathcal{N}^0_s\,\boldsymbol{\psi}_{k,s} -$$
$$\left(\frac{\alpha_{k-1/2}}{w_k} + \frac{\alpha_{k+1/2}}{w_k}\left(\frac{1-\alpha_k}{\alpha_k}\right)\right)\,\mathcal{N}^0_s\,\boldsymbol{\psi}_{k-1/2,s} + \Sigma_{t,s,g}\,\mathcal{N}^1_s\,\boldsymbol{\psi}_{k,s}$$
$$= \mathcal{N}^1_s\left[\sum_{g'=1}^{G}\sum_{k'}\mathcal{M}_{k,k'}\,\Sigma_{\text{scat},s,g'\to g,k'}\,\boldsymbol{\phi}_{s,g',k'} + \mathbf{q}_{s,g,k}\right], \tag{A.20c}$$

$$\mathbf{S}_{s,g,k} = \sum_{g'=1}^{G}\sum_{k'}\mathcal{M}_{k,k'}\,\Sigma_{\text{scat},s,g'\to g,k'}\begin{pmatrix}\phi_{s,g',k',L}\\ \\ \phi_{s,g',k',R}\end{pmatrix}, \tag{A.20d}$$

$$\mathbf{q}_{s,g,k} = \frac{\chi_{s,g}}{k_{\text{eff}}}\sum_{g'=1}^{G}\nu\Sigma_{f,s,g'}\begin{pmatrix}\phi_{s,g',0,L}\\ \\ \phi_{s,g',0,R}\end{pmatrix}, \tag{A.20e}$$

$$\left[\mathcal{B}^1_{s,k} - \mu_k\,\mathcal{C}^1_s + \Sigma_{t,s,g}\,\mathcal{N}^1_s + \left(\frac{\alpha_{k+1/2}}{w_k}\right)\left(\frac{1}{\alpha_k}\right)\mathcal{N}^0_s\right]\boldsymbol{\psi}_{s,g,k}$$
$$= \mathcal{N}^1_s\Big(\mathbf{S}_{s,g,k} + \mathbf{q}_{s,g,k}\Big) + \left[\frac{\alpha_{k-1/2}}{w_k} + \frac{\alpha_{k+1/2}}{w_k}\left(\frac{1-\alpha_k}{\alpha_k}\right)\right]\mathcal{N}^1_s\boldsymbol{\psi}_{s,g,k-1/2} + \mathbf{b}^1_{s-1,g,k}. \tag{A.20f}$$

In general, the terms on the right-hand side are lagged and the sweep is used to solve for $\boldsymbol{\psi}_{s,g,k}$ by inverting the left-hand side. Instead of storing all the $\boldsymbol{\psi}_{s,g,k}$, their contributions to the ϕ moments are kept. Further discussion may be found in the Solution Techniques section.

A.3.1 The Starting Directions

When the transport equation is solved, there is a sweep in both angle and space due to the advection terms on both variables. If k_l is the angular index on level l, solving for ψ_{k_l} requires knowledge of $\psi_{k_l-1/2}$. The WDD relation can then be used to solve for $\psi_{k_l+1/2}$, which replaces the current $\psi_{k_l-1/2}$ to minimize storage. This occurs for each level independently. However, this algorithm is a non-starter: it does not specify a way to get the first $\psi_{k_l+1/2}$ on a level.

On each polar level, the first $\psi_{k_l+1/2}$ is found using a special starting-direction sweep. Assuming that ψ is not singular in $\omega = -\pi$, it can be shown that streaming in $\omega = \pm\pi$ is equivalent to streaming in a slab. Specifically, the angular derivative term drops out both physically and mathematically (easiest to see by setting $\eta = 0$ in Eq. (A.2b)) and the cylindrical coordinates become degenerate to a slab. For these reasons, a slab sweep is performed for each level to start off the angular derivative.

An ending-direction sweep (per level) is sometimes performed as well to satisfy specular reflecting boundary conditions needed to begin the starting-direction sweep (per level). Since we use white boundary conditions, we ignore this option.

The flux at the center may be treated as specular reflection to allow correct modeling of central voids. Since we do not anticipate needing these voids and because technically the starting-direction flux is the only one (per level) which makes it to the center, we use the standard Morel treatment of isotropic center flux given by the starting direction.

The sources for the starting-direction flux are computed in two ways. For the scattering and distributed sources, a special $\mathcal{M}_{\text{start}}$ is used where the starting-direction ($\omega = -\pi$) $\mathbf{\Omega}_j$ are used in Eq. (A.5a). For the boundary source, a linear extrapolation is done using ψ_{k_l} and ψ_{k_l+1} on the boundary and Eq. (A.7a)-type weighting. For the

first level, since there is only one angle with $\mu_k < 0$, the starting-direction flux is simply set equal to ψ_{1_1}.

For the starting direction, there is no angular derivative and the proper FE integral is to weight with r^0 instead of r^1. The equivalent of Eq. (A.20f) for the starting direction may be found through similar technique to that shown above and is given by

$$\int_0^R \mathrm{d}r \begin{pmatrix} B_{s,L}(r) \\ \\ B_{s,R}(r) \end{pmatrix} \left\{ \mu_k \frac{\partial \psi_{g,k}}{\partial r}(r) + \Sigma_{t,g}(r)\, \psi_{g,k}(r) \right. \\ \left. = \sum_{g'=1}^{G} \sum_{k'} \mathcal{M}_{k,k'}\, \Sigma_{\mathrm{scat},g'\to g,k'}(r)\; \phi_{g',k'}(r) + q_{g,k}(r) \right\}, \tag{A.21a}$$

$$\left[\mathcal{B}^0_{s,k} - \mu_k\, \mathcal{C}^0_s + \Sigma_{t,s,g}\, \mathcal{N}^0_s \right] \boldsymbol{\psi}_{s,g,k} = \mathcal{N}^0_s \Big(\mathbf{S}_{s,g,k} + \mathbf{q}_{s,g,k} \Big) + \mathbf{b}^0_{s-1,g,k}. \tag{A.21b}$$

The starting-direction flux provides the initial point for the WDD in angle over the level.

A.3.2 The Center Cell

When sweeping in the non-starting-direction $\mu_k < 0$ directions, for the innermost cell (near $r = 0$), both the incoming flux and outgoing flux were known. The outgoing flux was simply the level-dependent starting-direction flux solved earlier: the only way to pass through $r = 0$ is through directions with $\omega = \pm\pi$; thus, the flux there must be isotropic and equal to the starting-direction flux. For this cell, instead of testing with the left- and right-basis functions, the transport equation was tested with unity. This gave a 1x1 system to solve for ψ_R. This formulation was chosen to preserve balance.

The cell-integrated averages over the first cell are defined as:

$$\overline{f^p} = \int_0^{r_{3/2}} \mathrm{d}r \;\; r^p f(r). \tag{A.22}$$

The streaming term contains an exact derivative:

$$\int_0^{r_{3/2}} \mathrm{d}r \;\; r \, \frac{\mu_k}{r} \frac{\partial}{\partial r}\big(r \, \psi_{g,k}(r)\big) = \mu_k \int_0^{r_{3/2}} \mathrm{d}r \;\; \frac{\partial}{\partial r}\big(r \, \psi_{g,k}(r)\big) \tag{A.23a}$$

$$= \mu_k \big(r_{3/2} \, \psi_{g,k}(r_{3/2}) - 0 \, \psi_{g,k}(0)\big) \tag{A.23b}$$

$$= \mu_k \, r_{3/2} \; \psi_{\mathrm{upwind},g,k}. \tag{A.23c}$$

The solution for the one unknown in the inner-most cell for $\mu_k < 0$ is:

$$\int_0^{r_{3/2}} \mathrm{d}r \;\; r \left\{ \frac{\mu_k}{r} \frac{\partial}{\partial r}\big(r \, \psi_{g,k}(r)\big) + \frac{1}{r} \frac{\alpha_{k+1/2} \, \psi_{g,k+1/2}(r) - \alpha_{k-1/2} \, \psi_{g,k-1/2}(r)}{w_k} + \right.$$

$$\left. \Sigma_{t,g}(r) \, \psi_{g,k}(r) = \sum_{g'=1}^{G} \sum_{k'} \mathcal{M}_{k,k'} \, \Sigma_{\mathrm{scat},g' \to g,k'}(r) \; \phi_{g',k'}(r) + q_{g,k}(r) \right\}, \tag{A.24a}$$

$$\mu_k \, r_{3/2} \; \psi_{\mathrm{upwind},g,k} + \frac{\alpha_{k+1/2} \, \overline{\psi^0}_{g,k+1/2} - \alpha_{k-1/2} \, \overline{\psi^0}_{g,k-1/2}}{w_k} + \Sigma_{t,1,g} \overline{\psi^1}_{g,k}$$

$$= \overline{S^1}_{1,g,k} + \overline{q^1}_{1,g,k}. \tag{A.24b}$$

Splitting up $\overline{\psi^1}_{g,k}$ into its (basis function) component parts and using the WDD

relation in angle yields:

$$\begin{aligned}
&\mu_k\, r_{3/2}\, \psi_{\text{upwind},g,k} - \\
&\quad \left[\frac{\alpha_{k-1/2}}{w_k} + \frac{\alpha_{k+1/2}}{w_k}\left(\frac{1-\alpha_k}{\alpha_k}\right)\right]\left(\overline{B^0}_{1,L}\,\psi_{1,g,k-1/2,L} + \overline{B^0}_{1,R}\,\psi_{1,g,k-1/2,R}\right) + \\
&\qquad \left(\frac{\alpha_{k+1/2}}{w_k}\right)\left(\frac{1}{\alpha_k}\right)\left(\overline{B^0}_{1,L}\,\psi_{1,g,k,L} + \overline{B^0}_{1,R}\,\psi_{1,g,k,R}\right) + \\
&\qquad\qquad \Sigma_{t,1,g}\left(\overline{B^1}_{1,L}\,\psi_{1,g,k,L} + \overline{B^1}_{1,R}\,\psi_{1,g,k,R}\right) \\
&\qquad\qquad\qquad = \overline{S^1}_{1,g,k} + \overline{q^1}_{1,g,k}.
\end{aligned} \tag{A.24c}$$

Solving for the unknown $\psi_{1,g,k,R}$ and realizing that $\psi_{1,g,k,L}$ is the center value from the starting direction sweep yields:

$$\begin{aligned}
&\psi_{1,g,k,R} \\
&\quad = \Bigg\{ \left[\frac{\alpha_{k-1/2}}{w_k} + \frac{\alpha_{k+1/2}}{w_k}\left(\frac{1-\alpha_k}{\alpha_k}\right)\right]\left(\overline{B^0}_{1,L}\,\psi_{1,g,k-1/2,L} + \overline{B^0}_{1,R}\,\psi_{1,g,k-1/2,R}\right) - \\
&\qquad \left[\Sigma_{t,1,g}\overline{B^1}_{1,L} + \left(\frac{\alpha_{k+1/2}}{w_k}\right)\left(\frac{1}{\alpha_k}\right)\overline{B^0}_{1,L}\right]\psi_{\text{center},g,k} - \mu_k\, r_{3/2}\, \psi_{\text{upwind}} + \\
&\qquad\qquad \overline{S^1}_{1,g,k} + \overline{q^1}_{1,g,k}\Bigg\}\Bigg\{\Sigma_{t,1,g}\overline{B^1}_{1,R} + \left(\frac{\alpha_{k+1/2}}{w_k}\right)\left(\frac{1}{\alpha_k}\right)\overline{B^0}_{1,R}\Bigg\}^{-1}.
\end{aligned} \tag{A.24d}$$

A.4 Angular Acceleration (cont.)

Acceleration of the scattering source is commonly used in S_n transport. Here we define S_2-synthetic acceleration (S2SA). While similar to diffusion-synthetic acceleration (DSA), S2SA uses the S_2 equations with Gauß-Legendre quadrature instead of the diffusion (or P_1) equation. In one-dimensional (Cartesian) slabs without time

dependence and for one group, these can be shown to be equivalent. We use S2SA because it allows us to reuse our angular and spatial discretization work.

To create the system, the 0^{th} and 1^{st} moments of transport equation, Eq. (A.1), were taken:

$$\frac{1}{4\pi}\int_{4\pi}\mathrm{d}\Omega\ \begin{pmatrix}1\\ \mu\end{pmatrix}\left\{\frac{\mu}{r}\frac{\partial}{\partial r}\big(r\,\psi_g(r,\mathbf{\Omega})\big)-\frac{1}{r}\frac{\partial}{\partial\omega}\big(\eta\,\psi_g(r,\mathbf{\Omega})\big)+\Sigma_{t,g}(r)\,\psi_g(r,\mathbf{\Omega})\right.$$
$$\left.=\frac{1}{4\pi}\int_{4\pi}\mathrm{d}\Omega'\ \Sigma_{\text{scat},g\to g}(r,\mathbf{\Omega}'\cdot\mathbf{\Omega})\,\phi_g(r,\mathbf{\Omega}')+q_g(r,\mathbf{\Omega})\right\}. \tag{A.25a}$$

Using the definitions of ϕ and J as the 0^{th} and 1^{st} moments of $\psi(\mathbf{\Omega})$, respectively, and the definition of the $Y_l^m(\mathbf{\Omega})$ moments of the scattering and extraneous source terms, the right-hand side simplifies. Additionally, we assume ψ to be linearly anisotropic in angle. The result is:

$$\frac{1}{4\pi}\int_{4\pi}\mathrm{d}\Omega\ \left\{\frac{1}{r}\frac{\partial}{\partial r}\left[r\begin{pmatrix}\mu\,\psi_g(r,\mathbf{\Omega})\\ \mu^2\,\psi_g(r,\mathbf{\Omega})\end{pmatrix}\right]+\frac{1}{r}\begin{pmatrix}1\\ \mu\end{pmatrix}\frac{\partial}{\partial\omega}\big(\eta\,\psi_g(r,\mathbf{\Omega})\big)\right.$$
$$\left.+\Sigma_{t,g}(r)\begin{pmatrix}\psi_g(r,\mathbf{\Omega})\\ \mu\,\psi_g(r,\mathbf{\Omega})\end{pmatrix}\right\}=\begin{pmatrix}\Sigma_{\text{scat},g\to g,0}(r)\,\phi(r)\\ \Sigma_{\text{scat},g\to g,1}(r)\,J(r)\end{pmatrix}+\begin{pmatrix}q_{g,0}(r)\\ q_{g,1}(r)\end{pmatrix} \tag{A.25b}$$

The spatial derivative term may be simplified by using the P_1 approximation for $\psi(\mathbf{\Omega})$,

$$\psi(\mathbf{\Omega})\simeq\phi+3\,\mu\,J, \tag{A.26a}$$

$$\psi(\mathbf{\Omega})\simeq\phi+3\,\sqrt{1-\xi^2}\,\cos\omega\,J. \tag{A.26b}$$

Then

$$\frac{1}{4\pi}\int\limits_{4\pi} \mathrm{d}\Omega \;\; \mu\,\psi(\mathbf{\Omega}) = \frac{1}{4\pi}\int\limits_{4\pi} \mathrm{d}\Omega \;\; \left(\mu\,\phi + 3\mu^2\,J\right)$$
$$\frac{1}{4\pi}\int\limits_{4\pi} \mathrm{d}\Omega \;\; \mu\,\psi(\mathbf{\Omega}) = J. \tag{A.27a}$$

and

$$\frac{1}{4\pi}\int\limits_{4\pi} \mathrm{d}\Omega \;\; \mu^2\,\psi(\mathbf{\Omega}) = \frac{1}{4\pi}\int\limits_{4\pi} \mathrm{d}\Omega \;\; \left(\mu^2\,\phi + 3\mu^3\,J\right)$$
$$\frac{1}{4\pi}\int\limits_{4\pi} \mathrm{d}\Omega \;\; \mu^2\,\psi(\mathbf{\Omega}) = \frac{1}{3}\phi. \tag{A.27b}$$

The angular derivative term may be simplified by expanding the integral or using integration by parts

$$\frac{1}{4\pi}\int\limits_{4\pi} \mathrm{d}\Omega \;\frac{\partial}{\partial\omega}\big(\eta\,\psi(\mathbf{\Omega})\big) = \frac{1}{4\pi}\int_{-1}^{1}\mathrm{d}\xi \;\int_{0}^{2\pi}\mathrm{d}\omega \;\frac{\partial}{\partial\omega}\big(\eta\,\psi(\mathbf{\Omega})\big) \tag{A.28a}$$
$$= \frac{1}{4\pi}\int_{-1}^{1}\mathrm{d}\xi \;\left(\eta\,\psi(\mathbf{\Omega})\right)\Big|_{\omega=0}^{\omega=2\pi} \tag{A.28b}$$
$$\frac{1}{4\pi}\int\limits_{4\pi} \mathrm{d}\Omega \;\frac{\partial}{\partial\omega}\big(\eta\,\psi(\mathbf{\Omega})\big) = 0, \tag{A.28c}$$

as $\eta(\omega = 0) = \eta(\omega = 2\pi) = 0$.

The first moment is

$$\frac{1}{4\pi}\int\limits_{4\pi} \mathrm{d}\Omega\ \mu\frac{\partial}{\partial\omega}\big(\eta\,\psi(\mathbf{\Omega})\big) = \frac{1}{4\pi}\int_{-1}^{1}\mathrm{d}\xi\ \int_{0}^{2\pi}\mathrm{d}\omega\ \mu\frac{\partial}{\partial\omega}\big(\eta\,\psi(\mathbf{\Omega})\big) \tag{A.28d}$$

$$= \frac{1}{4\pi}\int_{-1}^{1}\mathrm{d}\xi\ \left[\big(\mu\,\eta\,\psi(\mathbf{\Omega})\big)\Big|_{\omega=0}^{\omega=2\pi} - \int_{0}^{2\pi}\mathrm{d}\omega\ \frac{\partial\mu}{\partial\omega}\eta\,\psi(\mathbf{\Omega})\right] \tag{A.28e}$$

$$= \frac{1}{4\pi}\int_{-1}^{1}\mathrm{d}\xi\ \int_{0}^{2\pi}\mathrm{d}\omega\ \eta^2\psi(\mathbf{\Omega}) \tag{A.28f}$$

$$= \frac{1}{4\pi}\int_{-1}^{1}\mathrm{d}\xi\ \int_{0}^{2\pi}\mathrm{d}\omega\ \eta^2\left(\phi + 3\mu\,J\right) \tag{A.28g}$$

$$\frac{1}{4\pi}\int\limits_{4\pi} \mathrm{d}\Omega\ \mu\frac{\partial}{\partial\omega}\big(\eta\,\psi(\mathbf{\Omega})\big) = \frac{1}{3}\phi, \tag{A.28h}$$

where the facts that $\partial\mu/\partial\omega = -\eta$ and that $\mu\eta^2$ integrated over ω is zero were used.

Using the above procedures yields the system

$$\frac{1}{r}\frac{\partial}{\partial r}\left[r\begin{pmatrix}0 & 1\\ \frac{1}{3} & 0\end{pmatrix}\begin{pmatrix}\phi_g(r)\\ J_g(r)\end{pmatrix}\right] - \frac{1}{r}\begin{pmatrix}0 & 0\\ \frac{1}{3} & 0\end{pmatrix}\begin{pmatrix}\phi_g(r)\\ J_g(r)\end{pmatrix} + \Sigma_{t,g}(r)\begin{pmatrix}1 & 0\\ 0 & 1\end{pmatrix}\begin{pmatrix}\phi_g(r)\\ J_g(r)\end{pmatrix}$$
$$= \begin{pmatrix}\Sigma_{\mathrm{scat},g\to g,0} & 0\\ 0 & \Sigma_{\mathrm{scat},g\to g,1}\end{pmatrix}\begin{pmatrix}\phi_g(r)\\ J_g(r)\end{pmatrix} + \begin{pmatrix}q_{g,0}(r)\\ q_{g,1}(r)\end{pmatrix}. \tag{A.29}$$

This gives a system with ϕ and J basis. This basis is transformed into a ψ^+ and

ψ^- basis by applying a similarity transform.

$$\begin{pmatrix} \phi \\ J \end{pmatrix} = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2\sqrt{3}} & -\frac{1}{2\sqrt{3}} \end{pmatrix} \begin{pmatrix} \psi^+ \\ \psi^- \end{pmatrix}, \tag{A.30a}$$

$$\begin{pmatrix} \psi^+ \\ \psi^- \end{pmatrix} = \begin{pmatrix} 1 & \sqrt{3} \\ 1 & -\sqrt{3} \end{pmatrix} \begin{pmatrix} \phi \\ J \end{pmatrix}. \tag{A.30b}$$

The S_2 discrete-to-moment matrix, Eq. (A.30a) was used to define ϕ and J. The S_2 moment-to-discrete matrix, Eq. (A.30b), was multiplied on the left to complete the similarity transform. The similarity transform was done because it is known how to upwind the ψ; the upwinding of ϕ and J is not trivial. Additionally, it is expected that for diffusive systems, ψ^+ and ψ^- will be of similar magnitude, while ϕ will be much larger than J. Transforming to a basis with similar quantities in magnitude may decrease the condition number of the matrix, increasing the number of significant digits in the output.

The resulting system is

$$\frac{1}{r}\frac{\partial}{\partial r}\left[r\begin{pmatrix}\frac{1}{\sqrt{3}} & 0\\ 0 & -\frac{1}{\sqrt{3}}\end{pmatrix}\begin{pmatrix}\psi_g^+(r)\\ \psi_g^-(r)\end{pmatrix}\right]+ \frac{1}{r}\begin{pmatrix}-\frac{1}{2\sqrt{3}} & -\frac{1}{2\sqrt{3}}\\ \frac{1}{2\sqrt{3}} & \frac{1}{2\sqrt{3}}\end{pmatrix}\begin{pmatrix}\psi_g^+(r)\\ \psi_g^-(r)\end{pmatrix}+\Sigma_{t,g}(r)\begin{pmatrix}1 & 0\\ 0 & 1\end{pmatrix}\begin{pmatrix}\psi_g^+(r)\\ \psi_g^-(r)\end{pmatrix} = \left[\Sigma_{\text{scat},g\to g,0}\begin{pmatrix}\frac{1}{2} & \frac{1}{2}\\ \frac{1}{2} & \frac{1}{2}\end{pmatrix}+\sqrt{3}\,\Sigma_{\text{scat},g\to g,1}\begin{pmatrix}\frac{1}{2\sqrt{3}} & -\frac{1}{2\sqrt{3}}\\ -\frac{1}{2\sqrt{3}} & \frac{1}{2\sqrt{3}}\end{pmatrix}\right]\begin{pmatrix}\psi_g^+(r)\\ \psi_g^-(r)\end{pmatrix}+ \begin{pmatrix}q_{g,0}(r)+\sqrt{3}q_{g,1}(r)\\ q_{g,0}(r)-\sqrt{3}q_{g,1}(r)\end{pmatrix}. \tag{A.31}$$

was then treating using standard LD, including upwinding the $\psi^{\pm}$ as above.

The contribution to the full matrix from each cell is given below. Horizontal lines indicate zero entries of height two rows; zeros are explicitly given to show offsets and matrix sizes. Note that the $\mathcal{C}$ and $\mathcal{M}$ matrices are 2x2 matrices. The left-hand side is

$$
\begin{pmatrix}
0 & -\frac{1}{\sqrt{3}}\mathcal{C}_s^1 & \multicolumn{3}{c}{\multirow{3}{*}{———————}} \\
 & & & & \\
\multicolumn{2}{c}{\multirow{2}{*}{———————}} & \frac{1}{\sqrt{3}}\mathcal{C}_s^1 & 0 & \\
 & & & &
\end{pmatrix}
\begin{pmatrix}
\psi^+_{s-1,g,R} \\ \psi^+_{s,g,L} \\ \psi^+_{s,g,R} \\ \psi^-_{s,g,L} \\ \psi^-_{s,g,R} \\ \psi^-_{s+1,g,L}
\end{pmatrix} +
$$

$$
\begin{pmatrix}
0 & -\frac{1}{2\sqrt{3}}\mathcal{M}_s^0 & -\frac{1}{2\sqrt{3}}\mathcal{M}_s^0 & 0 \\
 & & & \\
0 & \frac{1}{2\sqrt{3}}\mathcal{M}_s^0 & \frac{1}{2\sqrt{3}}\mathcal{M}_s^0 & 0 \\
 & & &
\end{pmatrix}
\begin{pmatrix}
\psi^+_{s-1,g,R} \\ \psi^+_{s,g,L} \\ \psi^+_{s,g,R} \\ \psi^-_{s,g,L} \\ \psi^-_{s,g,R} \\ \psi^-_{s+1,g,L}
\end{pmatrix} +
$$

$$
\begin{pmatrix}
0 & \Sigma_{t,s,g}\mathcal{M}_s^1 & \multicolumn{3}{c}{\multirow{3}{*}{———————}} \\
 & & & & \\
\multicolumn{2}{c}{\multirow{2}{*}{———————}} & \Sigma_{t,s,g}\mathcal{M}_s^1 & 0 & \\
 & & & &
\end{pmatrix}
\begin{pmatrix}
\psi^+_{s-1,g,R} \\ \psi^+_{s,g,L} \\ \psi^+_{s,g,R} \\ \psi^-_{s,g,L} \\ \psi^-_{s,g,R} \\ \psi^-_{s+1,g,L}
\end{pmatrix} +
$$

$$
\begin{pmatrix}
-\frac{1}{\sqrt{3}}r_{i,L} & 0 & 0 & \multicolumn{3}{c}{\multirow{3}{*}{———————}} \\
0 & 0 & \frac{1}{\sqrt{3}}r_{i,R} & & & \\
 & & & & & \\
\multicolumn{3}{c}{\multirow{2}{*}{———————}} & \frac{1}{\sqrt{3}}r_{i,L} & 0 & 0 \\
 & & & 0 & 0 & -\frac{1}{\sqrt{3}}r_{i,R}
\end{pmatrix}
\begin{pmatrix}
\psi^+_{s-1,g,R} \\ \psi^+_{s,g,L} \\ \psi^+_{s,g,R} \\ \psi^-_{s,g,L} \\ \psi^-_{s,g,R} \\ \psi^-_{s+1,g,L}
\end{pmatrix}. \tag{A.32a}
$$

The right-hand side is

$$\begin{pmatrix} 0 & \frac{1}{2}\Sigma_{\text{scat},s,g\to g,0}\,\mathcal{M}_s^1 & \frac{1}{2}\Sigma_{\text{scat},s,g\to g,0}\,\mathcal{M}_s^1 & 0 \\ \\ 0 & \frac{1}{2}\Sigma_{\text{scat},s,g\to g,0}\,\mathcal{M}_s^1 & \frac{1}{2}\Sigma_{\text{scat},s,g\to g,0}\,\mathcal{M}_s^1 & 0 \\ \\ \end{pmatrix} \begin{pmatrix} \psi^+_{s-1,g,R} \\ \psi^+_{s,g,L} \\ \psi^+_{s,g,R} \\ \psi^-_{s,g,L} \\ \psi^-_{s,g,R} \\ \psi^-_{s+1,g,L} \end{pmatrix} +$$

$$\begin{pmatrix} 0 & \frac{1}{2}\Sigma_{\text{scat},s,g\to g,1}\,\mathcal{M}_s^1 & -\frac{1}{2}\Sigma_{\text{scat},s,g\to g,1}\,\mathcal{M}_s^1 & 0 \\ \\ 0 & -\frac{1}{2}\Sigma_{\text{scat},s,g\to g,1}\,\mathcal{M}_s^1 & \frac{1}{2}\Sigma_{\text{scat},s,g\to g,1}\,\mathcal{M}_s^1 & 0 \\ \\ \end{pmatrix} \begin{pmatrix} \psi^+_{s-1,g,R} \\ \psi^+_{s,g,L} \\ \psi^+_{s,g,R} \\ \psi^-_{s,g,L} \\ \psi^-_{s,g,R} \\ \psi^-_{s+1,g,L} \end{pmatrix} +$$

$$\begin{pmatrix} 0 & \mathcal{M}_s^1 & \text{———} \\ \\ \text{———} & \mathcal{M}_s^1 & 0 \\ \\ \end{pmatrix} \begin{pmatrix} q^+_{s-1,g,R} \\ q^+_{s,g,L} \\ q^+_{s,g,R} \\ q^-_{s,g,L} \\ q^-_{s,g,R} \\ q^-_{s+1,g,L} \end{pmatrix}, \qquad \text{(A.32b)}$$

where q is assumed isotropic. If the individual components from each cell are added together in the global matrix, a 7-banded system ensues.

The S2SA system is defined independently for each group. That is, it only converges within-group scattering. Across-group scattering is taken care of with a Gauß-Seidel (forward) sweep in energy.

The source for the S2SA was the difference between the scalar flux calculated after and before a sweep multiplied by the 0^{th} moment of the scattering cross section. In this way, as the solution converged, the source and solution of the S2SA would go to zero. If l is the iteration index for the scattering source iteration, then

$$q_{s,g,L/R} = \Sigma_{\text{scat},s,g\to g,0}\left(\phi^{l+1/2}_{s,g,L/R} - \phi^{l}_{s,g,L/R}\right) \tag{A.33}$$

Special care is needed on the boundaries. For the first cell, the instreaming term $\psi^{+}_{-1,g,R}$ does not exist. However, since this term is multiplied by $r_{0,L} = 0$, it does not contribute. For the last cell, the boundary is treated as a white boundary with albedo. The incident flux $\psi^{-}_{N+1,g,L}$ is given by the boundary relation in terms of $\psi^{+}_{N,g,R}$. For an albedo of zero, this is the vacuum condition; for an albedo of unity, this is similar to a reflective boundary.

The BC for the S2SA was the Marshak BC using quadrature to compute the half-integrals. It was desired that balance on the boundary be preserved after the update for white albedo boundaries. This condition is exact for vacuum and white boundaries, while being approximate for reflective and reflective albedo boundaries.

$$\sum_{(\mathbf{\Omega}_k\cdot\hat{\mathbf{n}})<0} |\mathbf{\Omega}_k\cdot\hat{\mathbf{n}}|\, w_k\, \psi^{l+1}_k(R) = \alpha \sum_{(\mathbf{\Omega}_k\cdot\hat{\mathbf{n}})>0} |\mathbf{\Omega}_k\cdot\hat{\mathbf{n}}|\, w_k\, \psi^{l+1}_k(R), \tag{A.34a}$$

$$\sum_{\mu_k<0} |\mu_k|\, w_k \left(\psi^{l+1}_k - \psi^{l+1/2}_k + \psi^{l+1/2}_k\right) = \alpha \sum_{\mu_k>0} |\mu_k|\, w_k \left(\psi^{l+1}_k - \psi^{l+1/2}_k + \psi^{l+1/2}_k\right), \tag{A.34b}$$

$$\sum_{\mu_k<0} |\mu_k|\, w_k\, \delta\psi_k - \alpha \sum_{\mu_k>0} |\mu_k|\, w_k\, \delta\psi_k = \alpha \sum_{\mu_k>0} |\mu_k|\, w_k\, \psi^{l+1/2}_k - \sum_{\mu_k<0} |\mu_k|\, w_k\, \psi^{l+1/2}_k. \tag{A.34c}$$

Defining

$$J_{\text{in}} = \alpha \sum_{\mu_k>0} |\mu_k|\, w_k\, \psi_k^{l+1/2} - \sum_{\mu_k<0} |\mu_k|\, w_k\, \psi_k^{l+1/2}, \tag{A.35a}$$

and taking the P_1 approximation for $\delta\psi_k$,

$$\delta\psi_k = \delta\phi + 3\mu_k \delta J, \tag{A.35b}$$

the BC becomes:

$$\left(\sum_{\mu_k<0} |\mu_k|\, w_k - \alpha \sum_{\mu_k>0} |\mu_k|\, w_k \right) \delta\phi +$$

$$3 \left(\sum_{\mu_k<0} |\mu_k| \mu_k\, w_k - \alpha \sum_{\mu_k>0} |\mu_k| \mu_k\, w_k \right) \delta J = J_{\text{in}}, \tag{A.36a}$$

$$\langle \mu \rangle (1-\alpha)\, \delta\phi - 3\langle \mu^2 \rangle (1+\alpha)\, \delta J = J_{\text{in}}, \tag{A.36b}$$

where

$$\langle \mu^2 \rangle = \sum_{\mu_k>0} \mu_k^2 w_k \tag{A.37a}$$

$$= \frac{1}{2} \sum_k \mu_k^2 w_k \tag{A.37b}$$

$$= \frac{1}{4\pi} \int_0^1 \mathrm{d}\mu\ \mu \int_0^{2\pi} \mathrm{d}\omega\ , \tag{A.37c}$$

$$\langle \mu^2 \rangle = \frac{1}{6}, \tag{A.37d}$$

$$\langle \mu \rangle = \sum_{\mu_k>0} \mu_k w_k, \tag{A.37e}$$

where Eq. (A.37b) follows by evenness and Eq. (A.37c) follows by exactness of all quadrature sets in integrating at least the quadratics.

Applying Eq. (A.30b) to Eq. (A.36b) yields the desired result,

$$J_{\text{in}} = \begin{pmatrix} \langle\mu\rangle(1-\alpha) & -\frac{1}{2}(1+\alpha) \end{pmatrix} \begin{pmatrix} \delta\phi \\ \\ \delta J \end{pmatrix}, \tag{A.38a}$$

$$J_{\text{in}} = \begin{pmatrix} \langle\mu\rangle(1-\alpha) & -\frac{1}{2}(1+\alpha) \end{pmatrix} \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \\ \\ \frac{1}{2\sqrt{3}} & -\frac{1}{2\sqrt{3}} \end{pmatrix} \begin{pmatrix} \delta\psi^{+} \\ \\ \delta\psi^{-} \end{pmatrix}, \tag{A.38b}$$

$$4\sqrt{3}\left(\frac{1}{1+\alpha}\right) J_{\text{in}} = \begin{pmatrix} 2\langle\mu\rangle\left(\frac{1-\alpha}{1+\alpha}\right) & -1 \end{pmatrix} \begin{pmatrix} \sqrt{3} & \sqrt{3} \\ \\ 1 & -1 \end{pmatrix} \begin{pmatrix} \delta\psi^{+} \\ \\ \delta\psi^{-} \end{pmatrix}, \tag{A.38c}$$

$$4\sqrt{3}\left(\frac{1}{1+\alpha}\right) J_{\text{in}} = \left(2\sqrt{3}\langle\mu\rangle\left(\frac{1-\alpha}{1+\alpha}\right) - 1\right)\delta\psi^{+} + \left(2\sqrt{3}\langle\mu\rangle\left(\frac{1-\alpha}{1+\alpha}\right) + 1\right)\delta\psi^{-}, \tag{A.38d}$$

$$\delta\psi^{-}_{\text{in}} = \frac{1 - 2\sqrt{3}\,\langle\mu\rangle\left(\frac{1-\alpha}{1+\alpha}\right)}{1 + 2\sqrt{3}\,\langle\mu\rangle\left(\frac{1-\alpha}{1+\alpha}\right)}\,\delta\psi^{+}_{\text{in}} + \frac{4\sqrt{3}\left(\frac{1}{1+\alpha}\right)}{1 + 2\sqrt{3}\,\langle\mu\rangle\left(\frac{1-\alpha}{1+\alpha}\right)} J_{\text{in}}. \tag{A.38e}$$

This BC preserves the partial currents but couples the ψ^{+} and ψ^{-}. Note that α is the true albedo used in the transport solution. In the k-eigenvalue problems studied, α was unity.

Once the S2SA matrix was inverted, the scalar flux it predicted was added to the current value of the scalar flux. To preserve balance, the outward angular fluxes on the right boundary were updated as well.

APPENDIX B

SOLUTION TECHNIQUES FOR THE TRANSPORT EQUATION

B.1 Traditional Schemes

We investigate the convergence criteria used when solving the transport equation. As a reminder, the traditional solution methodology is nested. Going from outside to inside the nesting, the solves look like: convergence of the fission source (and criticality eigenvalue) through power iteration, convergence of the across-group scattering source through Gauß-Seidel (GS), convergence of the within-group scattering source through source iteration (SI), and inversion of the streaming operator through sweeps over angle and space. Since the convergence of across-group scattering is a sweep over groups, the solution technique looks like nested sweeps. Going from outside to inside the nesting, these sweeps are over groups, levels, angles within a level, and space (in the direction of the angle).

B.1.1 Convergence of the fission source through power iteration (cont.)

The convergence criteria used were a set of relative L^∞ norms. As, for many problems, the eigenvector converges more slowly than the eigenvalue [2, 52], both the eigenvalue and eigenvector separately were checked for convergence. The convergence criteria were

$$\left| \frac{{k_{\text{eff}}}^z - {k_{\text{eff}}}^{z-1}}{{k_{\text{eff}}}^z} \right| < \frac{\epsilon_k}{1-\rho_k}, \tag{B.1a}$$

$$\left\| \frac{\phi_0^z - \phi_0^{z-1}}{\phi_0^z + \delta} \right\|_{L^\infty} < \frac{\epsilon_\phi}{1-\rho_\phi}. \tag{B.1b}$$

where the ρ_x are the estimated spectral radii used to prevent false convergence, the ϵ_x are the relative error criteria given by the user, ϕ_0 is the scalar flux, and δ is a smaller number ($\sim 10^{-12}$) used to prevent division by zero.

False convergence occurs when the difference between two successive iterates is small but the solution is not near the true / converged solution. This situation occurs when the method does not decrease the error in the solution much between iterates.

Define the error in k_{eff} at the z^{th} iteration as

$$e_k^z = {k_{\text{eff}}}^z - k_{\text{eff}}, \tag{B.2a}$$

where k_{eff} is the true eigenvalue and ${k_{\text{eff}}}^z$ is the approximation at the z^{th} iteration.

If it is assumed that each iteration decreases this error by a constant factor, ρ, then, for large z, the error has the form

$$e_k^z = c_0\,\rho^z. \tag{B.2b}$$

Some algebra yields

$$\frac{{k_{\text{eff}}}^z - {k_{\text{eff}}}^{z-1}}{{k_{\text{eff}}}^{z-1} - {k_{\text{eff}}}^{z-2}} = \frac{e_k^z - e_k^{z-1}}{e_k^{z-1} - e_k^{z-2}} \tag{B.2c}$$

$$= \frac{c_0\,\rho^z - c_0\,\rho^{z-1}}{c_0\,\rho^{z-1} - c_0\,\rho^{z-2}} \tag{B.2d}$$

$$= \frac{\rho^{z-1}\,(\rho - 1)}{\rho^{z-2}\,(\rho - 1)} \tag{B.2e}$$

$$\frac{{k_{\text{eff}}}^z - {k_{\text{eff}}}^{z-1}}{{k_{\text{eff}}}^{z-1} - {k_{\text{eff}}}^{z-2}} = \rho. \tag{B.2f}$$

Applying the same ansatz to the L^2 error norm of the eigenvector equation yields

$$\rho_k = \frac{{k_{\text{eff}}}^z - {k_{\text{eff}}}^{z-1}}{{k_{\text{eff}}}^{z-1} - {k_{\text{eff}}}^{z-2}}, \tag{B.3a}$$

$$\rho_\phi = \frac{\left\|\phi_0^z - \phi_0^{z-1}\right\|_{L^2}}{\left\|\phi_0^{z-1} - \phi_0^{z-2}\right\|_{L^2}}. \tag{B.3b}$$

The L^2 error norm was chosen because it is less sensitive to noise than the L^∞ norm is; thus, its estimate for ρ_ϕ is expected to converge more quickly and be stable. The L^1 and L^2 norms are both appropriate for this role.

Note that Eq. (B.3b) does not require the storage of ϕ_0^{z-2}. Instead, the numerator and denominator of Eq. (B.3b) should be separately stored. When the iteration index increments, the new value of the denominator will be the old value of the numerator.

When the scattering source is fully converged (i.e., the inversion in Eq. (5.3c) is done to a tight tolerance), the ρ_k computed from Eq. (B.3a) is equal to the dominance ratio.

B.1.2 Convergence of the across-group scattering source through GS (cont.)

The convergence criteria for the across-group scattering source iteration of upscattering is analogous to that used in power iteration. Scalar flux is the variable of interest, since it is used to generate reaction rates and equivalent-critical flux profiles. Instead of looking at a single group, the criteria are applied over all of the upscattering region.

$$\max_{g\in\text{upscat}} \left\| \frac{\phi_{g,0}^u - \phi_{g,0}^{u-1}}{\phi_{g,0}^u + \delta} \right\|_{L^\infty} < \frac{\epsilon_\phi}{1-\rho_\phi}, \tag{B.4a}$$

$$\rho_\phi = \left[\max_{g\in\text{upscat}} \frac{\left\|\phi_{g,0}^u - \phi_{g,0}^{u-1}\right\|_{L^1}}{\left\|\phi_{g,0}^{u-1} - \phi_{g,0}^{u-2}\right\|_{L^1}} \right]. \tag{B.4b}$$

As the upscattering can be converged in the outer iterations, ϵ_ϕ may be chosen to be large or the maximum number of upscatter iterations may be chosen to be small.

B.1.3 Convergence of the within-group scattering source through SI (cont.)

The convergence criteria for the within-group scattering source iteration is analogous to that used in power iteration. Scalar flux is the variable of interest, since it is used to generate reaction rates and equivalent-critical flux profiles. The iteration is terminated when

$$\left\| \frac{\phi_0^l - \phi_0^{l-1}}{\phi_0^l + \delta} \right\|_{L^\infty} < \frac{\epsilon_\phi}{1 - \rho_\phi}, \tag{B.5a}$$

$$\rho_\phi = \frac{\left\| \phi_0^l - \phi_0^{l-1} \right\|_{L^1}}{\left\| \phi_0^{l-1} - \phi_0^{l-2} \right\|_{L^1}}. \tag{B.5b}$$

The ϵ_ϕ above should start larger than the one for power iteration and decrease as the fission source is converged.

B.1.4 Inversion of the streaming operator through sweeps over angle and space (cont.)

B.1.4.1 The Starting-Direction Sweeps

The derivative in angle is over the azimuthal component, ω. Each level of the quadrature is decoupled in this way (Fig. A.1). For each level, a starting-direction sweep must be done to generate the BC for the angular derivative. The definition of direction in cylindrical coordinates makes the angular advection proceed from negative ω to positive ω. Thus, the starting-direction sweep should be at the negative-most ω for a given level (given l; given ξ). That is, the starting-direction sweep has $\omega = -\pi$ and $\xi = \xi_l$. This defines a unique $\mathbf{\Omega}$. The sweep hierarchy, from outside to inside, is over levels, then angles within the level, then space in the direction of the

angle.

Since all of the starting-direction sweeps have $\omega < 0$, the inner spatial sweeps proceed from $r = R$ to $r = 0$. The BC used is the extrapolated angular flux from the first two negative-ω directions on the level; for the first level, there is only one negative-ω direction in the quadrature, so its angular flux is used as the BC. The cells are swept inward, with Eq. (A.21b) used to determine the angular fluxes in a cell given the upwinded flux. The scattering source for the starting-direction sweeps uses the starting-direction moment-to-discrete matrix discussed earlier. The angular fluxes are stored to provide the starting condition, $\psi_{k-1/2}$, for subsequent sweeps. The angular flux at the center ($r = 0$) is stored, as that is the proper boundary condition for subsequent sweeps on the level with $\mu_k > 0$.

B.1.4.2 Subsequent Sweeps

An angular sweep is done over all angles in a level. A spatial sweep is done for each angle. For a given direction, the cells are swept in that direction, using Eq. (A.20f) to solve for the angular flux within a cell given the spatially upwinded flux from the previous cell and given the angularly "upwinded" flux from the previous direction (or the starting direction). The normal moment-to-discrete matrix is used to set the scattering source. The moments of the angular flux are stored, as only they are needed to define the scattering source, which is being iterated upon. Additionally, the WDD in angle relationship (Eq. (A.7a)) is used to define the angular source for the next direction, which is also stored.

Special care is needed on the boundaries. For negative μ, the center-cell balance equation (Eq. (A.24d)) is solved at the center instead of Eq. (A.20f). For positive μ, the angular flux on the outer boundary is kept to allow for implicit boundary conditions to be used and for surface partial currents to be calculated.

B.2 NKA

In contrast to the traditional solution technique, we introduce details for nonlinear Krylov acceleration (NKA), which breaks nesting by doing only one sweep at each level instead of iterating to convergence. Assuming a sufficiently accurate starting guess for eigenvector and eigenvalue, NKA may save computation time by not expending effort converging inner iterations when the system is relatively far from being converged.

B.2.1 Krylov methods

Krylov methods are a mathematical class of iterative methods that solve a linear system using subspaces. The linear problem,

$$\mathcal{A}\mathbf{x} = \mathbf{b}, \tag{B.6a}$$

may be recast using an approximate solution x_0:

$$\mathcal{A}\mathbf{x} - \mathbf{b} = \mathcal{A}\mathbf{x}_0 - \mathcal{A}\mathbf{x}_0, \tag{B.6b}$$

$$\mathcal{A}\left(\mathbf{x} - \mathbf{x}_0\right) = \mathbf{b} - \mathcal{A}\mathbf{x}_0, \tag{B.6c}$$

$$\mathcal{A}\mathbf{e} = \mathbf{r}, \tag{B.6d}$$

with

$$\mathbf{e} = \mathbf{x} - \mathbf{x}_0, \tag{B.6e}$$

$$\mathbf{r} = \mathbf{b} - \mathcal{A}\mathbf{x}_0. \tag{B.6f}$$

Here $\mathbf{r}$ is the residual and $\mathbf{e}$ is the error.

The Krylov subspace is the vector space consisting of the matrix $\mathcal{A}$ applied to the residual $\mathbf{r}$:

$$\mathcal{A} \in \mathbb{R}^{N,N}, \tag{B.7a}$$

$$\mathbb{K}_n = \left\{ \mathbf{r} \otimes \mathcal{A}\mathbf{r} \otimes \cdots \otimes \mathcal{A}^{n-1}\mathbf{r} \right\} \qquad n \in 1, \ldots, N. \tag{B.7b}$$

Once a subspace of size $n = N$ is achieved, the solution to the problem is within the subspace and the system is solved (see any graduate-level textbook on Krylov methods; this is proved for the GMRES routine by Saad and Schultz [53]). Normally this never occurs; instead, the problem is somehow restarted at the new estimate for x every $M \ll N$ iterations. This allows the subspace, which must be stored in memory, to be forgotten and keeps memory costs manageable. Further, finite machine precision leads to rounding errors that plague the normalization procedure generally applied to the subspace. For these reasons, more than N iterations may be required for the exact solution. More importantly, the method is not guaranteed to converge if restarting or finite precision arithmetic are used.

Using Krylov spaces has become increasingly popular recently, especially applied to nuclear engineering problems. Krylov methods may be wrapped around existing code to provide both linear [47] and nonlinear [52, 54, 55] acceleration.

A chief advantage of Krylov spaces is that only the action of the matrix is required, not the matrix itself. This saves storage costs and allows Krylov methods to be used on sweeps, where generally only the action of the matrix (inverse) is known or computed in the code.

B.2.2 Anderson Mixing method (cont.)

The NKA implementation of the Anderson Mixing Method follows for the n^{th} iteration. All vectors live in $\mathbb{R}^N$ and M states are saved. (The quantity $i = n-M+1$ is more formally $i = \max(0, n - M + 1)$.)

1. Given a solution vector $\mathbf{x}_n$ (here an eigenvector, eigenvalue concatenation), compute the residual $\mathbf{f}(\mathbf{x}_n)$. This involves doing one space-angle sweep for each group.

2. From previous iterations, $\mathbf{x}_n$, $\{\mathbf{v}_i\}_{i=n-M+1}^{i=n}$, $\{\mathbf{w}_i\}_{i=n-M+1}^{i=n-1}$, $\mathbf{f}(\mathbf{x}_{n-1})$, and now $\mathbf{f}(\mathbf{x}_n)$ are known. Compute

$$\mathbf{w_n} = \mathbf{f}(\mathbf{x}_{n-1}) - \mathbf{f}(\mathbf{x}_n). \tag{B.8a}$$

3. Compute

$$\mathbf{v}_{n+1} = \left(\sum_{i=n-M+1}^{n} z_i \, \mathbf{v}_i \right) + \left(\mathbf{f}(\mathbf{x}_n) - \sum_{i=n-M+1}^{n} z_i \, \mathbf{w}_i \right). \tag{B.8b}$$

   Determine the z_i by minimizing a norm of

$$\mathbf{f}(\mathbf{x}_n) \; - \sum_{i=n-M+1}^{n} z_i \, \mathbf{w}_i. \tag{B.8c}$$

4. Update the solution vector by

$$\mathbf{x}_{n+1} = \mathbf{x}_n - \mathbf{v}_{n+1}. \tag{B.8d}$$

5. Check convergence and do memory management. Remove the oldest state from

$\{\mathbf{v}_i\}$ and $\{\mathbf{w}_i\}$; fill in the hole in $\{\mathbf{w}_i\}$ with $\mathbf{f}(\mathbf{x}_n)$ so the $\mathbf{f}(\mathbf{x}_{n-1})$ mentioned in step 2 does not need to be stored explicitly.

6. If not converged, go to step 1.

The NKA algorithm with pointwise convergence checking requires storing $2M+3$ copies of the R^N solution vector. Both $\{\mathbf{v}_i\}$ and $\{\mathbf{w}_i\}$ must be kept for each saved state; $\mathbf{f}(\mathbf{x}_n)$, $\mathbf{x}_n$ and $\mathbf{x}_{n-1}$ must also be stored. Comparing $\mathbf{x}_n$ and $\mathbf{x}_{n-1}$ constitutes the pointwise convergence check. This is twice as many state vectors are as required for standard Krylov solvers, like GMRES (e.g., when used in JFNK), meaning we come up against memory limitations approximately twice as fast[1].

The NKA code used for this work was provided in the paper by Calef et al., and may be found at `http://sourceforge.net/projects/nlkain/`. It uses an L^2 norm in step 3 and a Cholesky factorization to solve the resulting least-squares problem.

Note that the two time-consuming parts of the algorithm are steps 2 (the sweep) and 3 (the minimization problem). More preconditioning than a simple sweep may be used when computing the residual, including S2SA and/or multiple sweeps.

B.2.3 Simultaneous convergence of the fission and scattering sources through NKA (cont.)

The theme of Krylov is that a small residual implies a small error. For linear systems,

$$\mathbf{r} \to 0 \Leftrightarrow \mathbf{e} \to 0. \tag{B.9}$$

While this is not true for matrices of large condition number — the error for linear systems can be shown to be bounded by the residual times the condition number

[1] Cf. flexible GMRES, which allows for changing preconditioners, which also stores order $2M$ solution vectors.

— it does provide a convergence criterion. An L^2 norm was used on the residual to check convergence.

False convergence may occur when the residual is small but the error is not. To guard against this persistent problem, additional criteria are required. For this system, the change in eigenvalue and eigenvector between NKA iterations were also checked:

$$\left| \frac{{k_{\text{eff}}}^z - {k_{\text{eff}}}^{z-1}}{{k_{\text{eff}}}^z} \right| < \epsilon_k, \tag{B.10a}$$

$$\left\| \frac{\phi_0^z - \phi_0^{z-1}}{\phi_0^z + \delta} \right\|_{L^\infty} < \epsilon_\phi. \tag{B.10b}$$

Note that no spectral radii were used in the above equations. The NKA method is not a linear method and is not even guaranteed to decrease the error every iteration, to say nothing of decreasing it by a constant factor. The idea of a spectral radius has no meaning for such schemes and can easily be negative or larger than unity.

B.3 SCDT (cont.)

B.3.1 Nomenclature

SCDT uses nomenclature similar to that PDT, the *Parallel Deterministic Transport* code developed by Dr. Adams and others at Texas A&M University [56]. The spatial domain is split up into divisions, cells and elements. A division has a consistent grid-spacing scheme (either linear or logarithmic), one material type and one distributed source type (for source-driven problems). A division consists of cells, which have a beginning and end radius, and an LD (linear discontinuous) FE representation in space. For LD, a cell has two degrees of freedom, called elements. These elements house the unknowns of scalar flux moments, new and old.

Cross sections are defined as constant within materials which are themselves

mixtures of components. Each component is a set of microscopic cross sections with major units of barns. The components are summed by atom fraction and multiplied by a density in atoms per barn-centimeter to yield a macroscopic cross section for each material. Divisions reference materials for their cross-section lookups. The macroscopic scattering transfer matrix and total cross section are computed as described above. The fission cross sections are stored as microscopic sets of χ and $\nu\sigma_f$ for each component in lieu of storing the entire fission transfer matrix. This lower-dimensional storage scheme preserves the differences in χ among disparate nuclides while minimizing memory overhead by storing the fission kernel as dyads.

B.3.2 Balance

An important aspect of code verification is the notion of particle — or in the case of radiative transfer, energy — balance, which characterizes the degrees to which particles are (energy is) preserved in the discrete system.

Balance is defined as the relative difference between sources and sinks integrated out over all phase space.

$$\text{sources} = \text{inflow} + \text{distributed sources}, \tag{B.11a}$$

$$\text{sinks} = \text{outflow} + \text{net absorption}, \tag{B.11b}$$

$$\text{balance} = \frac{\text{sources} - \text{sinks}}{\text{sources}}. \tag{B.11c}$$

Balance may be defined individually for each group (within-group balance) or summed over all groups (across-group balance). If summed over all groups, the inflow and outflow refer to currents across the spatial boundary of the problem. For per-group balance, inflow includes in-scattering and outflow includes the transfer rate through total interaction.

Balance is an important diagnostic for the code. The schemes used should conserve total particles (over all phase-space) upon convergence. Therefore, balance should approach 0 (to machine precision) as the scattering source is converged and should be exact when there is no scattering (and explicit BC). When S2SA is used on a fixed-source problem with no upscatter, it can be shown that balance should be exact for each iteration. If balance is not achieved, this signifies either a bug in the code or inexact angular quadrature in the form of moment-to-discrete and discrete-to-moment matrices that do not have the desired properties of row / column sums or inverse relationships.

B.3.3 Within-group balance

The within-group balance is given by

$$\text{sources}_g = J^-_{\text{in},g}(R) + \sum_{g' \neq g} \int_0^R \mathrm{d}r \;\; r \; \Sigma_{s,g' \to g} \; \phi^0_{0,g'}(r) + \int_0^R \mathrm{d}r \;\; r \; q_g(r), \tag{B.12a}$$

$$\text{sinks}_g = J^+_{\text{out},g}(R) + \int_0^R \mathrm{d}r \;\; r \; \Sigma_{tr,g}(r) \, \phi^0_{0,g}(r), \tag{B.12b}$$

where the transfer cross section for group g is defined as $\Sigma_{tr,g} = \Sigma_{t,g} - \Sigma_{s,g \to g}$. Within-group balance will be met when the within-group (within-band) scattering source is converged.

The fission source is included as part of the distributed sources. As the fission source is converged, balance increases. For the k-eigenvalue problem, the distributed source is

$$q_g(r) = \frac{1}{k_{\text{eff}}} \sum_{i \in \text{comp}} \chi_{i,g}(r) \; n_i(r) \sum_{g'=1}^{G} \nu\sigma_{f,i,g'}(r) \; \phi^0_{0,g'}(r), \tag{B.12c}$$

where the first sum is over components in a material.

B.3.4 Across-group balance

The across-group balance is given by

$$\text{sources} = \sum_{g=1}^{G} \left\{ J^{-}_{\text{in},g}(R) + \int_0^R \mathrm{d}r \;\; r \; q_g(r) \right\}, \tag{B.13a}$$

$$\text{sinks} = J^{+}_{\text{out},g}(R) + \int_0^R \mathrm{d}r \;\; r \; \Sigma_{a,g}(r) \, \phi^0_{0,g}(r), \tag{B.13b}$$

where the net absorption for group g is defined as $\Sigma_{a,g} = \Sigma_{t,g} - \sum_{\text{all } g'} \Sigma_{s,g\to g'}$. Note that if reactions like $(n, 2n)$ are included in $\Sigma_{s,g\to g'}$, then $\Sigma_{a,g}$ may be negative and is thus best known as the net absorption.

The fission source is included as part of the distributed sources. As the fission source is converged, balance increases. For the k-eigenvalue problem, using the normalization of the χ_g, the distributed source may be simplified to

$$\sum_{g=1}^{G} q_g(r) = \frac{1}{k_{\text{eff}}} \sum_{g'=1}^{G} \nu\Sigma_{f,g'}(r) \; \phi^0_{0,g'}(r). \tag{B.13c}$$

B.3.5 S2SA solver

To solve the S2SA system, in each cell the unknowns were ordered as

$$\begin{pmatrix} \delta\psi^{+}_{L} \\ \\ \delta\psi^{+}_{R} \\ \\ \delta\psi^{-}_{L} \\ \\ \delta\psi^{-}_{R} \end{pmatrix}. \tag{B.14}$$

When upwinding was used on the streaming terms, inter-cell coupling was introduced between the row corresponding to $\psi^{+}_{i,L}$ and the column corresponding to $\psi^{+}_{i-1,R}$ for the positive directions; for the negative directions, inter-cell coupling was introduced between the row corresponding to $\psi^{-}_{i,R}$ and the column corresponding to $\psi^{-}_{i+1,L}$. With this coupling, the S2SA matrix became 7-banded. The matrix was stored in Compressed Diagonal Storage (CDS) sparse format. A direct LU solver with pivoting was used to invert the sparse matrix. This solver required 10-bands of storage per row for the pivoting.

APPENDIX C

RESULTS ADDENDA

C.1 Problem 2 (cont.)

C.1.1 Reference-calculation weighting

C.1.2 Resonance-only resolution study

C.1.2.1 Resolution 1 with 2 bands per group

Results for an energy grid of resolution factor 1. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 2 bands per group were used.

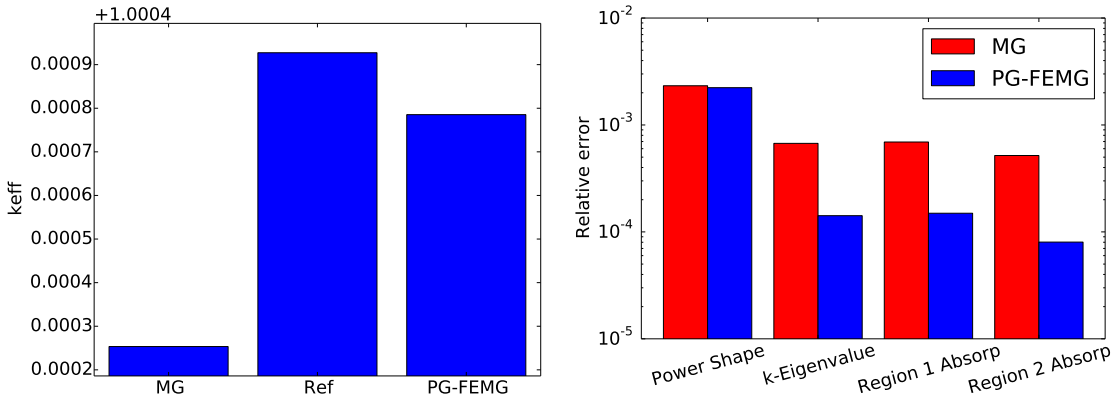

(a) k-Eigenvalues

(b) Relative Errors

Figure C.1: Comparison of eigenvalues and relative errors for problem 2 with high-fidelity weight for resolution 1 and 2 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

Fig. C.1 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.1a gives the k-eigenvalues for each of the calculations. Fig. C.1b gives the error in k-eigenvalue, power shape, and absorption rates.

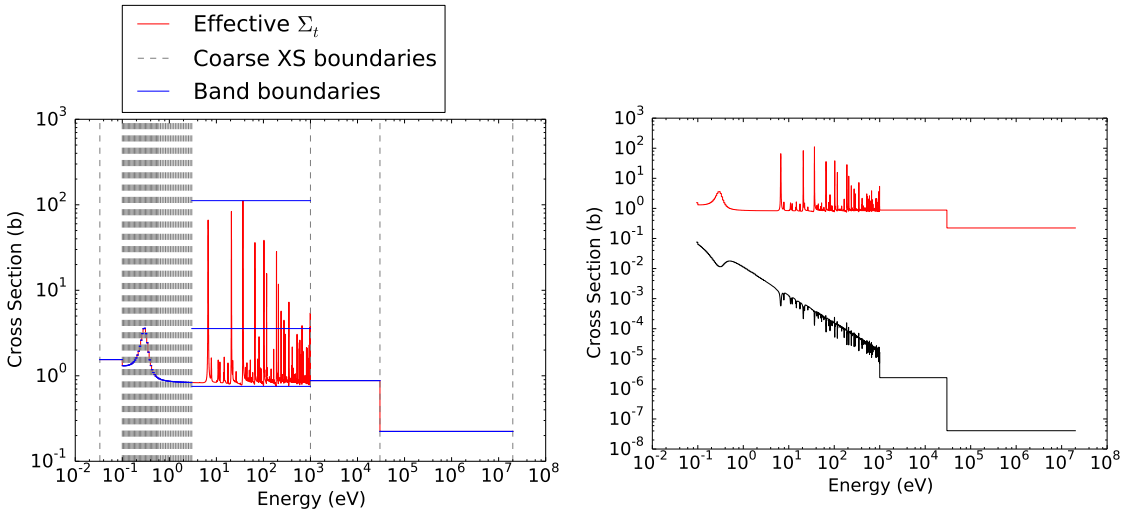

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.2: Effective total cross sections for the PG-FEMG band boundary calculation for problem 2 with high-fidelity weight for resolution 1 and 2 bands per group in the resonance region.

Fig. C.2a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.2b gives the effective total cross section and the weighting spectrum used. For this calculation, a high-fidelity weighting spectrum was used (note resonances in b).

C.1.2.2 Resolution 4 with 4 bands per group

Results for an energy grid of resolution factor 4. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution

factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 4 bands per group were used.

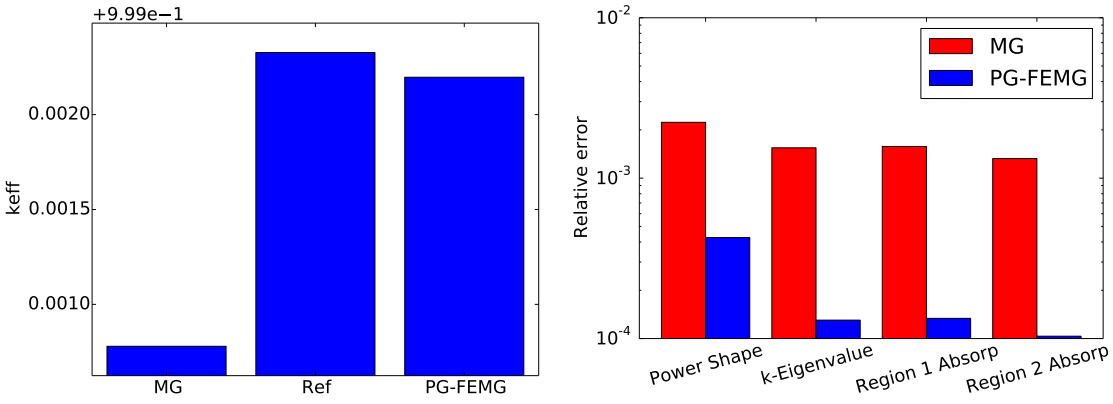

(a) k-Eigenvalues (b) Relative Errors

Figure C.3: Comparison of eigenvalues and relative errors for problem 2 with high-fidelity weight for resolution 4 and 4 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

Fig. C.3 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.3a gives the k-eigenvalues for each of the calculations. Fig. C.3b gives the error in k-eigenvalue, power shape, and absorption rates.

Fig. C.4a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.4b gives the effective total cross section and the weighting spectrum used. For this calculation, a high-fidelity weighting spectrum was used (note resonances in b).

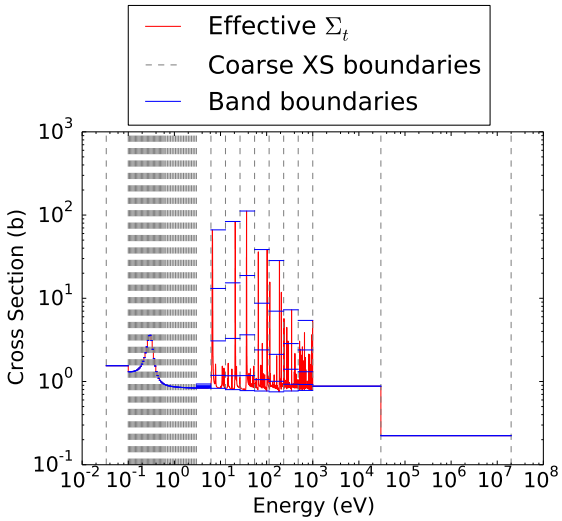

(a) Effective total cross section and band boundaries.

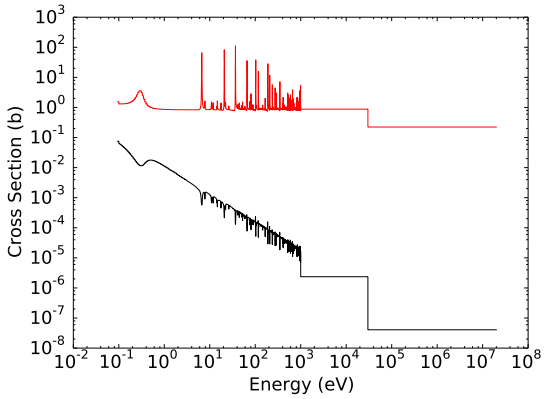

(b) Effective total cross section and the weighting spectrum.

Figure C.4: Effective total cross sections for the PG-FEMG band boundary calculation for problem 2 with high-fidelity weight for resolution 4 and 4 bands per group in the resonance region.

C.1.3 Full resolution study

C.1.3.1 Resolution 1 with 2 bands per group

Results for an energy grid of resolution factor 1. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 2 bands per group were used.

Fig. C.5 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.5a gives the k-eigenvalues for each of the calculations. Fig. C.5b gives the error in k-eigenvalue, power shape, and absorption rates.

Fig. C.6a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.6b

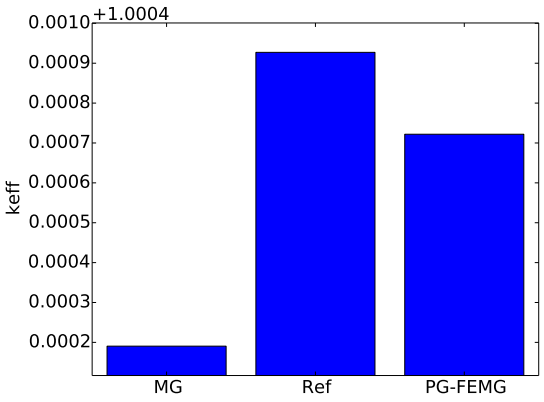

(a) k-Eigenvalues

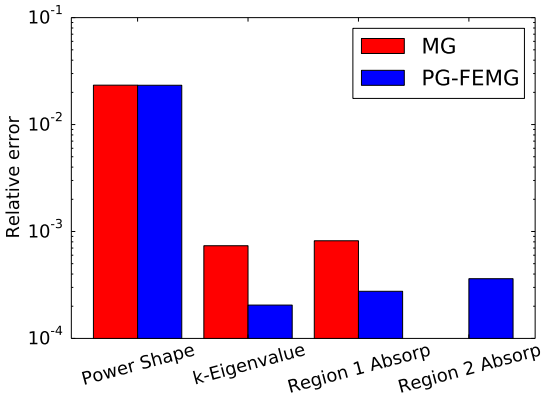

(b) Relative Errors

Figure C.5: Comparison of eigenvalues and relative errors for problem 2 with high-fidelity weight for resolution 1 and 2 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

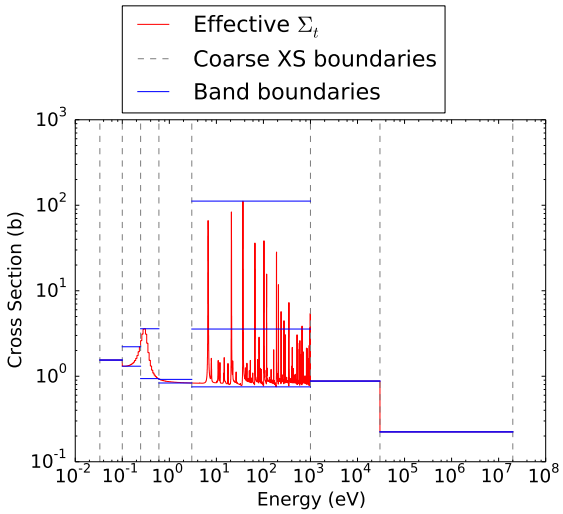

(a) Effective total cross section and band boundaries.

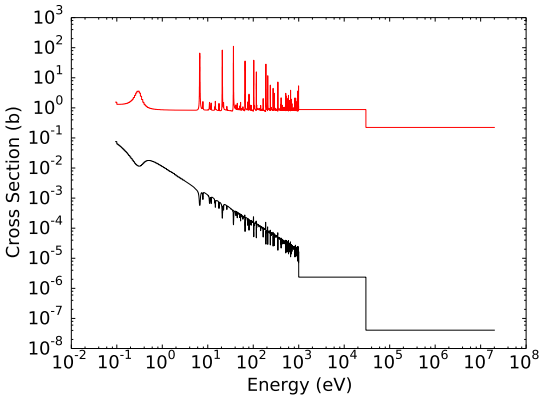

(b) Effective total cross section and the weighting spectrum.

Figure C.6: Effective total cross sections for the PG-FEMG band boundary calculation for problem 2 with high-fidelity weight for resolution 1 and 2 bands per group in the resonance region.

gives the effective total cross section and the weighting spectrum used. For this calculation, a high-fidelity weighting spectrum was used (note resonances in b).

C.1.3.2 Resolution 4 with 4 bands per group

Results for an energy grid of resolution factor 4. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 4 bands per group were used.

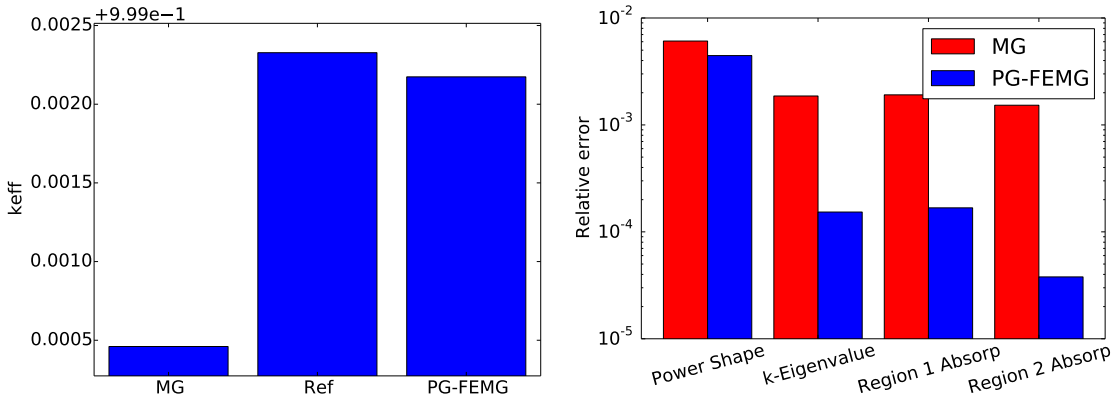

(a) k-Eigenvalues

(b) Relative Errors

Figure C.7: Comparison of eigenvalues and relative errors for problem 2 with high-fidelity weight for resolution 4 and 4 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

Fig. C.7 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.7a gives the k-eigenvalues for each of the calculations. Fig. C.7b gives the error in k-eigenvalue, power shape, and absorption rates.

Fig. C.8a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note

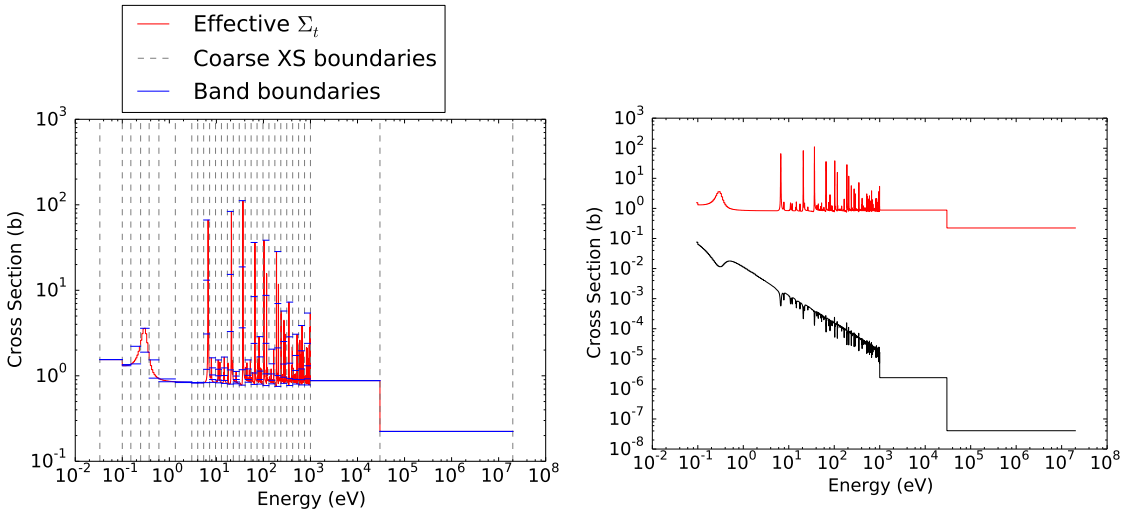

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.8: Effective total cross sections for the PG-FEMG band boundary calculation for problem 2 with high-fidelity weight for resolution 4 and 4 bands per group in the resonance region.

that multiple bands per coarse group were only used in the resonance region. Fig. C.8b gives the effective total cross section and the weighting spectrum used. For this calculation, a high-fidelity weighting spectrum was used (note resonances in b).

C.1.4 Generic weighting

C.1.5 Resonance-only resolution study

C.1.5.1 Resolution 1 with 2 bands per group

Results for an energy grid of resolution factor 1. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 2 bands per group were used.

Fig. C.9 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.9a gives the k-eigenvalues for each of the calculations.

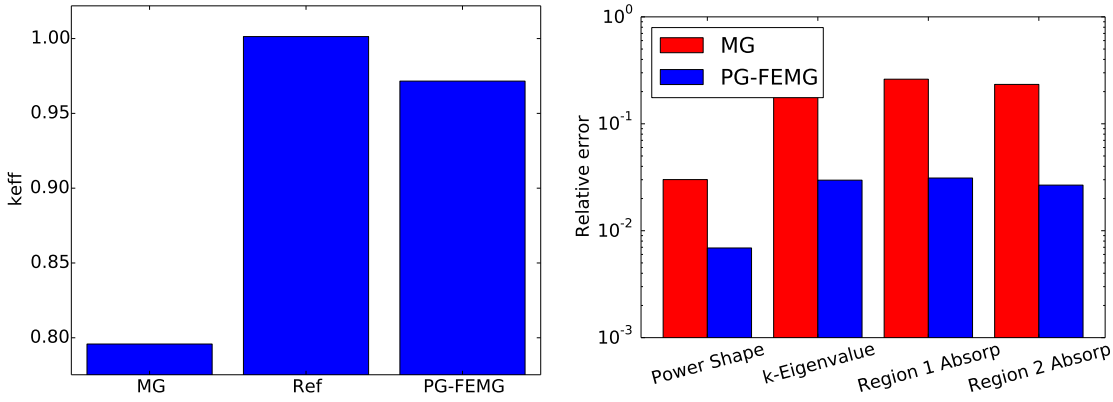

(a) k-Eigenvalues (b) Relative Errors

Figure C.9: Comparison of eigenvalues and relative errors for problem 2 with low-fidelity weight for resolution 1 and 2 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

Fig. C.9b gives the error in k-eigenvalue, power shape, and absorption rates.

Fig. C.10a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.10b gives the effective total cross section and the weighting spectrum used. For this calculation, a low-fidelity weighting spectrum was used (note smoothness in b).

C.1.5.2 Resolution 4 with 4 bands per group

Results for an energy grid of resolution factor 4. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 4 bands per group were used.

Fig. C.11 provides error measurements between the reference solution and the MG

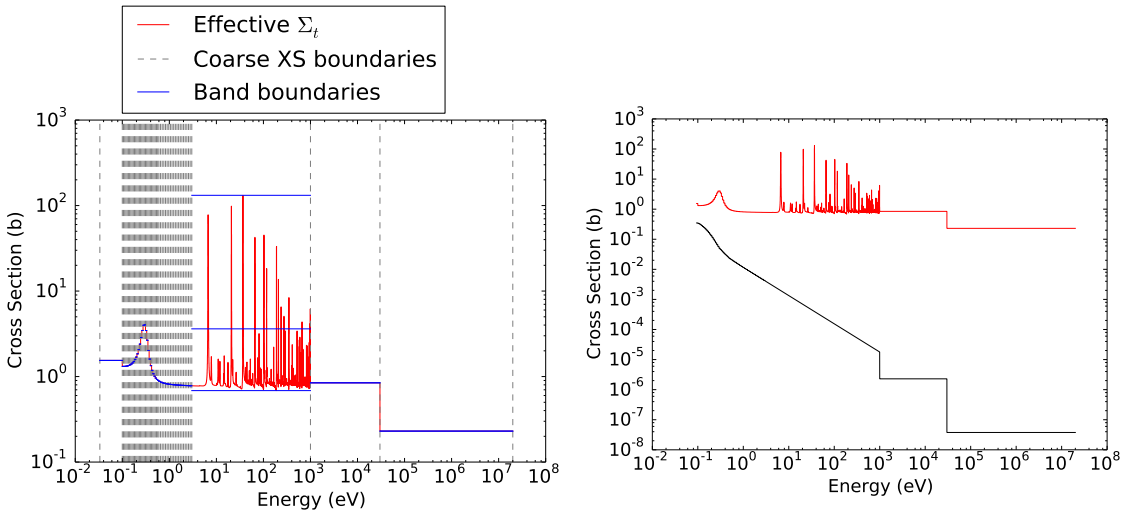

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.10: Effective total cross sections for the PG-FEMG band boundary calculation for problem 2 with low-fidelity weight for resolution 1 and 2 bands per group in the resonance region.

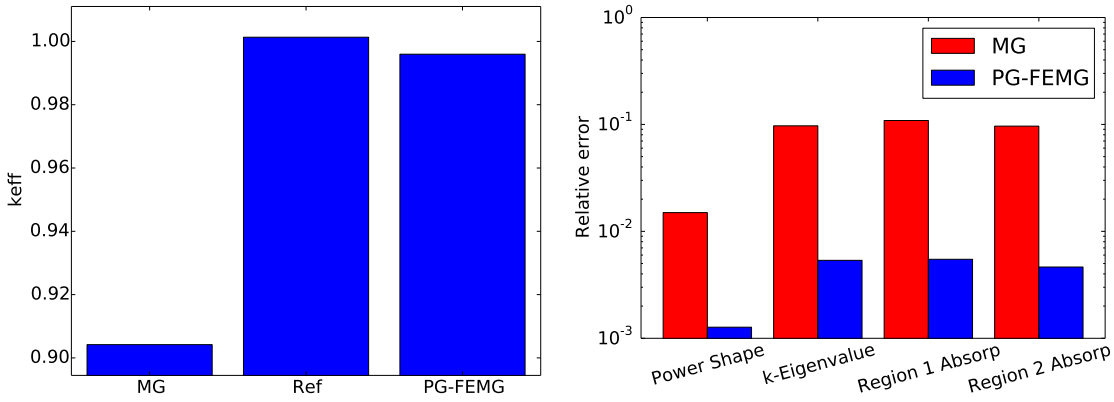

(a) k-Eigenvalues

(b) Relative Errors

Figure C.11: Comparison of eigenvalues and relative errors for problem 2 with low-fidelity weight for resolution 4 and 4 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

/ PG-FEMG solutions. Fig. C.11a gives the k-eigenvalues for each of the calculations. Fig. C.11b gives the error in k-eigenvalue, power shape, and absorption rates.

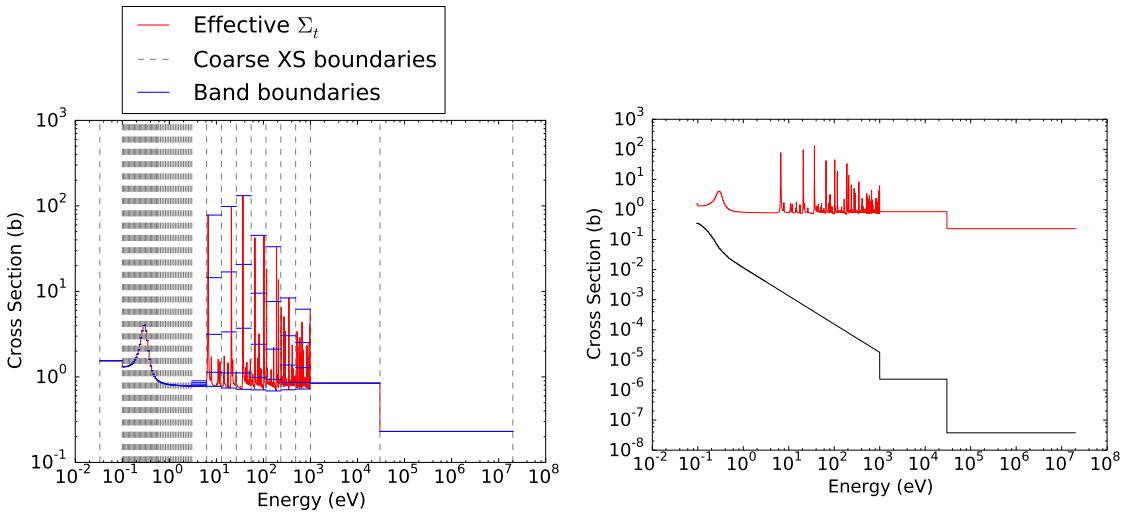

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.12: Effective total cross sections for the PG-FEMG band boundary calculation for problem 2 with low-fidelity weight for resolution 4 and 4 bands per group in the resonance region.

Fig. C.12a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.12b gives the effective total cross section and the weighting spectrum used. For this calculation, a low-fidelity weighting spectrum was used (note smoothness in b).

C.1.6 Full resolution study

C.1.6.1 Resolution 1 with 2 bands per group

Results for an energy grid of resolution factor 1. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 2 bands per group were used.

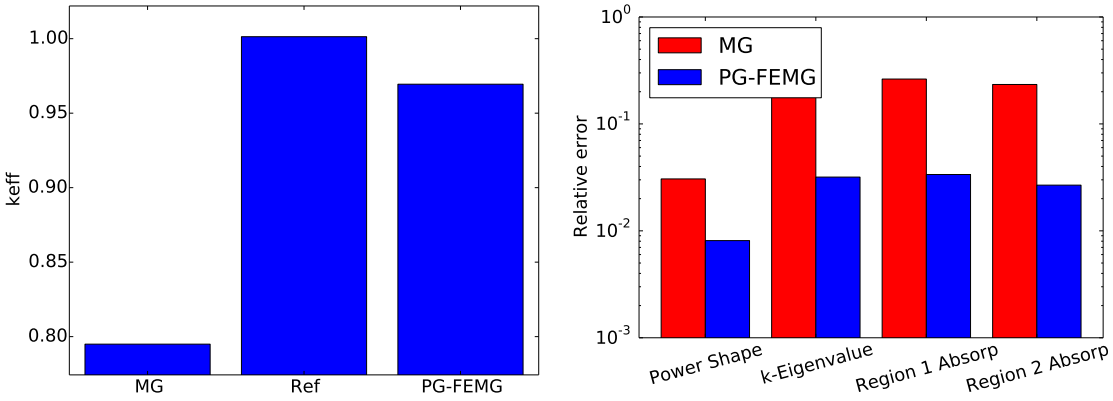

(a) k-Eigenvalues (b) Relative Errors

Figure C.13: Comparison of eigenvalues and relative errors for problem 2 with low-fidelity weight for resolution 1 and 2 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

Fig. C.13 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.13a gives the k-eigenvalues for each of the calculations. Fig. C.13b gives the error in k-eigenvalue, power shape, and absorption rates.

Fig. C.14a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region.

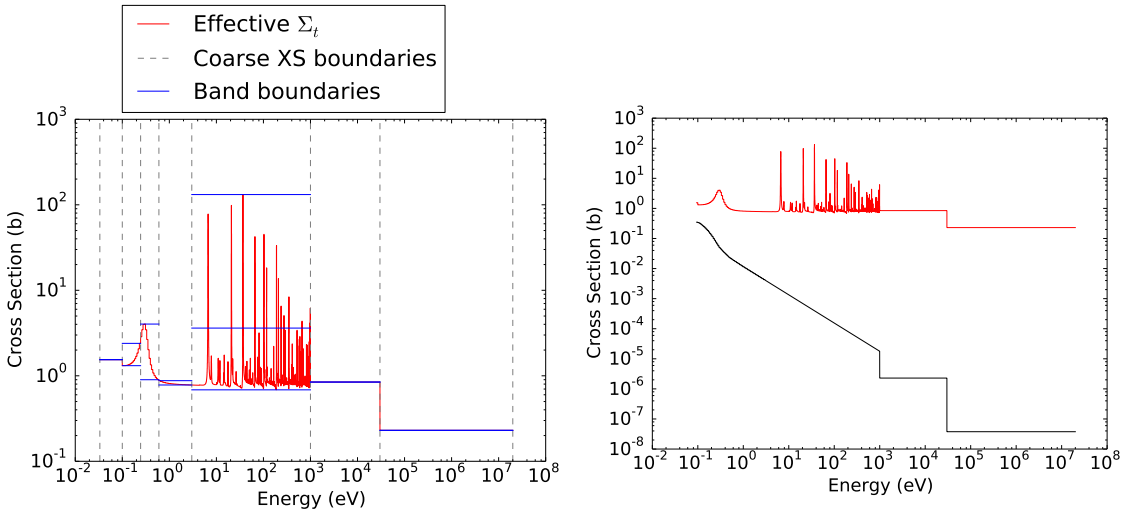

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.14: Effective total cross sections for the PG-FEMG band boundary calculation for problem 2 with low-fidelity weight for resolution 1 and 2 bands per group in the resonance region.

Fig. C.14b gives the effective total cross section and the weighting spectrum used. For this calculation, a low-fidelity weighting spectrum was used (note smoothness in b).

C.1.6.2 Resolution 4 with 4 bands per group

Results for an energy grid of resolution factor 4. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 4 bands per group were used.

Fig. C.15 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.15a gives the k-eigenvalues for each of the calculations. Fig. C.15b gives the error in k-eigenvalue, power shape, and absorption rates.

Fig. C.16a shows the effective total XS used to create the band boundaries

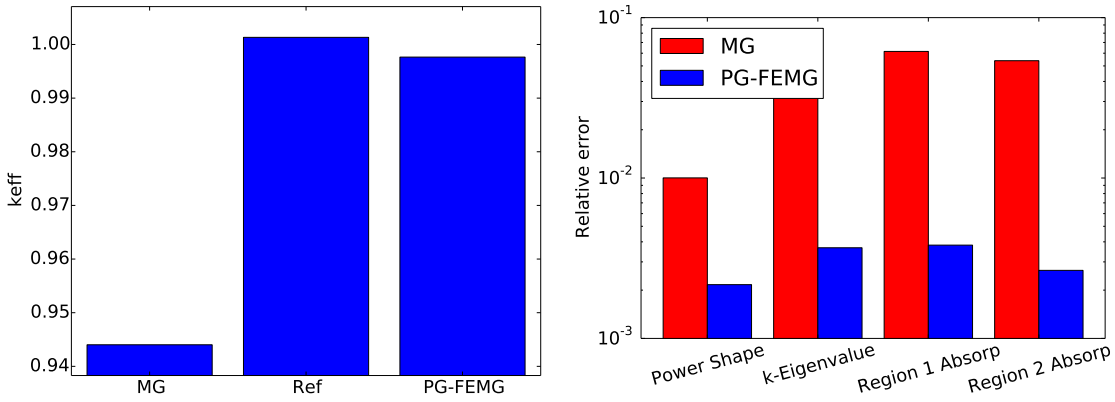

(a) k-Eigenvalues

(b) Relative Errors

Figure C.15: Comparison of eigenvalues and relative errors for problem 2 with low-fidelity weight for resolution 4 and 4 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

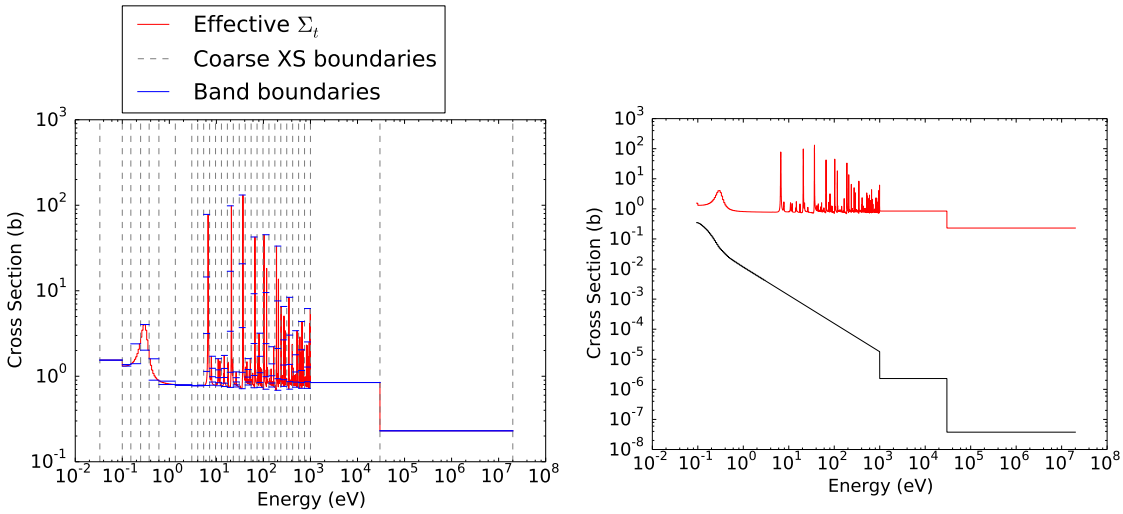

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.16: Effective total cross sections for the PG-FEMG band boundary calculation for problem 2 with low-fidelity weight for resolution 4 and 4 bands per group in the resonance region.

(red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.16b gives the effective total cross section and the weighting spectrum used. For this calculation, a low-fidelity weighting spectrum was used (note smoothness in b).

C.2 Problem 3 (cont.)

C.2.1 Reference-calculation weighting

C.2.2 Resonance-only resolution study

C.2.2.1 Resolution 1 with 2 bands per group

Results for an energy grid of resolution factor 1. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 2 bands per group were used.

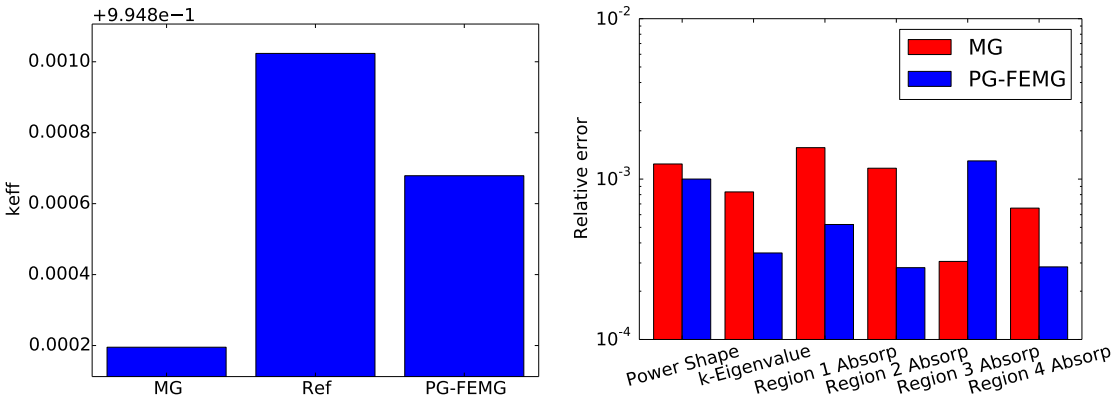

(a) k-Eigenvalues

(b) Relative Errors

Figure C.17: Comparison of eigenvalues and relative errors for problem 3 with high-fidelity weight for resolution 1 and 2 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

Fig. C.17 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.17a gives the k-eigenvalues for each of the calculations. Fig. C.17b gives the error in k-eigenvalue, power shape, and absorption rates.

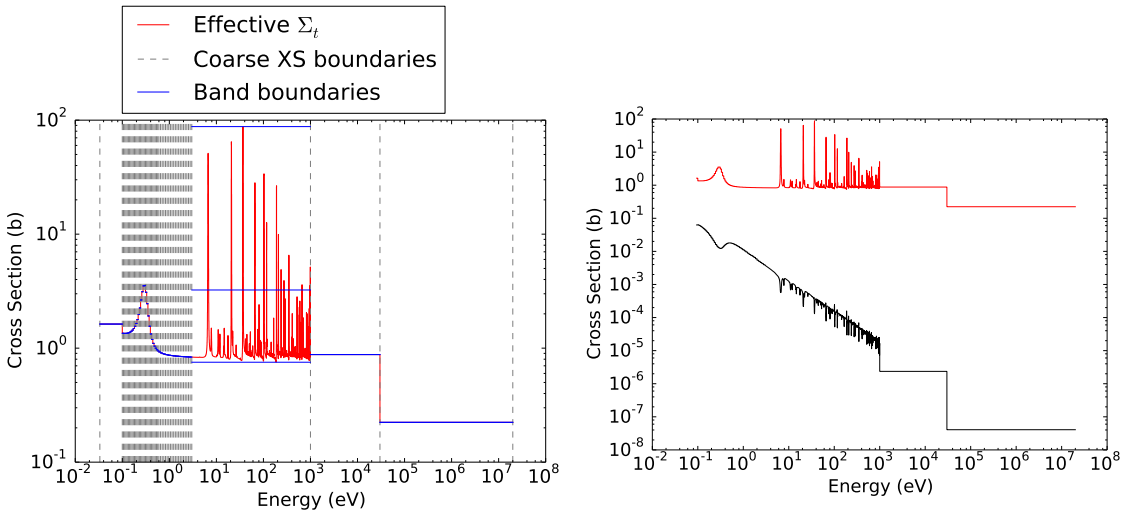

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.18: Effective total cross sections for the PG-FEMG band boundary calculation for problem 3 with high-fidelity weight for resolution 1 and 2 bands per group in the resonance region.

Fig. C.18a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.18b gives the effective total cross section and the weighting spectrum used. For this calculation, a high-fidelity weighting spectrum was used (note resonances in b).

C.2.2.2 Resolution 4 with 4 bands per group

Results for an energy grid of resolution factor 4. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 4 bands per group were used.

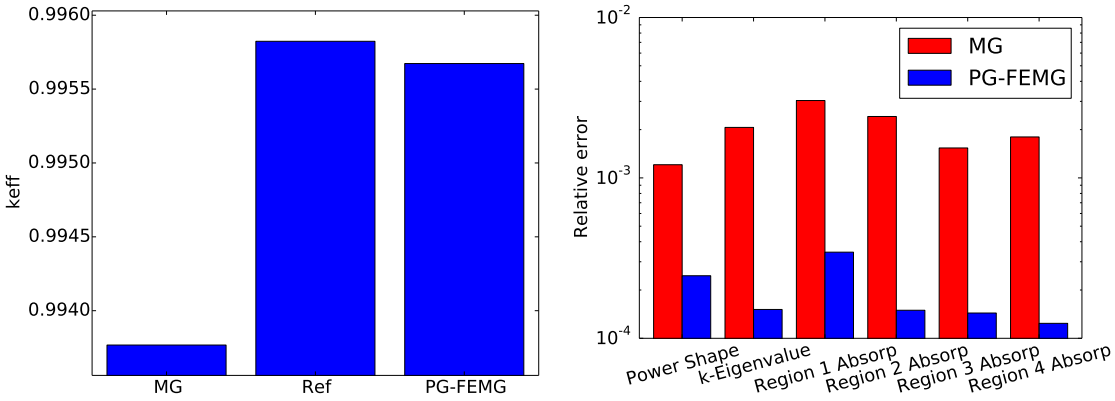

(a) k-Eigenvalues

(b) Relative Errors

Figure C.19: Comparison of eigenvalues and relative errors for problem 3 with high-fidelity weight for resolution 4 and 4 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

Fig. C.19 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.19a gives the k-eigenvalues for each of the calculations. Fig. C.19b gives the error in k-eigenvalue, power shape, and absorption rates.

Fig. C.20a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.20b gives the effective total cross section and the weighting spectrum used.

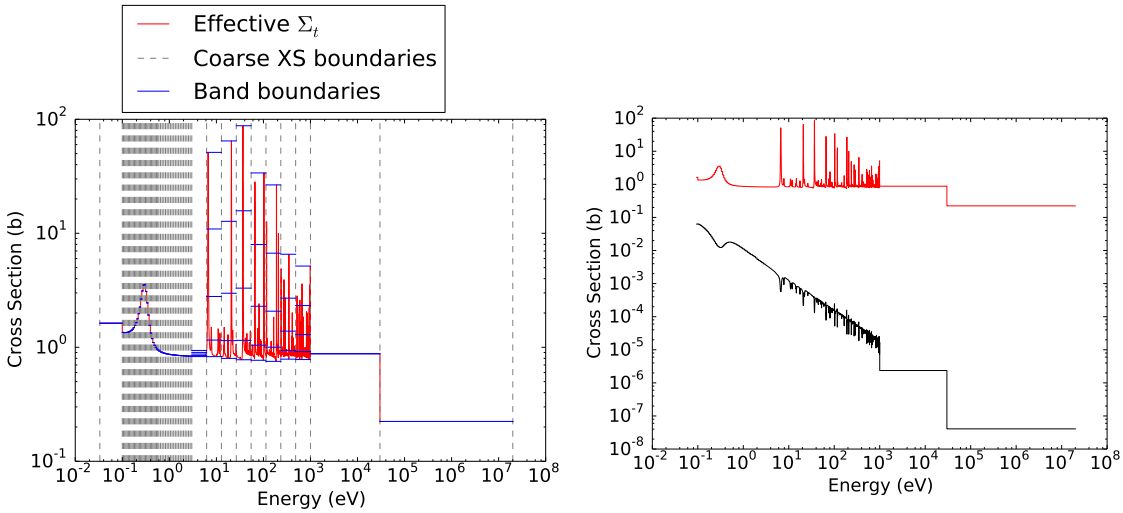

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.20: Effective total cross sections for the PG-FEMG band boundary calculation for problem 3 with high-fidelity weight for resolution 4 and 4 bands per group in the resonance region.

For this calculation, a high-fidelity weighting spectrum was used (note resonances in b).

C.2.3 Full resolution study

C.2.3.1 Resolution 1 with 2 bands per group

Results for an energy grid of resolution factor 1. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 2 bands per group were used.

Fig. C.21 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.21a gives the k-eigenvalues for each of the calculations. Fig. C.21b gives the error in k-eigenvalue, power shape, and absorption rates.

Fig. C.22a shows the effective total XS used to create the band boundaries

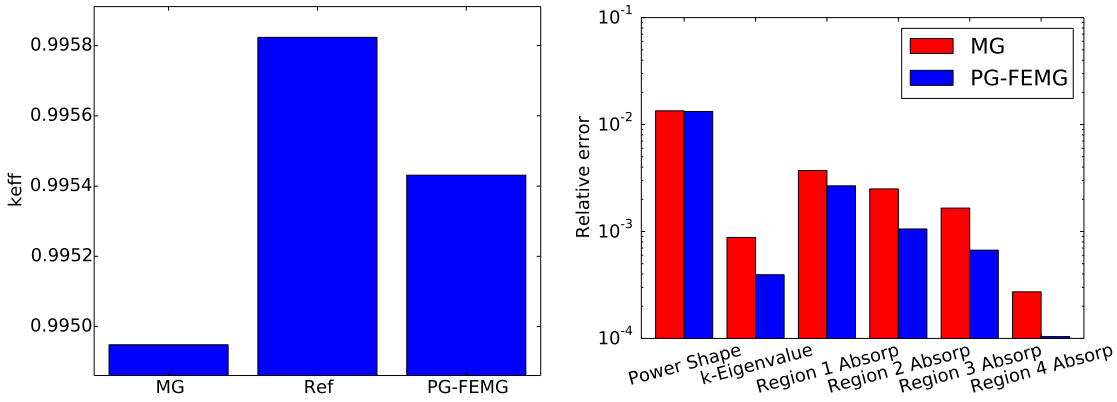

(a) k-Eigenvalues

(b) Relative Errors

Figure C.21: Comparison of eigenvalues and relative errors for problem 3 with high-fidelity weight for resolution 1 and 2 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

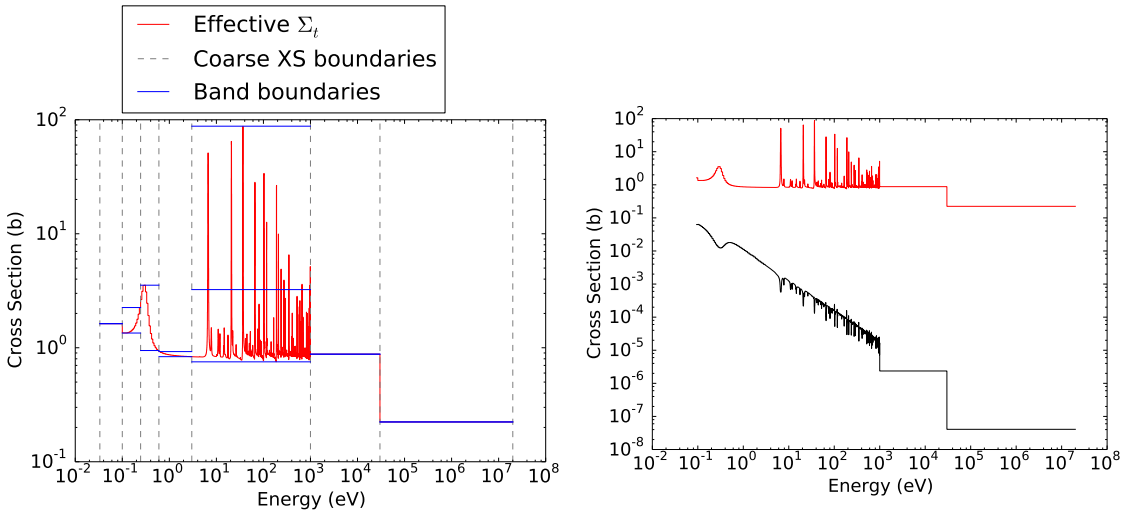

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.22: Effective total cross sections for the PG-FEMG band boundary calculation for problem 3 with high-fidelity weight for resolution 1 and 2 bands per group in the resonance region.

(red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.22b gives the effective total cross section and the weighting spectrum used. For this calculation, a high-fidelity weighting spectrum was used (note resonances in b).

C.2.3.2 Resolution 4 with 4 bands per group

Results for an energy grid of resolution factor 4. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 4 bands per group were used.

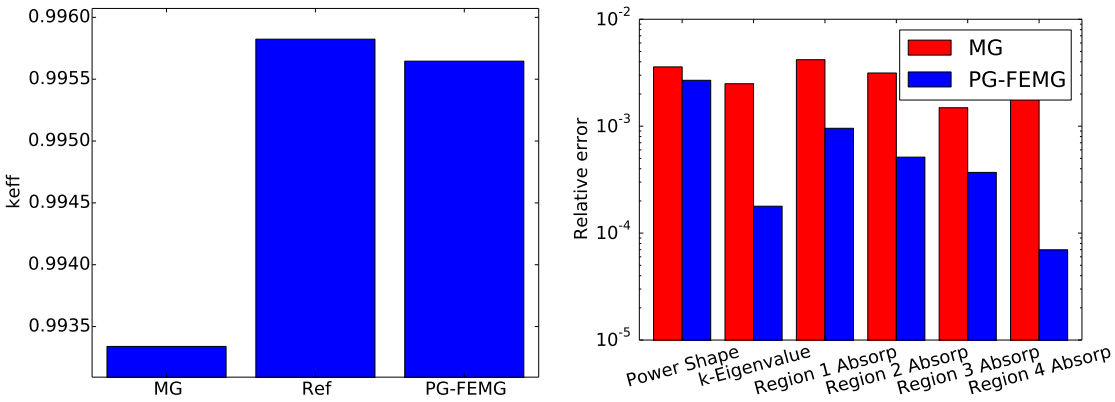

(a) k-Eigenvalues

(b) Relative Errors

Figure C.23: Comparison of eigenvalues and relative errors for problem 3 with high-fidelity weight for resolution 4 and 4 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

Fig. C.23 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.23a gives the k-eigenvalues for each of the calculations.

Fig. C.23b gives the error in k-eigenvalue, power shape, and absorption rates.

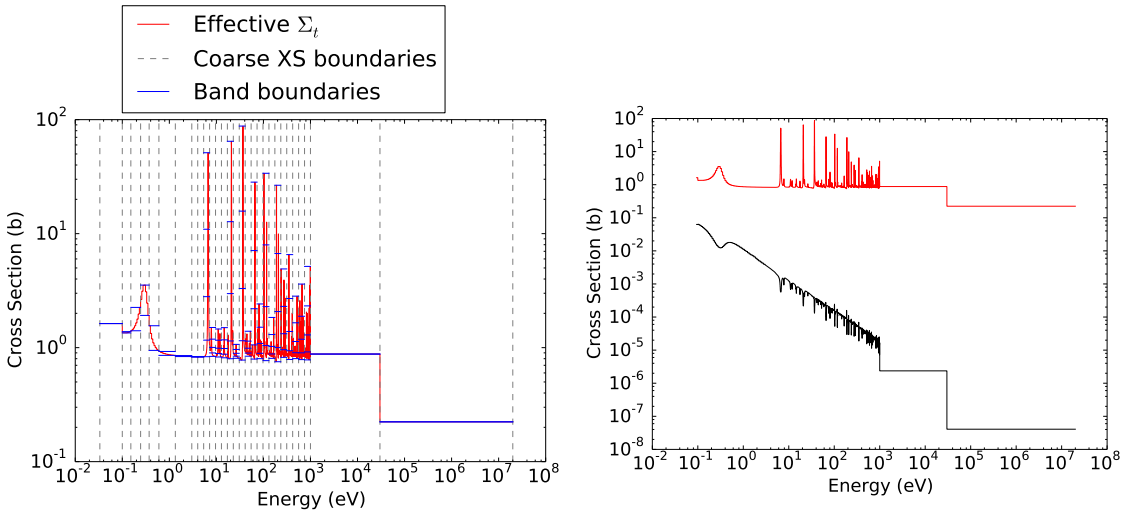

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.24: Effective total cross sections for the PG-FEMG band boundary calculation for problem 3 with high-fidelity weight for resolution 4 and 4 bands per group in the resonance region.

Fig. C.24a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.24b gives the effective total cross section and the weighting spectrum used. For this calculation, a high-fidelity weighting spectrum was used (note resonances in b).

C.2.4 Generic weighting

C.2.5 Resonance-only resolution study

C.2.5.1 Resolution 1 with 2 bands per group

Results for an energy grid of resolution factor 1. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 2 bands per group were used.

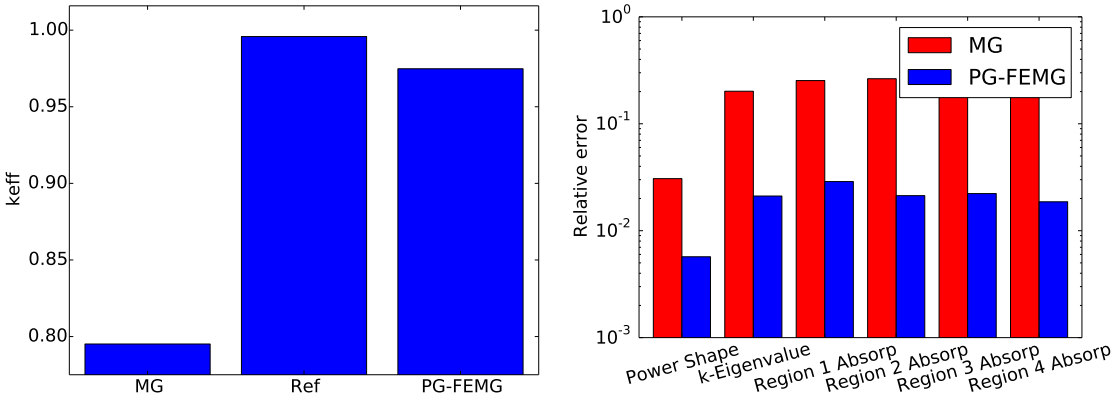

(a) k-Eigenvalues

(b) Relative Errors

Figure C.25: Comparison of eigenvalues and relative errors for problem 3 with low-fidelity weight for resolution 1 and 2 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

Fig. C.25 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.25a gives the k-eigenvalues for each of the calculations. Fig. C.25b gives the error in k-eigenvalue, power shape, and absorption rates.

Fig. C.26a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue).

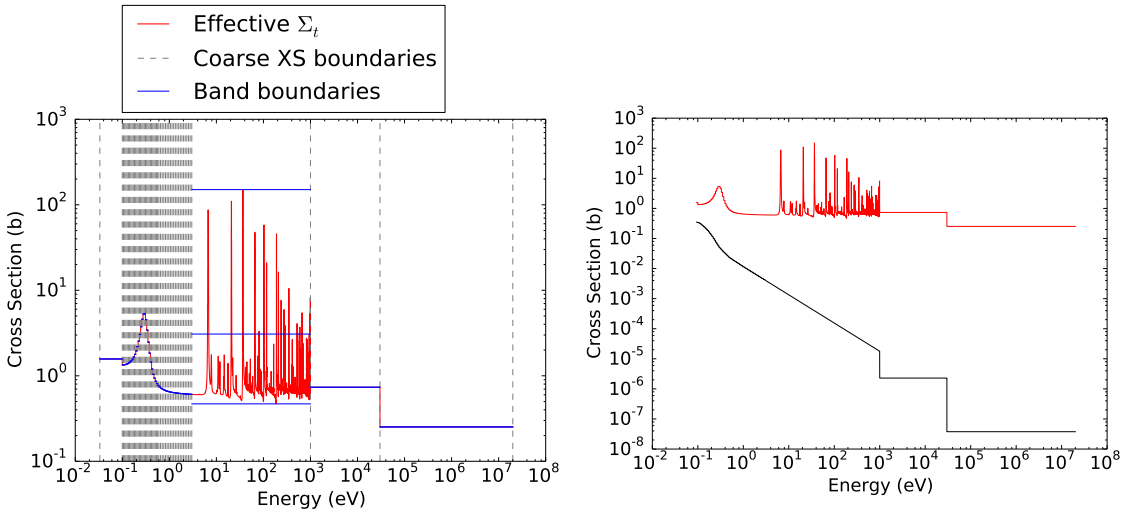

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.26: Effective total cross sections for the PG-FEMG band boundary calculation for problem 3 with low-fidelity weight for resolution 1 and 2 bands per group in the resonance region.

Note that multiple bands per coarse group were only used in the resonance region. Fig. C.26b gives the effective total cross section and the weighting spectrum used. For this calculation, a low-fidelity weighting spectrum was used (note smoothness in b).

C.2.5.2 Resolution 4 with 4 bands per group

Results for an energy grid of resolution factor 4. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 4 bands per group were used.

Fig. C.27 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.27a gives the k-eigenvalues for each of the calculations. Fig. C.27b gives the error in k-eigenvalue, power shape, and absorption rates.

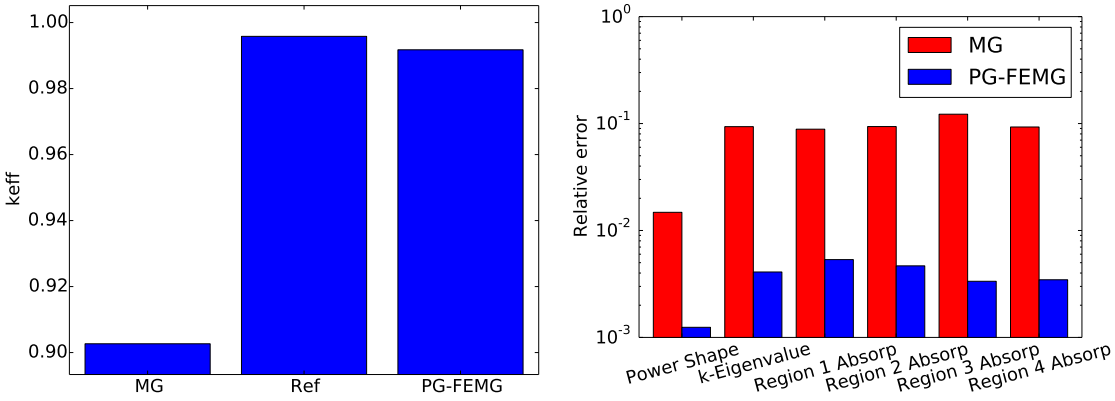

(a) k-Eigenvalues

(b) Relative Errors

Figure C.27: Comparison of eigenvalues and relative errors for problem 3 with low-fidelity weight for resolution 4 and 4 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

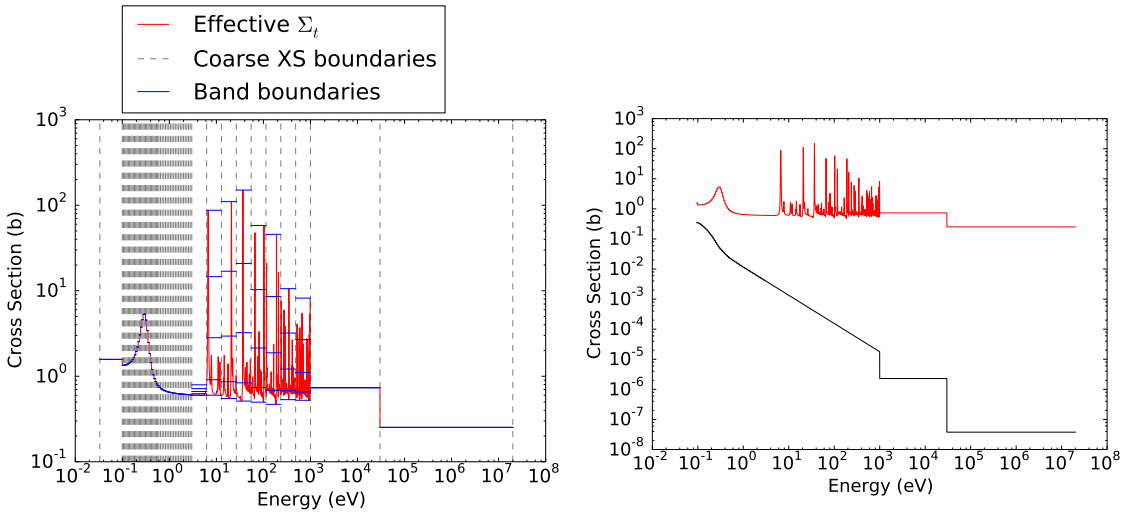

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.28: Effective total cross sections for the PG-FEMG band boundary calculation for problem 3 with low-fidelity weight for resolution 4 and 4 bands per group in the resonance region.

Fig. C.28a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.28b gives the effective total cross section and the weighting spectrum used. For this calculation, a low-fidelity weighting spectrum was used (note smoothness in b).

C.2.6 Full resolution study

C.2.6.1 Resolution 1 with 2 bands per group

Results for an energy grid of resolution factor 1. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 2 bands per group were used.

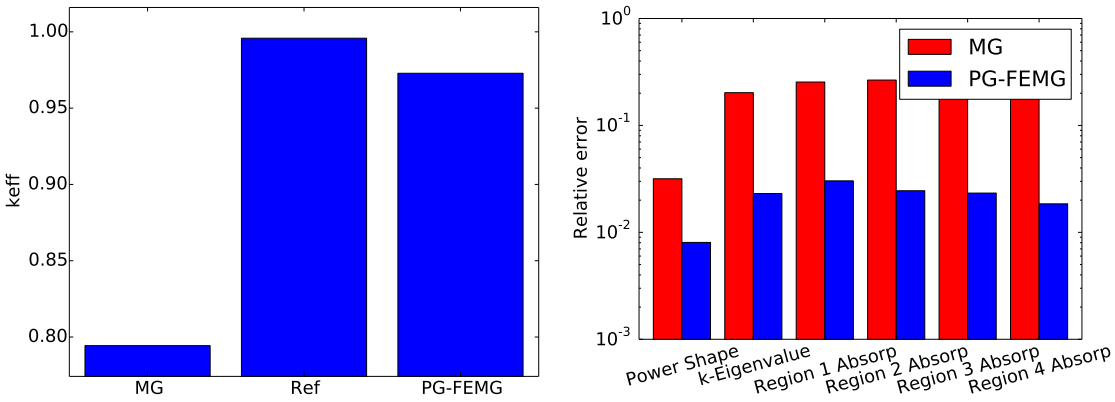

(a) k-Eigenvalues

(b) Relative Errors

Figure C.29: Comparison of eigenvalues and relative errors for problem 3 with low-fidelity weight for resolution 1 and 2 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

Fig. C.29 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.29a gives the k-eigenvalues for each of the calculations. Fig. C.29b gives the error in k-eigenvalue, power shape, and absorption rates.

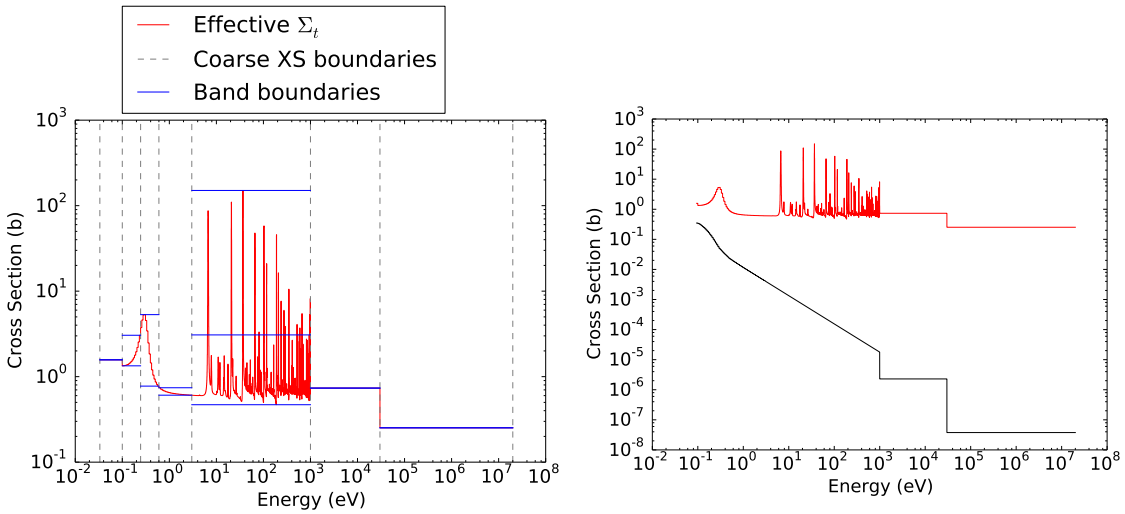

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.30: Effective total cross sections for the PG-FEMG band boundary calculation for problem 3 with low-fidelity weight for resolution 1 and 2 bands per group in the resonance region.

Fig. C.30a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.30b gives the effective total cross section and the weighting spectrum used. For this calculation, a low-fidelity weighting spectrum was used (note smoothness in b).

C.2.6.2 Resolution 4 with 4 bands per group

Results for an energy grid of resolution factor 4. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 4 bands per group were used.

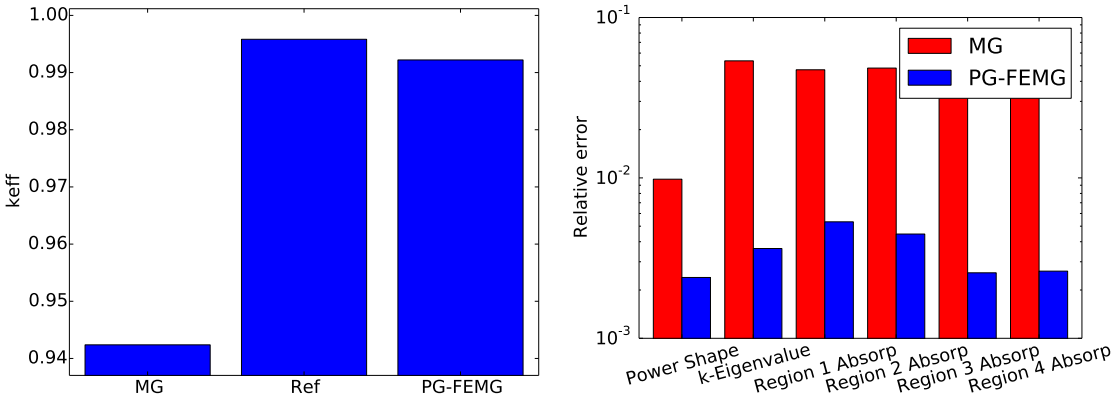

(a) k-Eigenvalues

(b) Relative Errors

Figure C.31: Comparison of eigenvalues and relative errors for problem 3 with low-fidelity weight for resolution 4 and 4 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

Fig. C.31 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.31a gives the k-eigenvalues for each of the calculations. Fig. C.31b gives the error in k-eigenvalue, power shape, and absorption rates.

Fig. C.32a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.32b gives the effective total cross section and the weighting spectrum used.

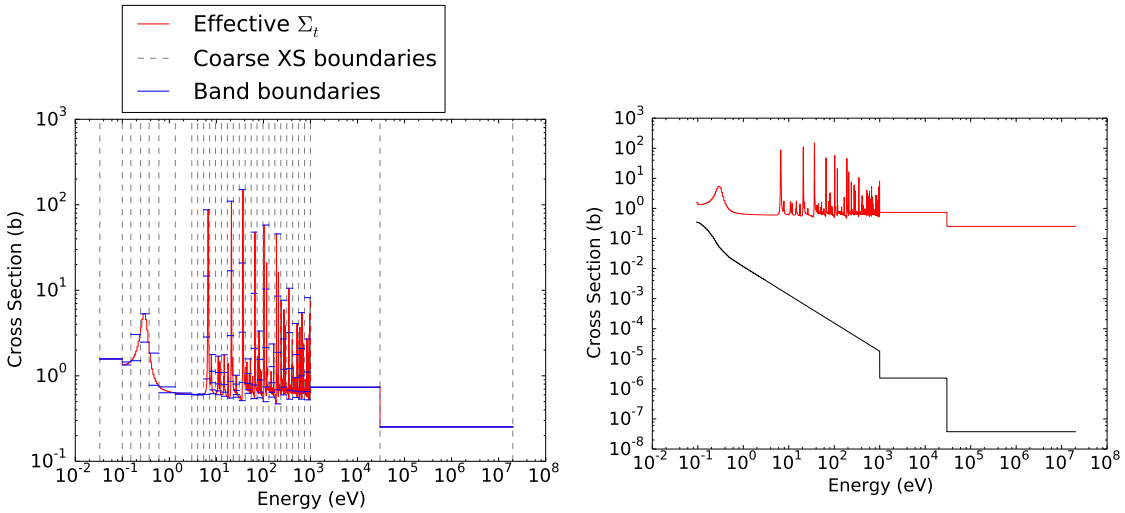

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.32: Effective total cross sections for the PG-FEMG band boundary calculation for problem 3 with low-fidelity weight for resolution 4 and 4 bands per group in the resonance region.

For this calculation, a low-fidelity weighting spectrum was used (note smoothness in b).

C.3 Problem 4 (cont.)

C.3.1 Reference-calculation weighting

C.3.2 Resonance-only resolution study

C.3.2.1 Resolution 1 with 2 bands per group

Results for an energy grid of resolution factor 1. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 2 bands per group were used.

Fig. C.33 provides error measurements between the reference solution and the MG

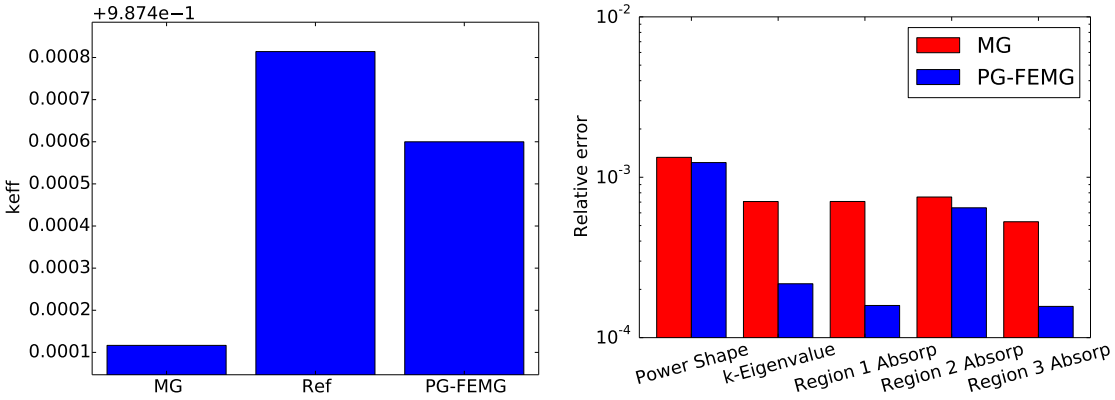

(a) k-Eigenvalues

(b) Relative Errors

Figure C.33: Comparison of eigenvalues and relative errors for problem 4 with high-fidelity weight for resolution 1 and 2 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

/ PG-FEMG solutions. Fig. C.33a gives the k-eigenvalues for each of the calculations. Fig. C.33b gives the error in k-eigenvalue, power shape, and absorption rates.

Fig. C.34a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.34b gives the effective total cross section and the weighting spectrum used. For this calculation, a high-fidelity weighting spectrum was used (note resonances in b).

C.3.2.2 Resolution 4 with 4 bands per group

Results for an energy grid of resolution factor 4. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 4 bands per group were used.

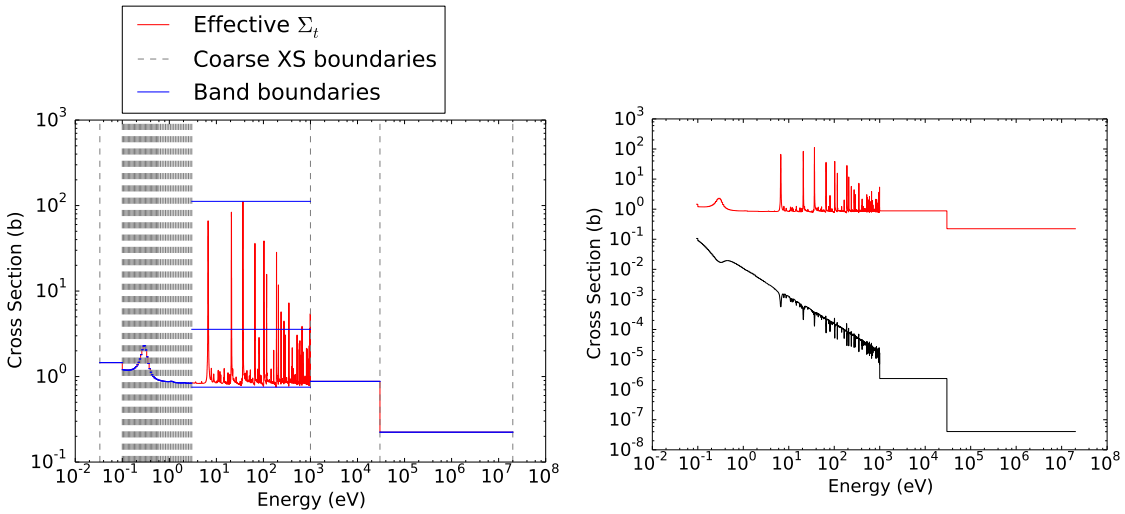

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.34: Effective total cross sections for the PG-FEMG band boundary calculation for problem 4 with high-fidelity weight for resolution 1 and 2 bands per group in the resonance region.

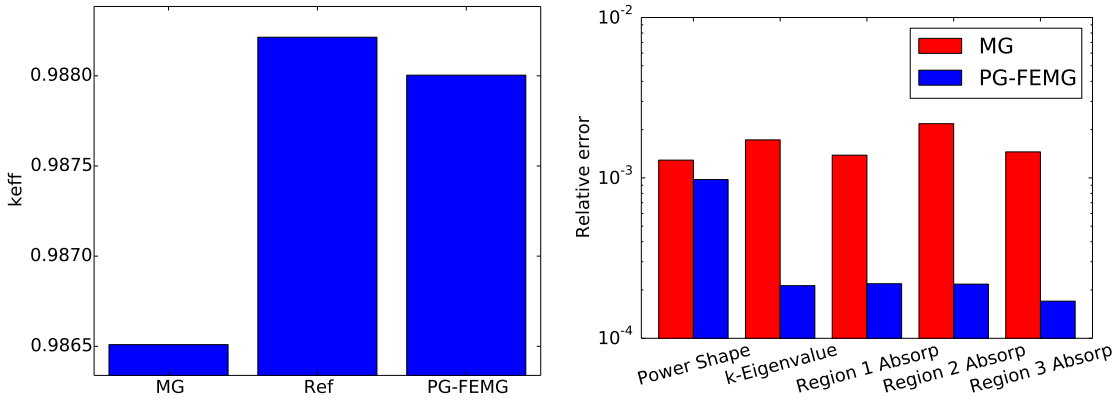

(a) k-Eigenvalues

(b) Relative Errors

Figure C.35: Comparison of eigenvalues and relative errors for problem 4 with high-fidelity weight for resolution 4 and 4 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

Fig. C.35 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.35a gives the k-eigenvalues for each of the calculations. Fig. C.35b gives the error in k-eigenvalue, power shape, and absorption rates.

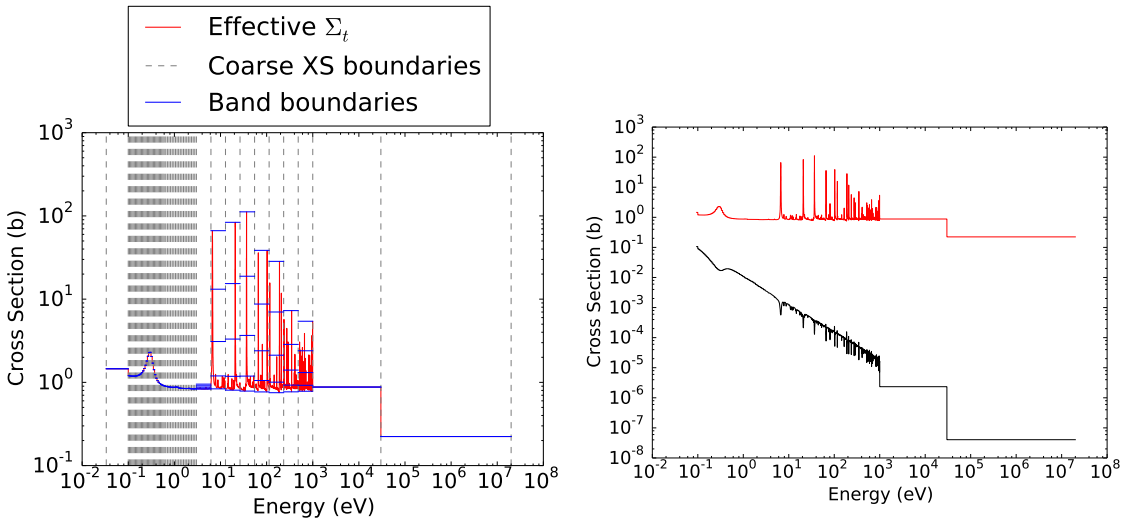

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.36: Effective total cross sections for the PG-FEMG band boundary calculation for problem 4 with high-fidelity weight for resolution 4 and 4 bands per group in the resonance region.

Fig. C.36a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.36b gives the effective total cross section and the weighting spectrum used. For this calculation, a high-fidelity weighting spectrum was used (note resonances in b).

C.3.3 Full resolution study

C.3.3.1 Resolution 1 with 2 bands per group

Results for an energy grid of resolution factor 1. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 2 bands per group were used.

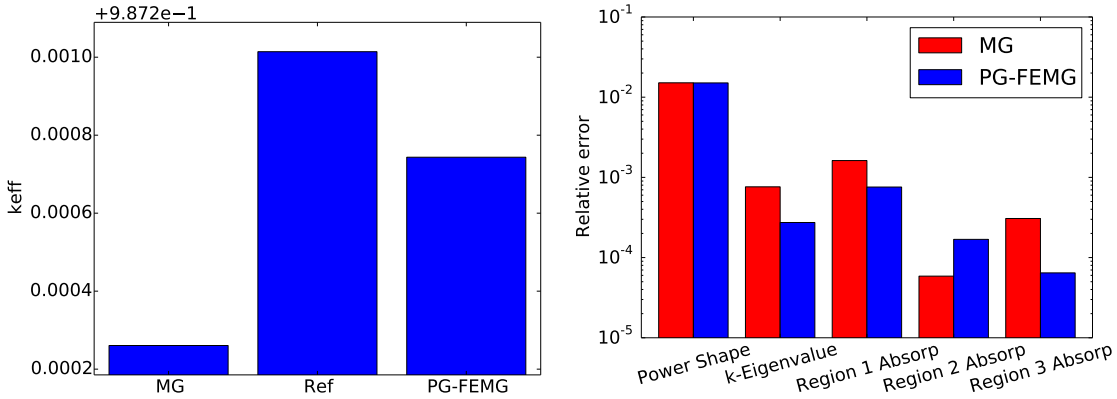

(a) k-Eigenvalues (b) Relative Errors

Figure C.37: Comparison of eigenvalues and relative errors for problem 4 with high-fidelity weight for resolution 1 and 2 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

Fig. C.37 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.37a gives the k-eigenvalues for each of the calculations. Fig. C.37b gives the error in k-eigenvalue, power shape, and absorption rates.

Fig. C.38a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region.

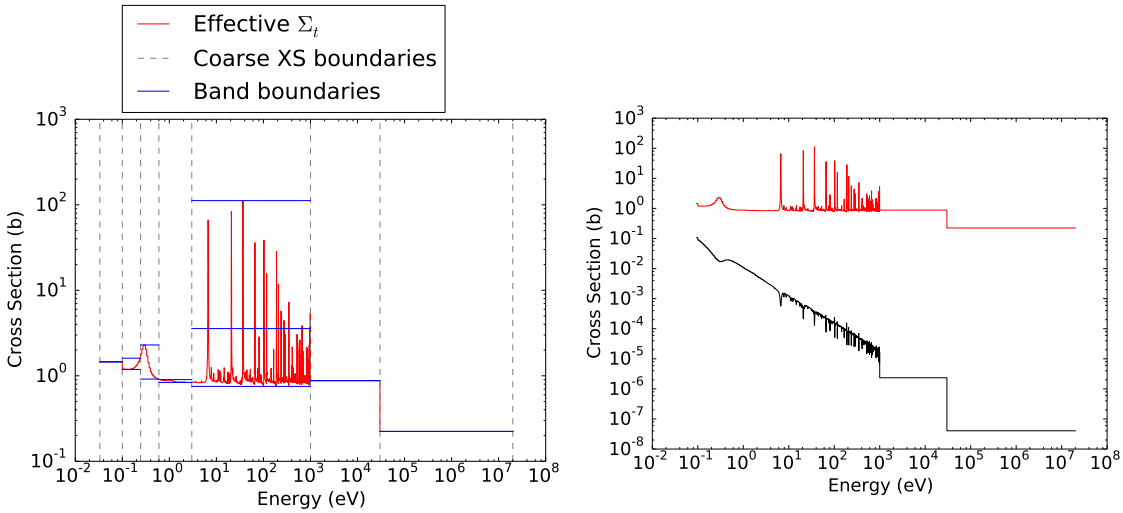

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.38: Effective total cross sections for the PG-FEMG band boundary calculation for problem 4 with high-fidelity weight for resolution 1 and 2 bands per group in the resonance region.

Fig. C.38b gives the effective total cross section and the weighting spectrum used. For this calculation, a high-fidelity weighting spectrum was used (note resonances in b).

C.3.3.2 Resolution 4 with 4 bands per group

Results for an energy grid of resolution factor 4. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 4 bands per group were used.

Fig. C.39 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.39a gives the k-eigenvalues for each of the calculations. Fig. C.39b gives the error in k-eigenvalue, power shape, and absorption rates.

Fig. C.40a shows the effective total XS used to create the band boundaries

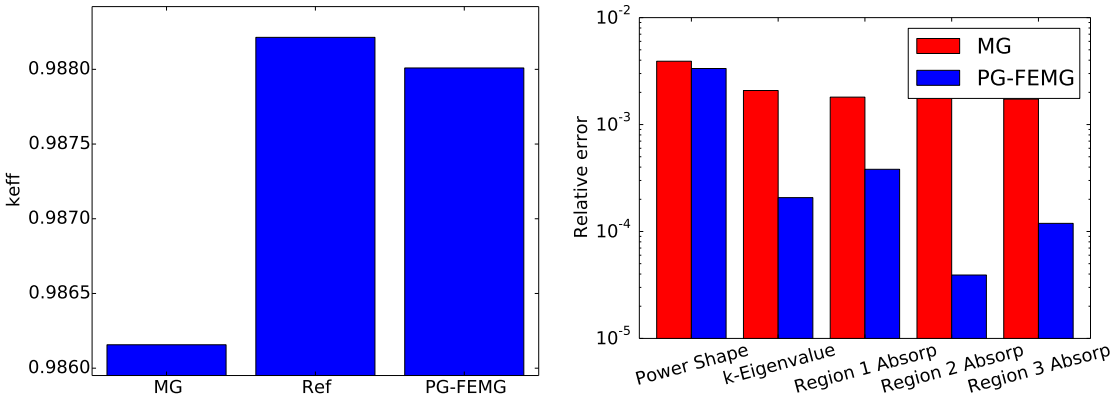

(a) k-Eigenvalues

(b) Relative Errors

Figure C.39: Comparison of eigenvalues and relative errors for problem 4 with high-fidelity weight for resolution 4 and 4 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

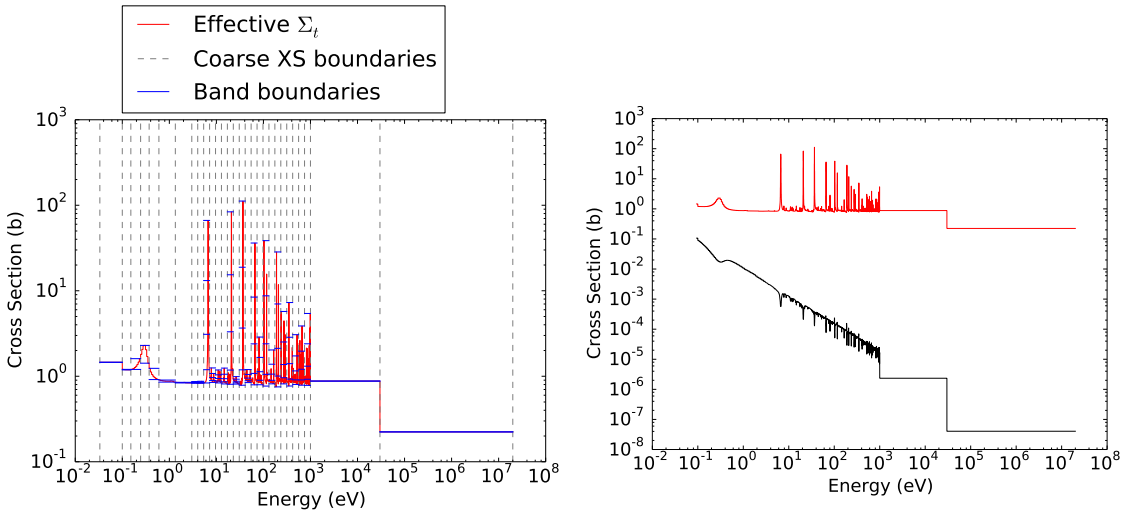

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.40: Effective total cross sections for the PG-FEMG band boundary calculation for problem 4 with high-fidelity weight for resolution 4 and 4 bands per group in the resonance region.

(red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.40b gives the effective total cross section and the weighting spectrum used. For this calculation, a high-fidelity weighting spectrum was used (note resonances in b).

C.3.4 Generic weighting

C.3.5 Resonance-only resolution study

C.3.5.1 Resolution 1 with 2 bands per group

Results for an energy grid of resolution factor 1. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 2 bands per group were used.

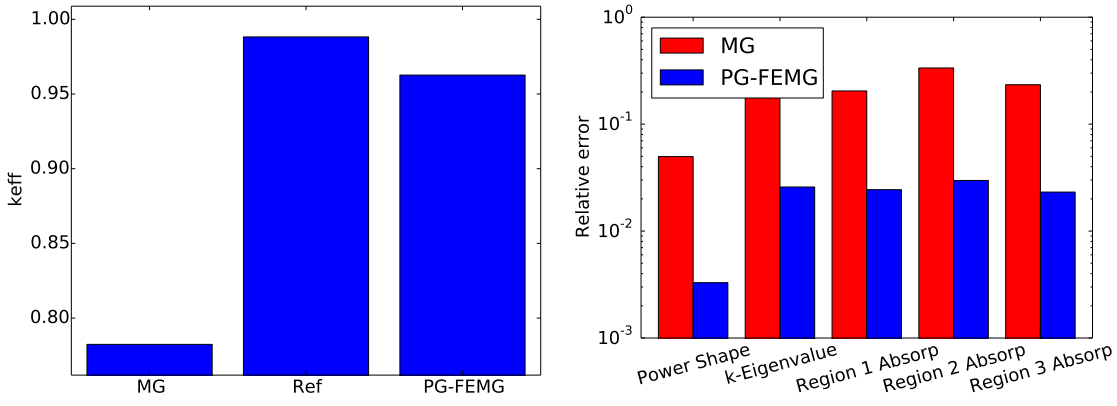

(a) k-Eigenvalues

(b) Relative Errors

Figure C.41: Comparison of eigenvalues and relative errors for problem 4 with low-fidelity weight for resolution 1 and 2 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

Fig. C.41 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.41a gives the k-eigenvalues for each of the calculations. Fig. C.41b gives the error in k-eigenvalue, power shape, and absorption rates.

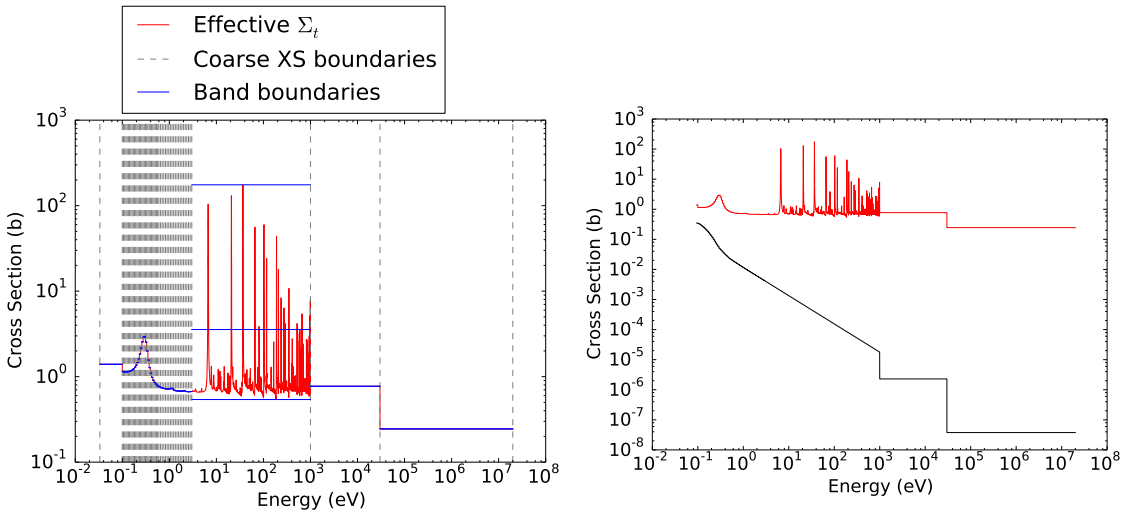

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.42: Effective total cross sections for the PG-FEMG band boundary calculation for problem 4 with low-fidelity weight for resolution 1 and 2 bands per group in the resonance region.

Fig. C.42a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.42b gives the effective total cross section and the weighting spectrum used. For this calculation, a low-fidelity weighting spectrum was used (note smoothness in b).

C.3.5.2 Resolution 4 with 4 bands per group

Results for an energy grid of resolution factor 4. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 4 bands per group were used.

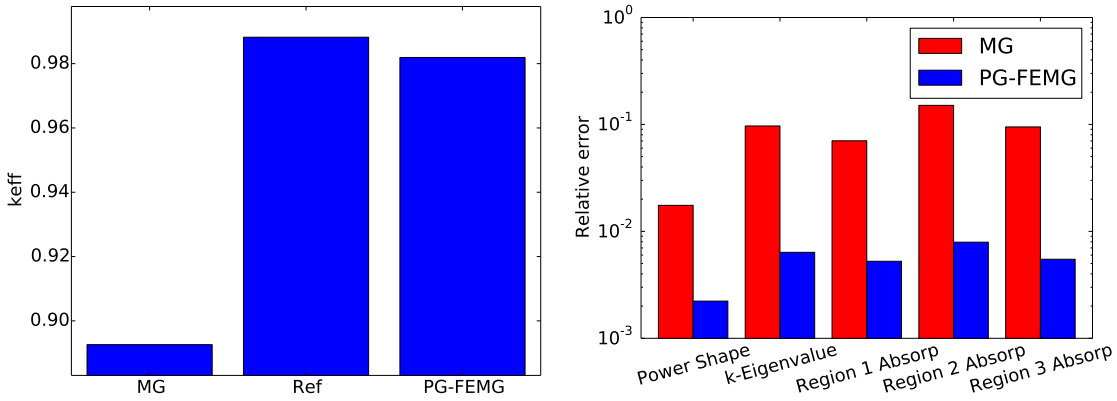

(a) k-Eigenvalues

(b) Relative Errors

Figure C.43: Comparison of eigenvalues and relative errors for problem 4 with low-fidelity weight for resolution 4 and 4 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

Fig. C.43 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.43a gives the k-eigenvalues for each of the calculations. Fig. C.43b gives the error in k-eigenvalue, power shape, and absorption rates.

Fig. C.44a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.44b gives the effective total cross section and the weighting spectrum used.

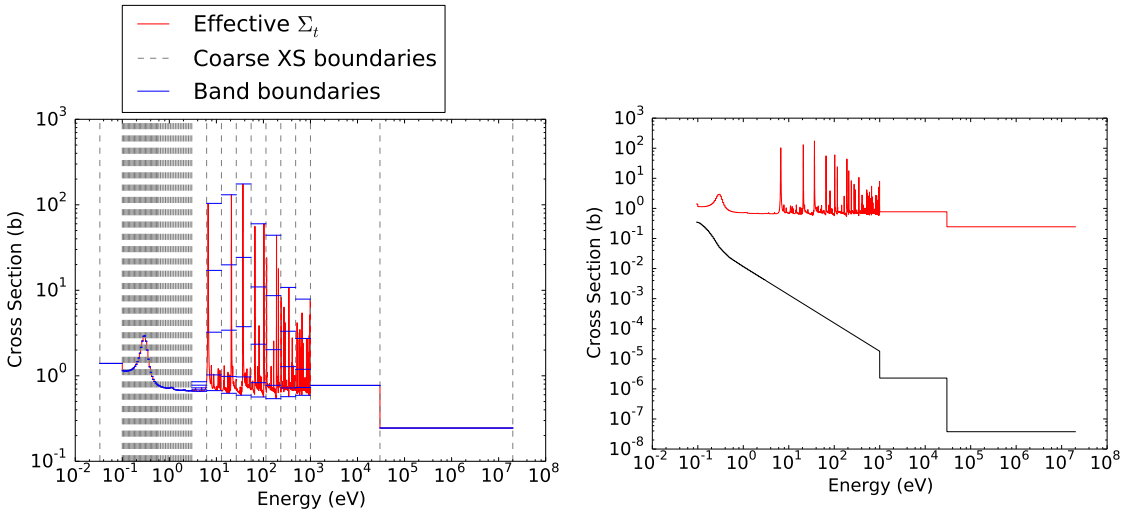

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.44: Effective total cross sections for the PG-FEMG band boundary calculation for problem 4 with low-fidelity weight for resolution 4 and 4 bands per group in the resonance region.

For this calculation, a low-fidelity weighting spectrum was used (note smoothness in b).

C.3.6 Full resolution study

C.3.6.1 Resolution 1 with 2 bands per group

Results for an energy grid of resolution factor 1. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 2 bands per group were used.

Fig. C.45 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.45a gives the k-eigenvalues for each of the calculations. Fig. C.45b gives the error in k-eigenvalue, power shape, and absorption rates.

Fig. C.46a shows the effective total XS used to create the band boundaries

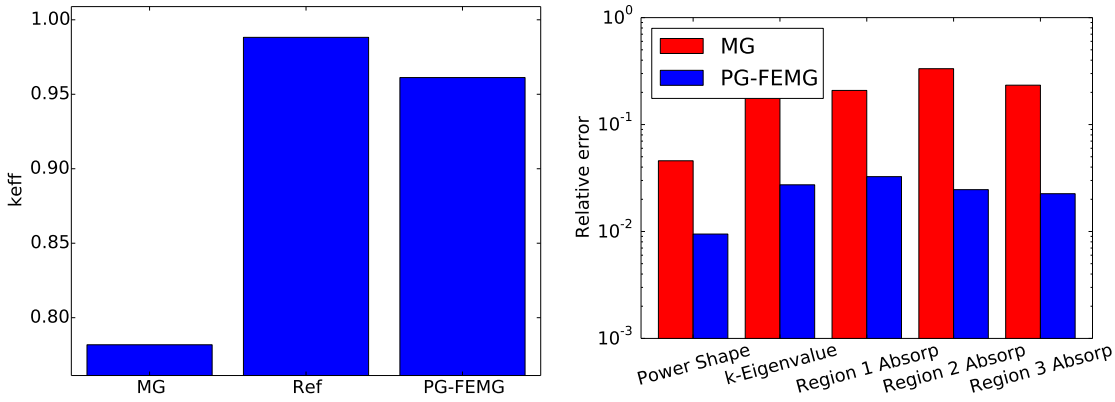

(a) k-Eigenvalues

(b) Relative Errors

Figure C.45: Comparison of eigenvalues and relative errors for problem 4 with low-fidelity weight for resolution 1 and 2 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

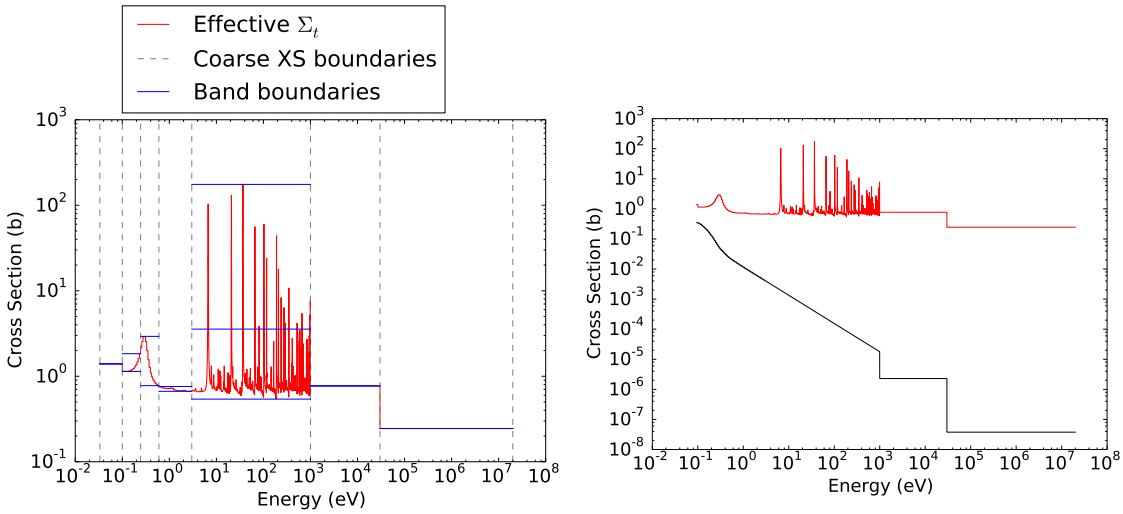

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.46: Effective total cross sections for the PG-FEMG band boundary calculation for problem 4 with low-fidelity weight for resolution 1 and 2 bands per group in the resonance region.

(red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.46b gives the effective total cross section and the weighting spectrum used. For this calculation, a low-fidelity weighting spectrum was used (note smoothness in b).

C.3.6.2 Resolution 4 with 4 bands per group

Results for an energy grid of resolution factor 4. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 4 bands per group were used.

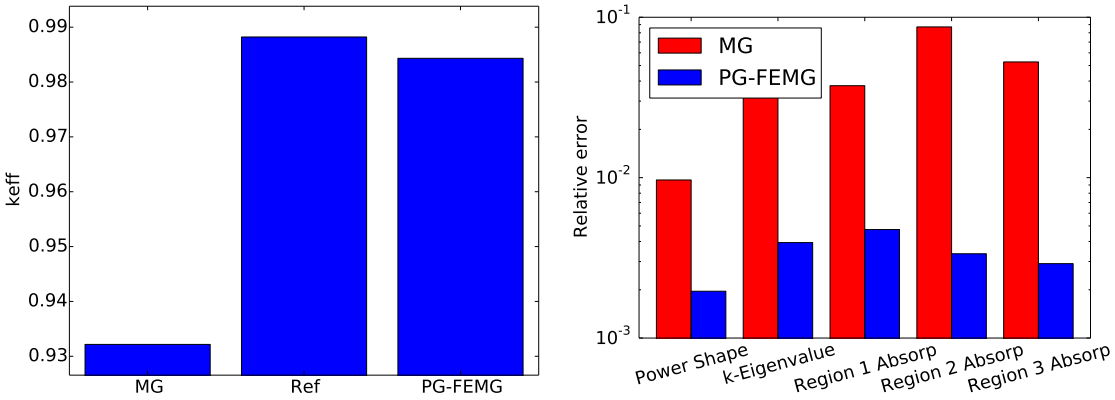

(a) k-Eigenvalues (b) Relative Errors

Figure C.47: Comparison of eigenvalues and relative errors for problem 4 with low-fidelity weight for resolution 4 and 4 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

Fig. C.47 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.47a gives the k-eigenvalues for each of the calculations.

Fig. C.47b gives the error in k-eigenvalue, power shape, and absorption rates.

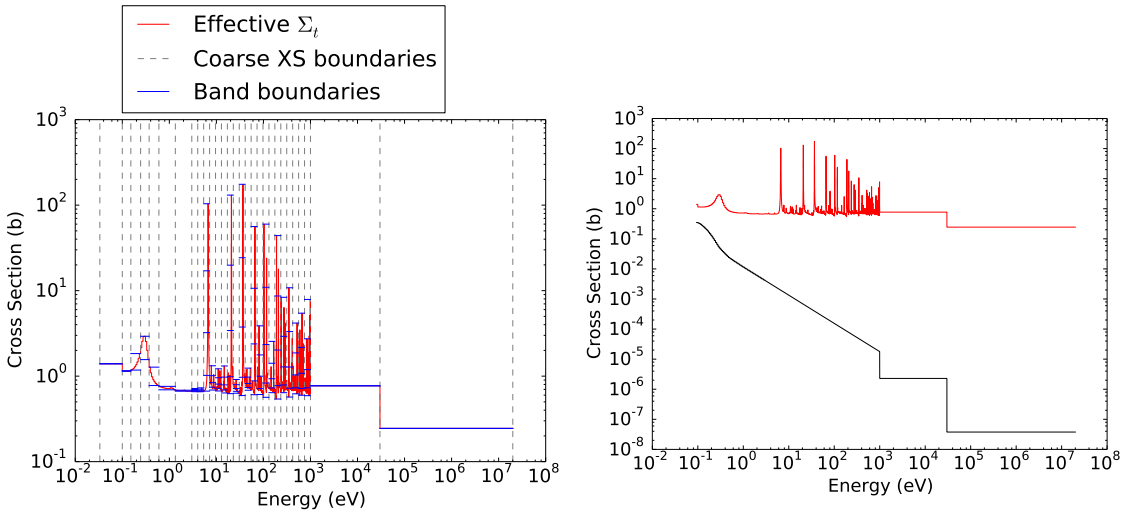

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.48: Effective total cross sections for the PG-FEMG band boundary calculation for problem 4 with low-fidelity weight for resolution 4 and 4 bands per group in the resonance region.

Fig. C.48a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.48b gives the effective total cross section and the weighting spectrum used. For this calculation, a low-fidelity weighting spectrum was used (note smoothness in b).

C.4 Problem 5 (cont.)

C.4.1 Reference-calculation weighting

C.4.2 Resonance-only resolution study

C.4.2.1 Resolution 1 with 2 bands per group

Results for an energy grid of resolution factor 1. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 2 bands per group were used.

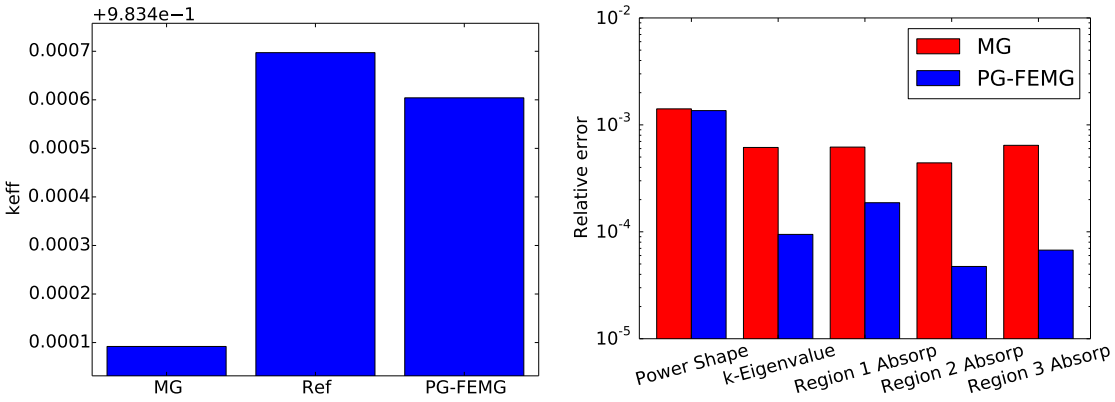

(a) k-Eigenvalues (b) Relative Errors

Figure C.49: Comparison of eigenvalues and relative errors for problem 5 with high-fidelity weight for resolution 1 and 2 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

Fig. C.49 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.49a gives the k-eigenvalues for each of the calculations. Fig. C.49b gives the error in k-eigenvalue, power shape, and absorption rates.

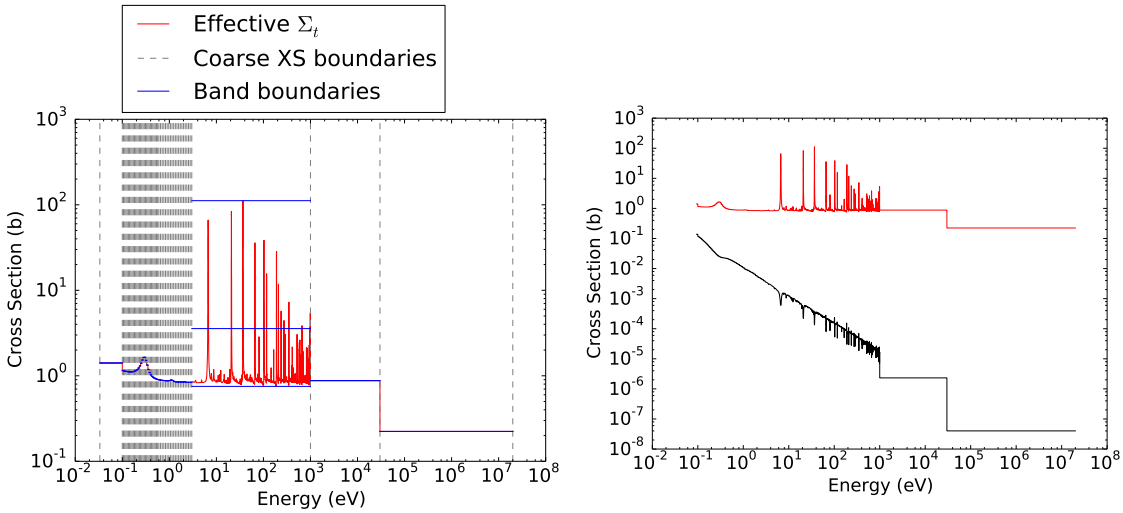

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.50: Effective total cross sections for the PG-FEMG band boundary calculation for problem 5 with high-fidelity weight for resolution 1 and 2 bands per group in the resonance region.

Fig. C.50a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.50b gives the effective total cross section and the weighting spectrum used. For this calculation, a high-fidelity weighting spectrum was used (note resonances in b).

C.4.2.2 Resolution 4 with 4 bands per group

Results for an energy grid of resolution factor 4. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 4 bands per group were used.

Fig. C.51 provides error measurements between the reference solution and the MG

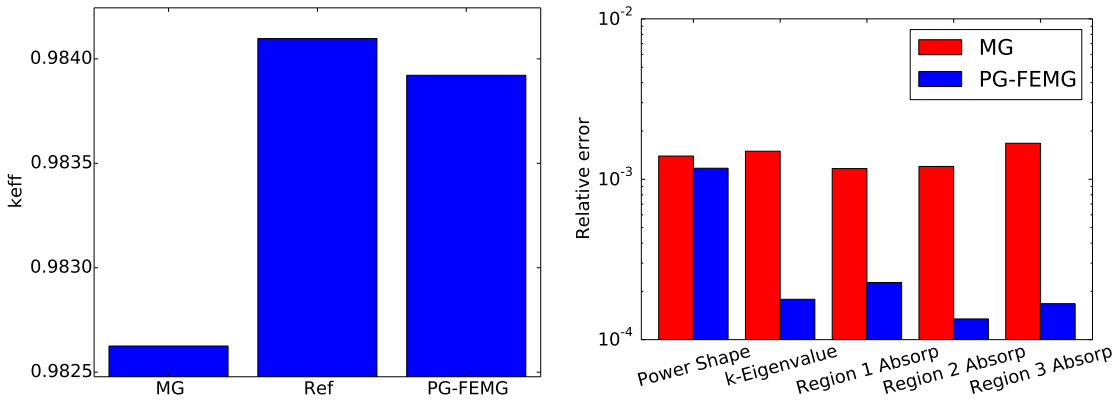

(a) k-Eigenvalues (b) Relative Errors

Figure C.51: Comparison of eigenvalues and relative errors for problem 5 with high-fidelity weight for resolution 4 and 4 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

/ PG-FEMG solutions. Fig. C.51a gives the k-eigenvalues for each of the calculations. Fig. C.51b gives the error in k-eigenvalue, power shape, and absorption rates.

Fig. C.52a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.52b gives the effective total cross section and the weighting spectrum used. For this calculation, a high-fidelity weighting spectrum was used (note resonances in b).

C.4.3 Full resolution study

C.4.3.1 Resolution 1 with 2 bands per group

Results for an energy grid of resolution factor 1. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 2 bands per group

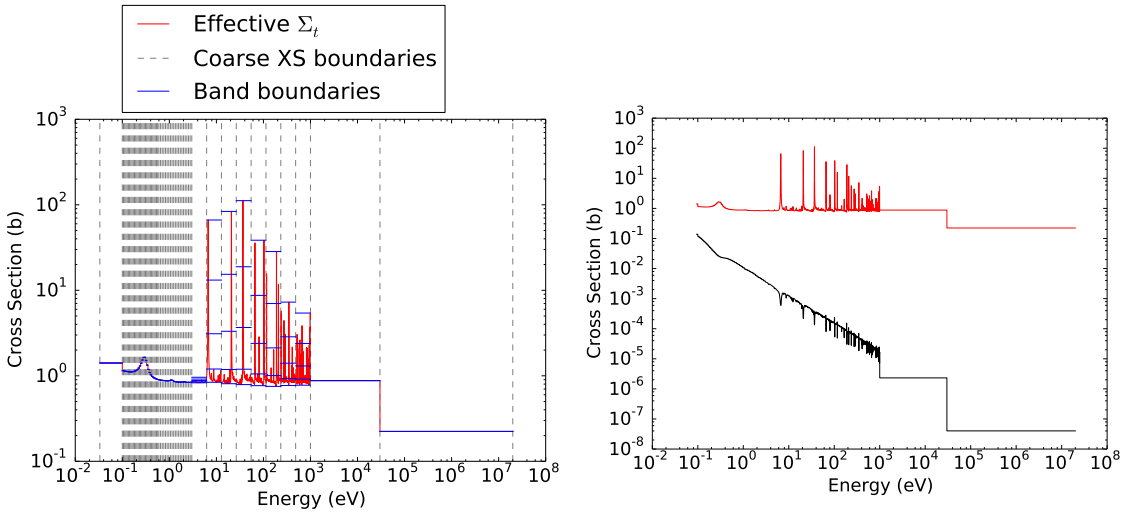

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.52: Effective total cross sections for the PG-FEMG band boundary calculation for problem 5 with high-fidelity weight for resolution 4 and 4 bands per group in the resonance region.

were used.

Fig. C.53 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.53a gives the k-eigenvalues for each of the calculations. Fig. C.53b gives the error in k-eigenvalue, power shape, and absorption rates.

Fig. C.54a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.54b gives the effective total cross section and the weighting spectrum used. For this calculation, a high-fidelity weighting spectrum was used (note resonances in b).

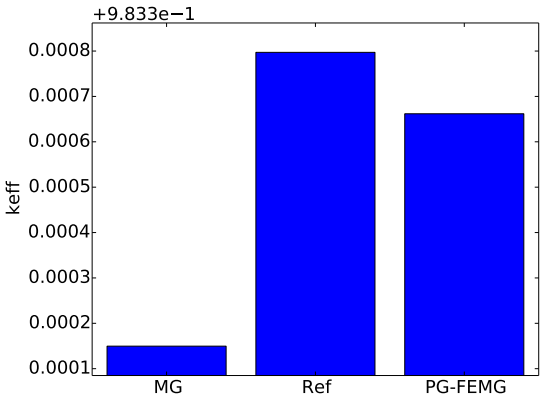

(a) k-Eigenvalues

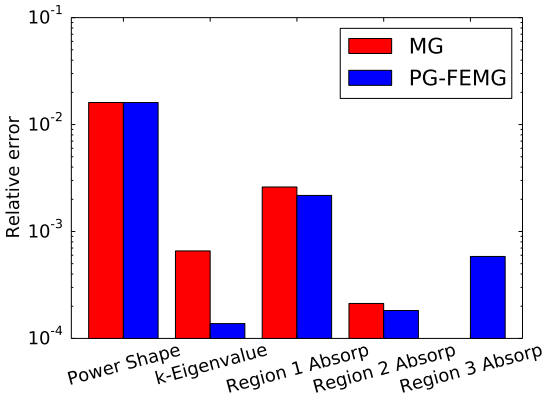

(b) Relative Errors

Figure C.53: Comparison of eigenvalues and relative errors for problem 5 with high-fidelity weight for resolution 1 and 2 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

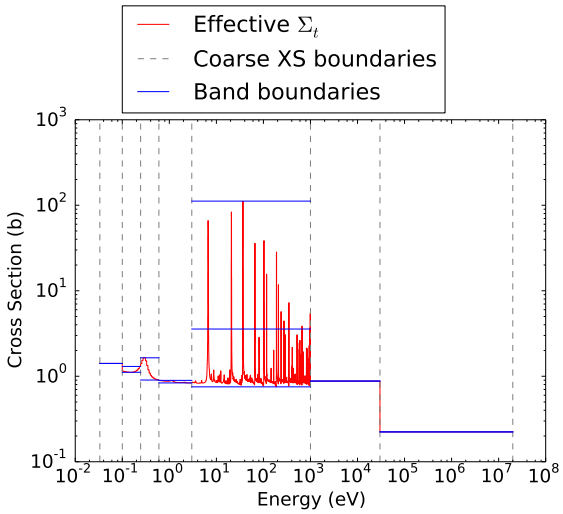

(a) Effective total cross section and band boundaries.

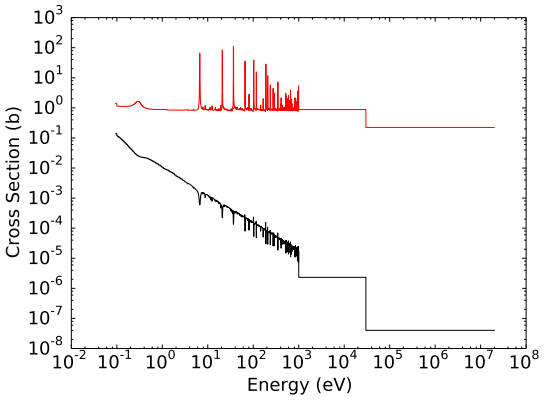

(b) Effective total cross section and the weighting spectrum.

Figure C.54: Effective total cross sections for the PG-FEMG band boundary calculation for problem 5 with high-fidelity weight for resolution 1 and 2 bands per group in the resonance region.

C.4.3.2 Resolution 4 with 4 bands per group

Results for an energy grid of resolution factor 4. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 4 bands per group were used.

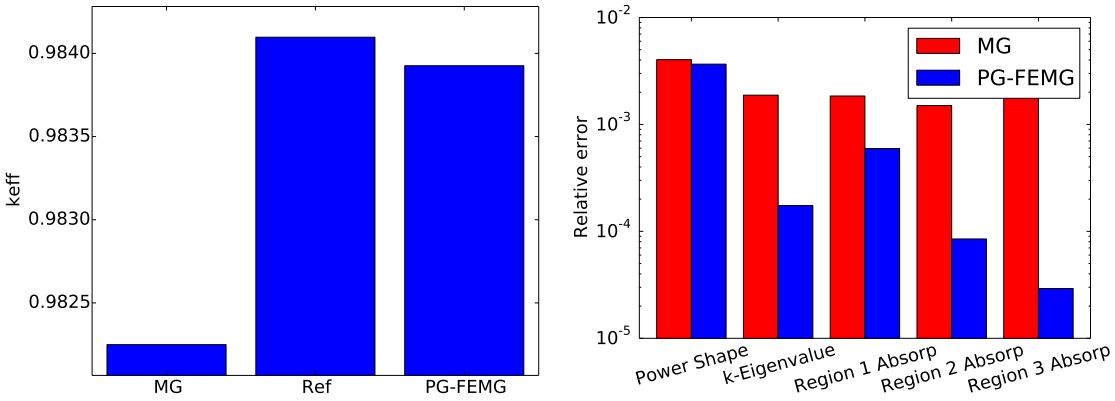

(a) k-Eigenvalues

(b) Relative Errors

Figure C.55: Comparison of eigenvalues and relative errors for problem 5 with high-fidelity weight for resolution 4 and 4 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

Fig. C.55 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.55a gives the k-eigenvalues for each of the calculations. Fig. C.55b gives the error in k-eigenvalue, power shape, and absorption rates.

Fig. C.56a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.56b gives the effective total cross section and the weighting spectrum used.

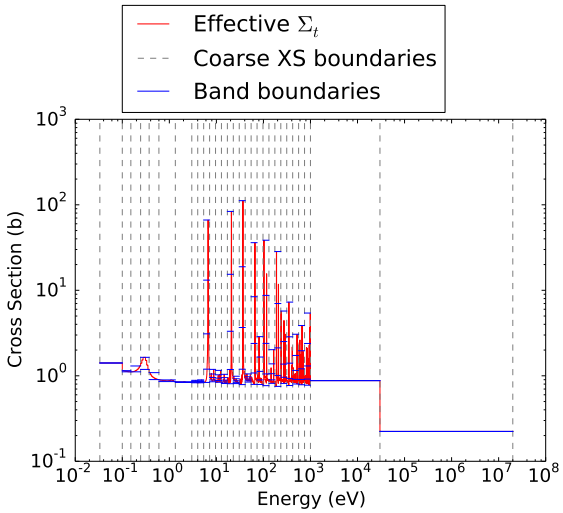

(a) Effective total cross section and band boundaries.

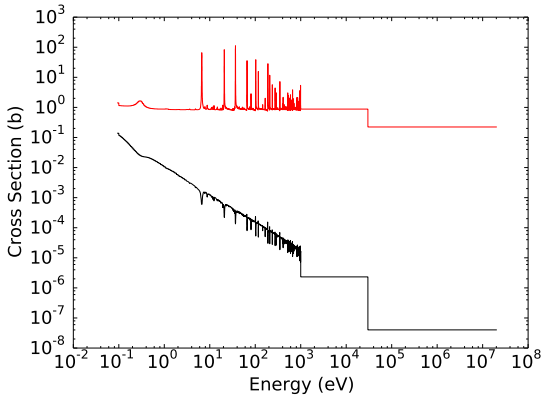

(b) Effective total cross section and the weighting spectrum.

Figure C.56: Effective total cross sections for the PG-FEMG band boundary calculation for problem 5 with high-fidelity weight for resolution 4 and 4 bands per group in the resonance region.

For this calculation, a high-fidelity weighting spectrum was used (note resonances in b).

C.4.4 Generic weighting

C.4.5 Resonance-only resolution study

C.4.5.1 Resolution 1 with 2 bands per group

Results for an energy grid of resolution factor 1. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 2 bands per group were used.

Fig. C.57 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.57a gives the k-eigenvalues for each of the calculations. Fig. C.57b gives the error in k-eigenvalue, power shape, and absorption rates.

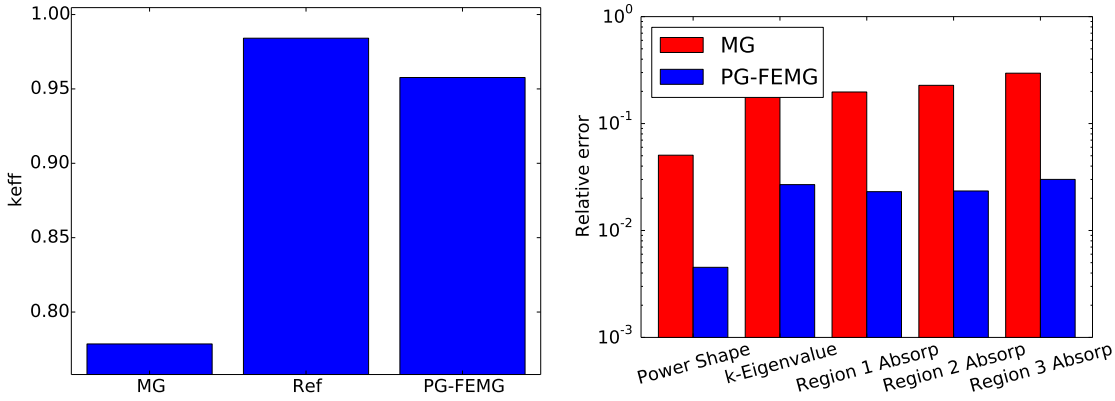

(a) k-Eigenvalues

(b) Relative Errors

Figure C.57: Comparison of eigenvalues and relative errors for problem 5 with low-fidelity weight for resolution 1 and 2 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

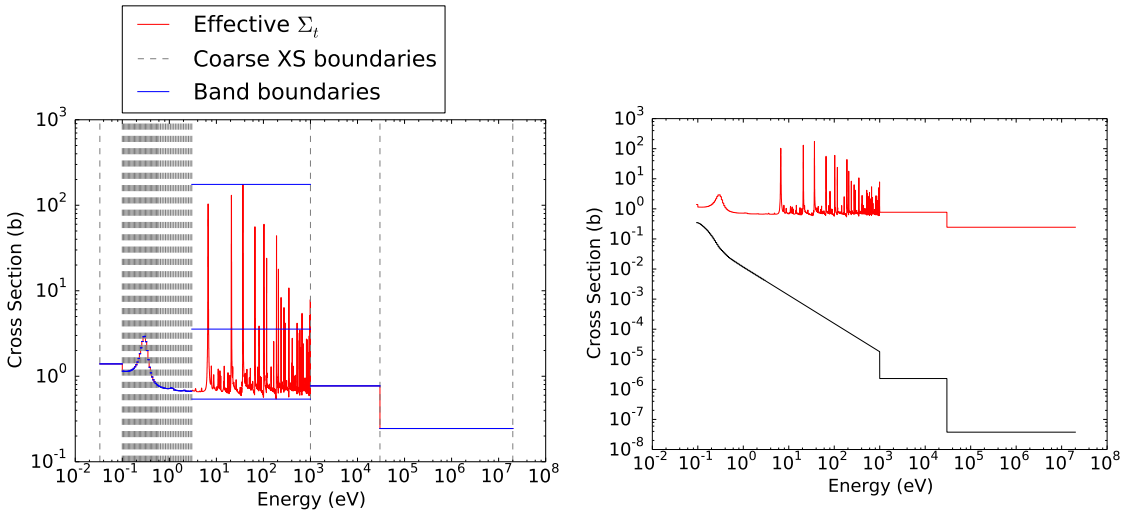

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.58: Effective total cross sections for the PG-FEMG band boundary calculation for problem 5 with low-fidelity weight for resolution 1 and 2 bands per group in the resonance region.

Fig. C.58a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.58b gives the effective total cross section and the weighting spectrum used. For this calculation, a low-fidelity weighting spectrum was used (note smoothness in b).

C.4.5.2 Resolution 4 with 4 bands per group

Results for an energy grid of resolution factor 4. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 4 bands per group were used.

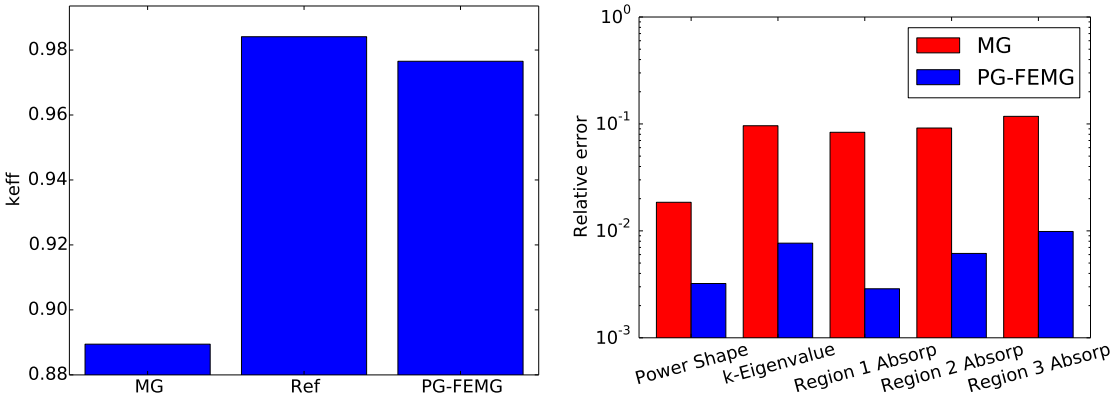

(a) k-Eigenvalues

(b) Relative Errors

Figure C.59: Comparison of eigenvalues and relative errors for problem 5 with low-fidelity weight for resolution 4 and 4 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

Fig. C.59 provides error measurements between the reference solution and the MG

/ PG-FEMG solutions. Fig. C.59a gives the k-eigenvalues for each of the calculations. Fig. C.59b gives the error in k-eigenvalue, power shape, and absorption rates.

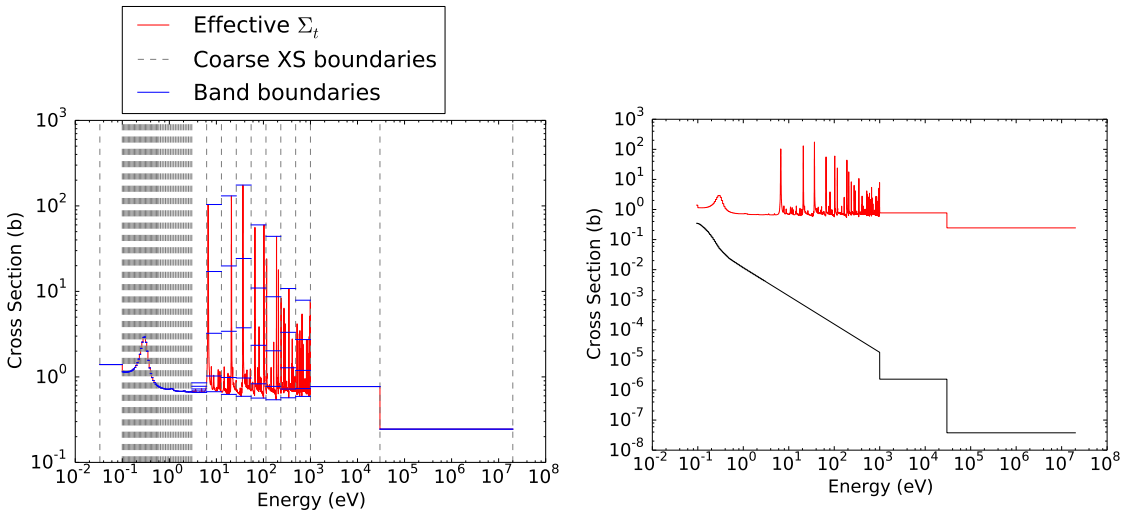

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.60: Effective total cross sections for the PG-FEMG band boundary calculation for problem 5 with low-fidelity weight for resolution 4 and 4 bands per group in the resonance region.

Fig. C.60a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.60b gives the effective total cross section and the weighting spectrum used. For this calculation, a low-fidelity weighting spectrum was used (note smoothness in b).

C.4.6 Full resolution study

C.4.6.1 Resolution 1 with 2 bands per group

Results for an energy grid of resolution factor 1. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 2 bands per group were used.

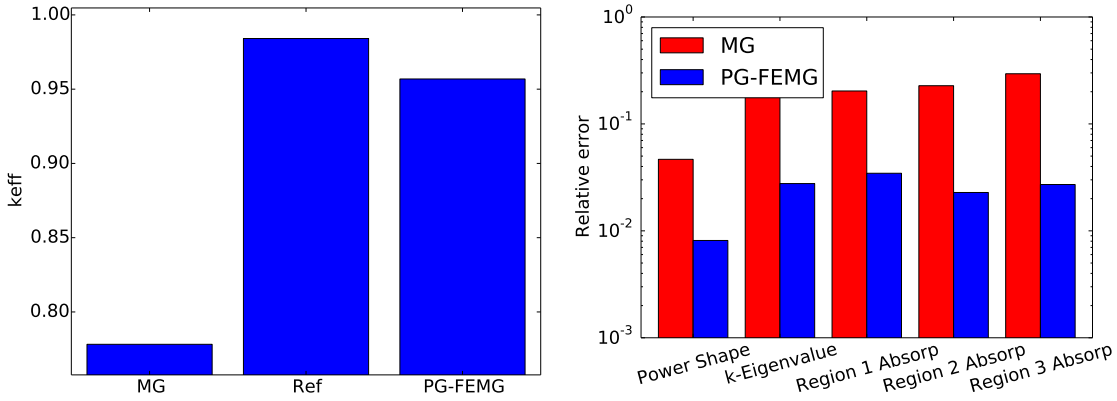

(a) k-Eigenvalues (b) Relative Errors

Figure C.61: Comparison of eigenvalues and relative errors for problem 5 with low-fidelity weight for resolution 1 and 2 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

Fig. C.61 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.61a gives the k-eigenvalues for each of the calculations. Fig. C.61b gives the error in k-eigenvalue, power shape, and absorption rates.

Fig. C.62a shows the effective total XS used to create the band boundaries (red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region.

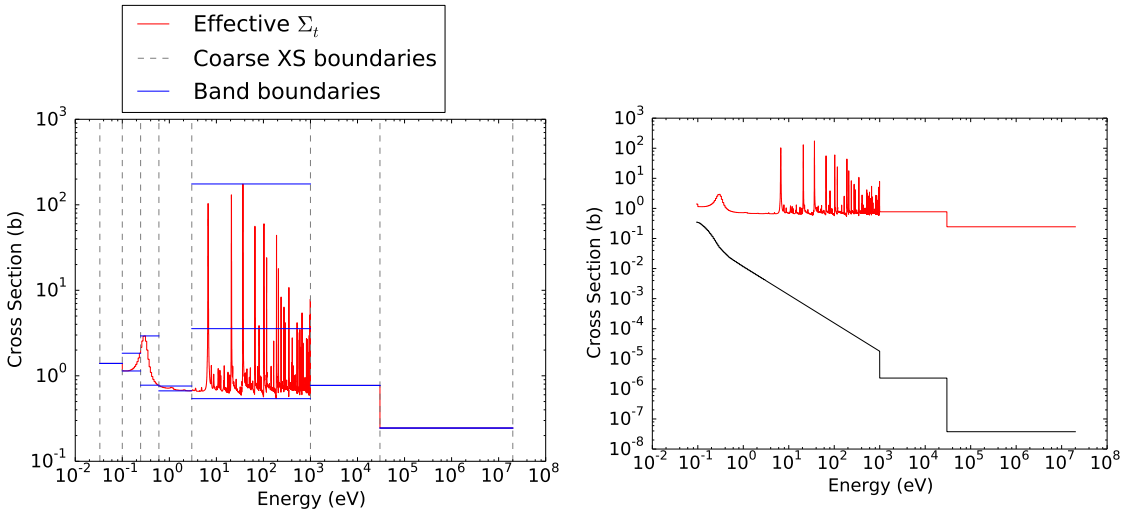

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.62: Effective total cross sections for the PG-FEMG band boundary calculation for problem 5 with low-fidelity weight for resolution 1 and 2 bands per group in the resonance region.

Fig. C.62b gives the effective total cross section and the weighting spectrum used. For this calculation, a low-fidelity weighting spectrum was used (note smoothness in b).

C.4.6.2 Resolution 4 with 4 bands per group

Results for an energy grid of resolution factor 4. An energy resolution factor of n implies the resonance region has n times as many energy DOF as the resolution factor of 1 case. For the generalized-multigroup (PG-FEMG) case, 4 bands per group were used.

Fig. C.63 provides error measurements between the reference solution and the MG / PG-FEMG solutions. Fig. C.63a gives the k-eigenvalues for each of the calculations. Fig. C.63b gives the error in k-eigenvalue, power shape, and absorption rates.

Fig. C.64a shows the effective total XS used to create the band boundaries

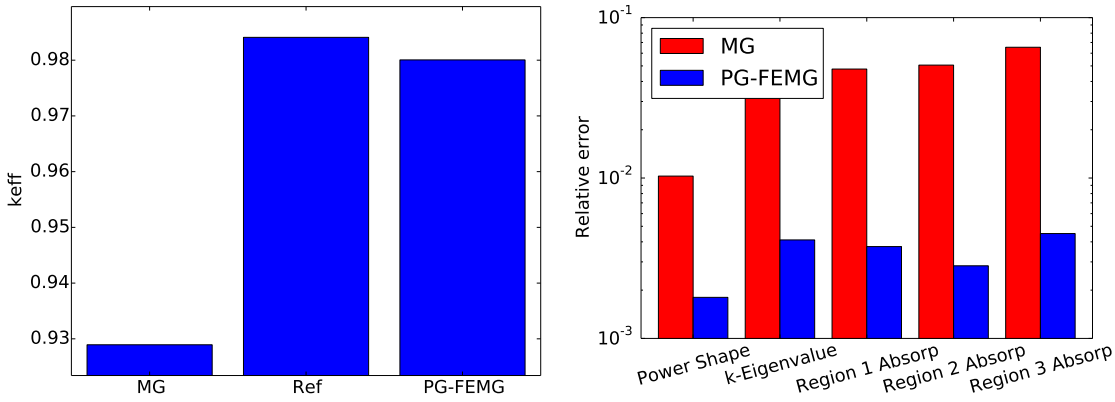

(a) k-Eigenvalues

(b) Relative Errors

Figure C.63: Comparison of eigenvalues and relative errors for problem 5 with low-fidelity weight for resolution 4 and 4 bands per group in the resonance region. Values are given for the reference, MG and PG-FEMG methods.

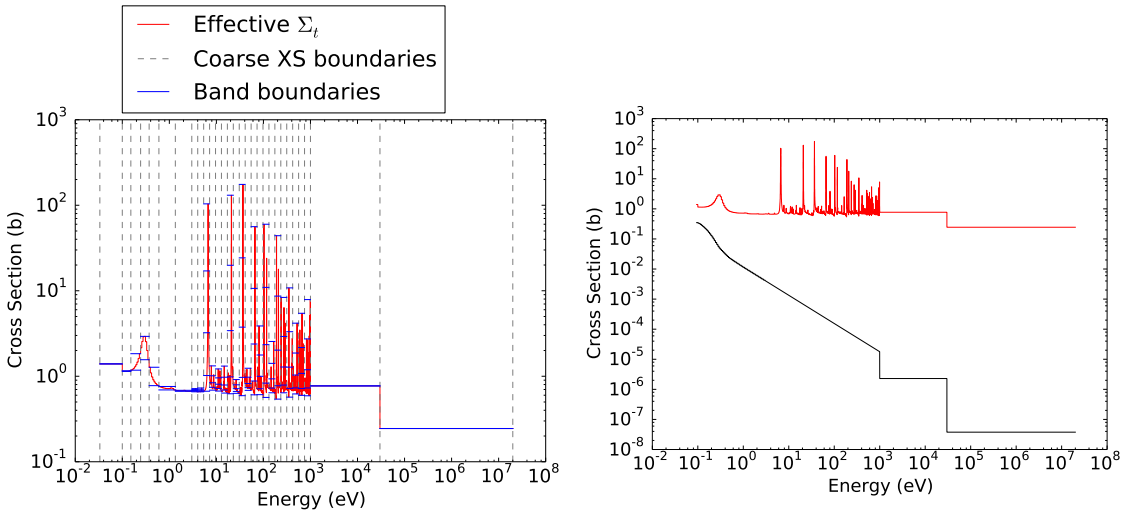

(a) Effective total cross section and band boundaries.

(b) Effective total cross section and the weighting spectrum.

Figure C.64: Effective total cross sections for the PG-FEMG band boundary calculation for problem 5 with low-fidelity weight for resolution 4 and 4 bands per group in the resonance region.

(red), the coarse-groups (dotted lines) and the band boundaries themselves (blue). Note that multiple bands per coarse group were only used in the resonance region. Fig. C.64b gives the effective total cross section and the weighting spectrum used. For this calculation, a low-fidelity weighting spectrum was used (note smoothness in b).