
APPLICATIONS OF THE
KARHUNEN-LOÉVE TRANSFORM FOR
BASIS GENERATION IN THE
RESPONSE MATRIX METHOD

by

RICHARD REED

B.S., Kansas State University, 2011

A THESIS

submitted in partial fulfillment of the requirements for the degree

MASTER OF SCIENCE

Department of Mechanical and Nuclear Engineering

College of Engineering

KANSAS STATE UNIVERSITY

Manhattan, KS

2015

Approved by:

Major Professor

Jeremy Roberts

Abstract

A novel approach based on the Karhunen-Loève Transform (KLT) is presented for treatment of the energy variable in response matrix methods, which are based on the partitioning of global domains into independent nodes linked by approximate boundary conditions. These conditions are defined using truncated expansions of nodal boundary fluxes in each phase-space variable (i.e., space, angle, and energy). There are several ways in which to represent the dependence on these variables, each of which results in a trade-off between accuracy and speed. This work provides a method to expand in energy that can reduce the number of energy degrees of freedom needed for sub-0.1% errors in nodal fission densities by up to an order of magnitude. The Karhunen-Loève Transform is used to generate basis sets for expansion in the energy variable that maximize the amount of physics captured by low-order moments, thus permitting low-order expansions with less error than basis sets previously studied, e.g., the Discrete Legendre Polynomials (DLP) or modified DLPs. To test these basis functions, two 1-D test problems were developed: (1) a 10-pin representation of the junction between two heterogeneous fuel assemblies, and (2) a 70-pin representation of a boiling water reactor. Each of these problems utilized two cross-section libraries based on a 44-group and 238-group structure. Furthermore, a 2-D test problem based on the C5G7 benchmark is used to show applicability to higher dimensions.

Table of Contents

List of Figures	vi
List of Tables	xi
List of Abbreviations	xii
List of Symbols	xiii
List of Symbols	xiii
Acknowledgements	xiv
Dedication	xv
1 Introduction and Background	1
1.1 Motivation	2
1.2 Eigenvalue Response Matrix Method	4
1.2.1 Overview	4
1.2.2 Mathematical Formulation of ERMM	6
1.2.3 Projection onto a Space and Angle Subspace	8
1.2.4 Projection onto an Energy Group Subspace	9
1.2.5 Response Matrix Formalism	11
1.3 Objectives	12
2 Summary of Previous Work	13
2.1 Expanding in Energy	13
2.2 Discrete Legendre Polynomials	15
2.3 Modified Discrete Legendre Polynomials	18
2.4 Summary	21
3 Karhunen-Loève Transform	23
4 Test Problems and Models	30
4.1 Test Problems	30
4.1.1 1-D Test Problems	31
4.1.2 2-D Test Problems	34
4.2 Snapshot Generating Models	36

4.2.1	Snapshots for 10-Pin Test Problem	36
4.2.2	Snapshots for BWR Test Problem	37
4.2.3	Snapshots for C5G7 Test Problem	38
4.3	SERMENT and DETRAN	41
5	Results for 1-D Studies	42
5.1	mDLP Comparison	43
5.2	Energy Spectra for the Test Problems	44
5.3	10-Pin Test Problem	47
5.3.1	44-Group Results	47
5.3.2	238-Group Results	49
5.4	BWR Test Problem	51
5.4.1	Configuration 0	52
5.4.1.1	44-Group Results	52
5.4.1.2	238-Group Results	54
5.4.2	Configuration 1	55
5.4.2.1	44-Group Results	55
5.4.2.2	238-Group Results	57
5.4.3	Configuration 2	59
5.4.3.1	44-Group Results	59
5.4.3.2	238-Group Results	60
5.5	Higher-Order Angular Moments	62
5.5.1	10-pin Problem	63
5.5.1.1	44-Group Results	63
5.5.1.2	238-Group Results	64
5.6	BWR Test Problem	65
5.6.1	Configuration 0	66
5.6.1.1	44-Group Results	66
5.6.1.2	238-Group Results	66
5.6.2	Configuration 1	68
5.6.2.1	44-Group Results	68
5.6.2.2	238-Group Results	69
5.6.3	Configuration 2	70
5.6.3.1	44-Group Results	70
5.6.3.2	238-Group Results	71
5.7	Conclusion	73
6	Results for 2-D Studies	74
6.1	Energy Spectra for the Test Problem	75
6.2	Nodal Fission Density Results	75
6.3	Pin Power Results	77
6.4	Conclusion	82
7	Conclusions and Future Work	85
7.1	Summary	85

7.2	Future Work	88
	Bibliography	89
A	Parametric Studies	93
A.1	Snapshot Parameters	93
A.1.1	Sensitivity to Snapshot Selection	93
A.1.2	Sensitivity to Fineness of Spatial Mesh	95
A.2	Database Parameters	96
A.2.1	Sensitivity to Number of k Values in Databases	96
A.2.2	Sensitivity to Range of k Values in Databases	96

List of Figures

1.1	Example of 1-D ERMM node connections. The outgoing conditions J for a given node are the incident conditions on the adjacent node. Each node is computed individually, and the global solution is found by sweeping across all nodes.	5
2.1	Discrete Legendre Polynomials shown through 3^{rd} order using 40 points	16
2.2	Using 40 discrete points, $\cos(x)$ $x \in [0, 2\pi]$ expanded by Discrete Legendre Polynomials to 4^{th} order	18
2.3	Example shape vector for creating mDLP basis functions	19
2.4	Example basis functions for mDLP-1 shown through 3^{rd} order	20
2.5	Example basis functions for mDLP-2 shown through 3^{rd} order	20
2.6	Relative error in the L_2 norm as a function of expansion order for various basis sets	21
3.1	Basis functions for the KLT applied to an example problem	28
3.2	Relative error in the L_2 norm as a function of expansion order for KLT compared to other basis sets	29
4.1	Configuration for 10-pin Test Problem	32
4.2	Configuration for BWR Test Problem	33
4.3	Configuration for Full-Core. Each square represents the area of a 17×17 pin assembly	34
4.4	Configuration for UO_2 fuel bundle. The green represents a UO_2 pincell, while the blue represents a guide tube modeled as a pincell filled with moderator	35
4.5	Configuration for MOX bundle. The light red represents 4.3% MOX fuel, the medium red represents 7.0 % MOX fuel, and the dark red represents 8.7% MOX fuel. The blue represents moderator (i.e., light water)	35
4.6	Configuration for pincell. The circular fuel element had a radius of 0.54 cm and was homogenized with cladding for this model.	36
4.7	Configuration for pincell junction; all Reflect BC; Every combination is used	39
4.8	Configuration for small UO_2 fuel bundle. The green represents a UO_2 pincell, while the blue represents a guide tube modeled as a pincell filled with moderator	39

4.9	Configuration for small MOX bundle. The light red represents 4.3% MOX fuel, the medium red represents 7.0 % MOX fuel, and the dark red represents 8.7% MOX fuel. The blue represents moderator, light water	40
4.10	Configuration for 1D approximation to the C5G7 benchmark. The green represents UO ₂ , the light red represents 4.3% MOX fuel, the medium red represents 7.0 % MOX fuel, and the dark red represents 8.7% MOX fuel, and the blue represents moderator	40
5.1	Comparison of mDLP applied to the 10-pin problem using 44-group cross-section library.	44
5.2	Comparison of mDLP applied to the 10-pin problem using the 238-group cross-section library.	44
5.3	Flux spectrum for the 10-pin problem using 44-group cross-section library	45
5.4	Flux spectrum for the 10-pin problem using 238-group cross-section library	45
5.5	Flux spectrum for the BWR-Core 0 problem using 238-group cross-section library	46
5.6	Flux spectrum for the BWR-Core 1 problem using 238-group cross-section library	46
5.7	Flux spectrum for the BWR-Core 2 problem using 238-group cross-section library	47
5.8	Performance of the KLT when applied to the 10-pin test problem with snapshots of only ϕ	48
5.9	Performance of the KLT when applied to the 10-pin test problem with snapshots of only J_{left}	49
5.10	Performance of the KLT when applied to the 10-pin test problem with snapshots of both ϕ and leftward partial current.	49
5.11	Relative error for 238-group, 10-pin test problem using snapshots of only ϕ	50
5.12	Relative error for 238-group, 10-pin test problem using snapshots of only J_{left}	50
5.13	Relative error for 238-group, 10-pin test problem using snapshots of both ϕ and J_{left}	51
5.14	Relative error for the 44-group, BWR-Core 0 test problem using snapshots of only ϕ	52
5.15	Relative error for the 44-group, BWR-Core 0 test problem using snapshots of only J_{left}	53
5.16	Relative error for the 44-group, BWR-Core 0 test problem using snapshots of both ϕ and J_{left}	53
5.17	Relative error for 238-group, BWR-Core 0 test problem using snapshots of only ϕ	54
5.18	Relative error for 238-group, BWR-Core 0 test problem using snapshots of only J_{left}	55

5.19	Relative error for 238-group, BWR-Core 0 test problem using snapshots of both ϕ and J_{left}	55
5.20	Relative error for the 44-group, BWR-Core 1 test problem using snapshots of only ϕ	56
5.21	Relative error for the 44-group, BWR-Core 1 test problem using snapshots of only J_{left}	56
5.22	Relative error for the 44-group, BWR-Core 1 test problem using snapshots of both ϕ and J_{left}	57
5.23	Relative error for 238-group, BWR-Core 1 test problem using snapshots of only ϕ	58
5.24	Relative error for 238-group, BWR-Core 1 test problem using snapshots of only J_{left}	58
5.25	Relative error for 238-group, BWR-Core 1 test problem using snapshots of both ϕ and J_{left}	59
5.26	Relative error for the 44-group, BWR-Core 2 test problem using snapshots of only ϕ	60
5.27	Relative error for the 44-group, BWR-Core 2 test problem using snapshots of only J_{left}	60
5.28	Relative error for the 44-group, BWR-Core 2 test problem using snapshots of both ϕ and J_{left}	61
5.29	Relative error for 238-group, BWR-Core 2 test problem using snapshots of only ϕ	61
5.30	Relative error for 238-group, BWR-Core 2 test problem using snapshots of only J_{left}	62
5.31	Relative error for 238-group, BWR-Core 2 test problem using snapshots of both ϕ and J_{left}	62
5.32	Relative error for 44-group, 10-pin test problem using snapshots from the 10-pin model. Sets of snapshots are used separately for basis generation	63
5.33	Relative error for 44-group, 10-pin test problem using snapshots from the 10-pin model. Sets of snapshots are combined together for basis generation	64
5.34	Relative error for 238-group, 10-pin test problem using snapshots from the 10-pin model. Sets of snapshots are used separately for basis generation	65
5.35	Relative error for 238-group, 10-pin test problem using snapshots from the 10-pin model. Sets of snapshots are combined together for basis generation	65
5.36	Relative error for 44-group, BWR-Core 0 test problem using snapshots from the Full-Core model. Sets of snapshots are used separately for basis generation	66
5.37	Relative error for 44-group, BWR-Core 0 test problem using snapshots from the Full-Core model. Sets of snapshots are combined together for basis generation	67

5.38	Relative error for 238-group, BWR-Core 0 test problem using snapshots from the Full-Core model. Sets of snapshots are used separately for basis generation	67
5.39	Relative error for 238-group, BWR-Core 0 test problem using snapshots from the Full-Core model. Sets of snapshots are combined together for basis generation	68
5.40	Relative error for 44-group, BWR-Core 1 test problem using snapshots from the Full-Core model. Sets of snapshots are used separately for basis generation	68
5.41	Relative error for 44-group, BWR-Core 1 test problem using snapshots from the Full-Core model. Sets of snapshots are combined together for basis generation	69
5.42	Relative error for 238-group, BWR-Core 1 test problem using snapshots from the Full-Core model. Sets of snapshots are used separately for basis generation	70
5.43	Relative error for 238-group, BWR-Core 1 test problem using snapshots from the Full-Core model. Sets of snapshots are combined together for basis generation	70
5.44	Relative error for 44-group, BWR-Core 2 test problem using snapshots from the Full-Core model. Sets of snapshots are used separately for basis generation	71
5.45	Relative error for 44-group, BWR-Core 2 test problem using snapshots from the Full-Core model. Sets of snapshots are combined together for basis generation	71
5.46	Relative error for 238-group, BWR-Core 2 test problem using snapshots from the Full-Core model. Sets of snapshots are used separately for basis generation	72
5.47	Relative error for 238-group, BWR-Core 2 test problem using snapshots from the Full-Core model. Sets of snapshots are combined together for basis generation	72
6.1	Flux spectrum for the C5G7 problem using 44-group cross-section library	75
6.2	Relative error in fission density for 44-group, C5G7 test problem using snapshot of only ϕ	76
6.3	Relative error in fission density for 44-group, C5G7 test problem using snapshot of J_{up} and J_{down}	77
6.4	Relative error in fission density for 44-group, C5G7 test problem using snapshot of ϕ , J_{up} , and J_{down}	78
6.5	Relative error in pin power for 44-group, C5G7 test problem using snapshot of only ϕ	79
6.6	Relative error in pin power for 44-group, C5G7 test problem using snapshot of both J_{up} and J_{down}	79
6.7	Relative error in pin power for 44-group, C5G7 test problem using snapshot of both ϕ , J_{up} , and J_{down}	80

6.8	Pin power heat map for the DETRAN reference solution. The upper left corner is the center of the core	80
6.9	Error of pin powers in the SERMENT reference solution relative to the DETRAN reference solution. The upper left corner is the center of the core	81
6.10	Error in the pin powers of the best performing KLT case (9th order, Reduced Small-Core, snapshots of ϕ) relative to the SERMENT reference solution. The upper left corner is the center of the core	82
6.11	Error in the pin powers of the best performing KLT case (9th order, Reduced Small-Core, snapshots of J_{up} and J_{down}) relative to the SERMENT reference solution. The upper left corner is the center of the core	83
6.12	Error in the pin powers of the best performing KLT case (9th order, Reduced Small-Core, snapshots of ϕ , J_{up} , and J_{down}) relative to the SERMENT reference solution. The upper left corner is the center of the core	84
A.1	Relative error for 238-group, BWR test problem when using various snapshot selection schemes. The schemes converge as the number of snapshots increases. The legend numbers correspond to the number of snapshots selected for KLT basis generation.	94
A.2	Relative error for 44-group, 10-pin test problem using snapshots from a spatially reduced model. The legend numbers correspond to the total number of available snapshots. The problem was reduced from 280 total snapshots.	95
A.3	Relative error for 238-group, BWR test problem a database of responses filled with a number of k values as indicated by the legend. The values spanned the range of ± 0.15 of the the true value of k . The number of k values had no effect of the results.	97
A.4	Relative error for 238-group, BWR test problem a database of responses filled with eight k values. The values spanned the range of \pm the value in the legend of the the true value of k . The size of the range no effect of the results.	97

List of Tables

4.1	Summary snapshot models for 10-pin Test Problem	37
4.2	Summary of snapshot models for BWR Test Problem	38
4.3	Summary of snapshot models for C5G7 test Problem	39

List of Abbreviations

ERMM	E igenvalue R esponse M atrix M ethod
RMM	R esponse M atrix M ethod
DLP	D iscrete L egendre P olynomials
mDLP	m odified D iscrete L egendre P olynomials
KLT	K arhunen- L oève T ransform
POD	P roper- O rthogonal D ecomposition
PCA	P rincipal C omponent A nalysis
SVD	S ingular V alue D ecomposition
BWR	B oiling W ater R eactor
PWR	P ressurized W ater R eactor
MOX	M ixed O Xide fuel
SERMENT	S olving E igenvalue R esponse M atrix E quations using N onlinear T echniques
DETRAN	DE terministic TRAN sport

List of Symbols

∇	del operator	
\mathbf{r}	position vector	cm
g	energy group	
k	eigenvalue	
\mathcal{T}	operator for transport processes	
\mathcal{F}	operator for neutron generation	
$J_{+/-}$	Current density in $+/-$ direction	$\text{cm}^{-1}\text{s}^{-1}$
\mathbf{j}	angular current	$\text{cm}^{-1}\text{s}^{-1}\text{sr}^{-1}$
P^m	m th vector of basis set P	
δ_{mn}	Kronecker delta	
\mathbf{R}	matrix of responses	
\mathbf{n}	vector normal to surface	
ν	neutrons produced per fission	
$\boldsymbol{\rho}$	phase-space vector	
Σ	macroscopic cross section	cm^{-1}
ϕ	scalar flux	$\text{cm}^{-1}\text{s}^{-1}$
χ	PDF for fission emission energy	
ψ	angular flux	$\text{cm}^{-1}\text{s}^{-1}\text{sr}^{-1}$
$\hat{\Omega}$	solid angle vector	sr

Acknowledgements

I would take this time to thank just a few of those that have had a hand in shaping this work. First among those would be my advisor Professor Jeremy Roberts. His support and guidance has been instrumental in bettering myself and my work.

In addition, I owe many thanks to Professor Larry Erickson, who has been a steadfast mentor and advisor throughout the years.

I would like to thank Professors Ken Shultis and Hitesh Bindra for serving on my thesis committee and for always making time for any of my questions.

I would also like to thank my colleagues, specifically Lyric, Kevin and Richard, for providing advice, encouragement, and much more.

For the funding for this effort, I acknowledge the Kansas State University Nuclear Research Fellowship Program, which is generously sponsored by the U.S. Nuclear Regulatory Commission.

Dedication

“Engineering is the art of modeling materials we do not wholly understand, into shapes we cannot precisely analyze so as to withstand forces we cannot properly assess, in such a way that the public has no reason to suspect the extent of our ignorance.”

A. R. Dyke

Chapter 1

Introduction and Background

Nuclear reactors are extremely complicated systems, and modeling their dynamic behavior requires the solution of the neutron transport equation, a seven-dimensional $(\mathbf{r}, \mathbf{v}, t)$ equation that describes the population of neutrons in the system. Unfortunately, to solve the complete neutron transport equation is all but impossible for realistic problems without first reducing the problem space in some way. For much of reactor analysis, it is reasonable to assume steady-state conditions, thus eliminating all time dependence from the problem. In addition, it is very difficult to account for the energy dependence directly, and a common approach is to use the multigroup method, in which the energy variable is discretized into G energy bins or “groups”. After these assumptions and the further assumption of isotropic scattering, the neutron transport equation reduces to

$$\hat{\Omega} \cdot \nabla \psi_g(\hat{\Omega}, \mathbf{r}) + \Sigma_{t,g} \psi_g(\hat{\Omega}, \mathbf{r}) = \frac{1}{4\pi} \left[\sum_{g'=1}^G \Sigma_{s,g' \rightarrow g} \phi_{g'}(\mathbf{r}) + \frac{\chi_g}{k} \sum_{g'=1}^G \nu \Sigma_{f,g'} \phi_{g'}(\mathbf{r}) \right], \quad (1.1)$$

where ψ_g represents the group-dependent angular flux and ϕ_g is the group-dependent scalar flux. Furthermore, $\Sigma_{t,g}$, $\Sigma_{s,g' \rightarrow g}$, and $\Sigma_{f,g}$ represent the group dependent cross sections for total, in-scattering, and fission respectively. In addition, χ_g is the fission spectrum, ν is the average number of neutrons emitted per fission, and the k -eigenvalue (or “multiplication factor”) represents the balance of neutron gains

(by fission) to losses (by absorption and leakage). A reactor is “critical,” i.e., the neutron population is steady, when $k = 1$.

Equation (1.1) may be used with any number of groups for which cross-section data is available. In any group structure, a system of G equations must be solved. Clearly, a greater number of groups will increase the difficulty of the problem, but it will also increase the accuracy, provided that the group structure is chosen wisely.

Equation (1.1) may be cast in operator notation, i.e.,

$$\mathcal{T}\phi(\boldsymbol{\rho}) = \frac{1}{k}\mathcal{F}\phi(\boldsymbol{\rho}), \quad (1.2)$$

where \mathcal{T} represents the transport processes, ϕ is the neutron flux, \mathcal{F} represents the neutron generation, $\boldsymbol{\rho}$ represents the relevant phase-space, and k is the eigenvalue.

A variety of approaches have been used to solve Eq. (1.2), which may broadly be categorized as deterministic (e.g., discrete ordinates) methods and stochastic (e.g., Monte Carlo) methods. This work has used a deterministic method exclusively, although much of the theory presented should also apply to stochastic methods.

1.1 Motivation

To model a full reactor core with high-fidelity resolution in space, angle, and energy requires an enormous amount of computational power and memory. To illustrate the challenge, consider a typical pressurized water reactor (PWR) core with 193 assemblies, each with a 17×17 array of fuel pins. Some reasonable parameters for resolving the localized pin powers are approximately 50 spatial cells in the x-y plane and 300 axial mesh points (i.e., 1 cm axial resolution) per pincell. Localized pin power resolution in energy requires approximately 100 groups, and similarly, to resolve the localized pin powers in angle requires approximately 100 angles. For discretization with just one unknown per cell, group, and angle, the total number

of unknowns is

$$\begin{aligned}
 N &= (193 \text{ assemblies}) \times (17^2 \frac{\text{pins}}{\text{assembly}}) \times (300 \text{ axial points}) \\
 &\quad \times (50 \text{ radial points}) \times (100 \text{ groups}) \times (100 \text{ angles})^2 \\
 &\approx 1 \times 10^{14},
 \end{aligned}
 \tag{1.3}$$

for a single problem. For realistic analyses, a problem of this size would need to be solved repeatedly to account for thermal-hydraulic feedback effects or to model the change of material compositions over time. For such realistic problems, each spatial cell can be assumed to have a unique material, each of which is defined by $O(100^2)$ floating point values. Hence, as shown, a large problem can quickly amass too many unknowns to solve in a reasonable amount of time on modern computers as storage requirements quickly rise above PByte levels. While the problems can be solved in principle, the solution requires too much time for production-level analyses, i.e., those for design of actual systems. Thus, the number of degrees of freedom must be reduced by reducing the spatial, angular, or energy resolution.

There are, of course, repercussions for each type of reduction, but the goal is to minimize the effect of the reduction on the final solution of the neutron transport equation. If the problem is reduced spatially, we are no longer solving the same problem, while if energy or angle is reduced the fidelity of the model is reduced because some of the underlying physics is muted. In lattice physics, a complete solution is attained by solving the problem of interest several times to identify the dependence of each phase-space variable, then combine the separate solutions in such a way as to predict the global solution with some accuracy.

It is common in lattice physics to first solve the problem with a 0-D representation in space and angle while continuous in energy to solve for the energy dependence. Next, a lattice physics solution will reduce the energy dependence to the multigroup approximation while using a 1-D or 2-D mesh in space along with high angular resolution to capture the angular dependence of the problem. These solutions are used to create approximate multigroup cross sections to be used in conjunction

with models with only a few energy groups, a coarse angular dependence, and a coarse 3-D spatial mesh.

Many traditional methods for solving the neutron transport equation focus on a quickly computed solution by approximating the energy dependence, which simplifies the problem and in turn reduces the fidelity of the solution. However, in lattice physics, upwards of 200 energy groups or more are used to solve for a high-fidelity energy spectra. As such, a method that can bridge the gap between the quick, low-fidelity methods and the accurate, high-fidelity methods would be immensely valuable to the future of large-scale, reactor-physics simulations and analyses. Not only would the model be appropriately sized for modern computers, but it could also be more accurate than the current methods of whole core-analysis.

1.2 Eigenvalue Response Matrix Method

1.2.1 Overview

This work applies the eigenvalue response matrix method (ERMM). Response matrix methods are not new, and have been used in various formulations since the 1960's, e.g., the work of Shimizu [1] and Shimizu et al. [2]. ERMM solves the reactor eigenvalue equation by decomposing the domain of a problem into independent nodes linked through approximate boundary conditions between each node. The boundary conditions are typically incident angular flux or current conditions. For ERMM, these conditions are represented by truncated, orthogonal basis expansions at the nodal boundaries. This approach effectively converts a large problem space into a large number of small, independent, transport problems, which creates many opportunities for parallelization of the algorithm.

An example of ERMM is depicted in Fig. 1.1, where a sample problem is broken into M nodes. The m th node, where $1 \leq m \leq M$, can be solved independently. In this way, the outgoing information (here denoted by J) from one cell becomes the incident conditions for the adjacent cell. The problem is then solved by assuming a

value for each of the initial incident boundary information. The cells are then solved for the outgoing information, which then becomes the new incident information for another solution. In this way, the method will continue until convergence criteria are met.




FIGURE 1.1: Example of 1-D ERMM node connections. The outgoing conditions J for a given node are the incident conditions on the adjacent node. Each node is computed individually, and the global solution is found by sweeping across all nodes.

In this case, the outgoing information is a boundary flux or current. The key to response matrix methods is that the outgoing conditions of a cell are not passed completely, but rather are projected onto a finite, orthogonal basis, and the coefficients of the resulting expansion become unknowns for the problem space. This projection reduces the number of degrees of freedom and, thus, reduces the size of the problem. An expansion is completed for each phase-space variable (i.e., space, angle, and energy), and the success of these expansions depends primarily on the selection of appropriate orthogonal bases for each phase-space variable. Response matrix methods are expensive in general, unless the problem space is greatly reduced in terms of the degrees of freedom as compared to more direct solutions. Thus, proper basis selection is paramount if ERMM is to compare in speed and accuracy to other methods for solving the transport equation.

In short, basis sets that capture the detail of a high-fidelity transport solution with low-order expansions are ideal for ERMM because fewer degrees of freedom are needed to achieve the desired accuracy, which reduces the problem space, and leads to an easier problem to solve computationally. This work builds on the previous effort presented by Roberts and Forget [3], which explored spatial and angular expansions in ERMM. This work extends their progress by utilizing a new, highly successful energy basis for use in ERMM based on the Karhunen-Loève transform,

which is discussed in detail in Chapter 3. These basis sets are applied to several illustrative problems, which provide insight to the success of the method.

It is possible (and typical) to use different basis expansions for each variable; thus the expansion for each variable can be studied mostly independently of the other variables. However, there exists coupling between each phase-space variable. For example, the error caused by the expansion in energy is influenced by the expansion in space. This work used reference cases designed to isolate the effect of only one variable to the best extent possible. These reference cases are described in Chapter 4.

1.2.2 Mathematical Formulation of ERMM

The following is a presentation of the time-independent, eigenvalue response matrix method. Although the following method is most closely linked to the work of Roberts and Forget [3], a presentation akin to that of Roberts [4] has been adapted for this work. For time-independent, multigroup problems, Eq. (1.2) may be rewritten as

$$\mathcal{T}\psi^{\text{global}}(\mathbf{r}, \boldsymbol{\Omega}, g) = \frac{1}{k} \mathcal{F}\psi^{\text{global}}(\mathbf{r}, \boldsymbol{\Omega}, g), \quad (1.4)$$

where ρ has been replaced by the spatial coordinate \mathbf{r} , the direction of travel $\boldsymbol{\Omega}$, and the energy group g .

The “global” superscript in Eq. (1.4) indicates that the equation defines balance for the full problem of interest, and thus is the complete or global solution. ERMM works by breaking the global problem into small, independent or local nodes. To do so, let the global volume V be decomposed into I disjoint, convex, nodal sub-volumes V_i that satisfy $V = V_1 \cup V_2 \cup \dots \cup V_I$. Furthermore, let the surface ∂V_i of V_i be composed of S disjoint, planar surfaces ∂V_{is} that satisfy $\partial V_i = \partial V_{i1} \cup \dots \cup \partial V_{iS}$. While unnecessary, the condition of planar surfaces permits a compact notation. Finally, let \mathbf{r}_i and \mathbf{r}_{is} be shorthand for the variable \mathbf{r} confined to values $\mathbf{r} \in V_i$ and $\mathbf{r} \in \partial V_{is}$, respectively.

Thus, the equivalent “local” transport equation for the i th node is

$$\mathcal{T}\psi^{\text{local}}(\mathbf{r}_i, \boldsymbol{\Omega}, g) = \frac{1}{k}\mathcal{F}\psi^{\text{local}}(\mathbf{r}_i, \boldsymbol{\Omega}, g), \quad (1.5)$$

subject to S incident-flux boundary conditions

$$\psi^{\text{local}}(\mathbf{r}_{is}, \boldsymbol{\Omega}, g) = \psi^{\text{global}}(\mathbf{r}_{is}, \boldsymbol{\Omega}, g), \quad \boldsymbol{\Omega} \cdot \mathbf{n}_{is} < 0, \quad (1.6)$$

where the unit vector \mathbf{n}_{is} is the outward normal of surface ∂V_{is} . Equations (1.5) and (1.6) can be combined by casting incident-flux conditions as external sources such as,

$$\left(\mathcal{T} - \frac{1}{k}\mathcal{F}\right)\psi^{\text{local}}(\mathbf{r}_i, \boldsymbol{\Omega}, g) = j^{\text{global}}(\mathbf{r}_i, \boldsymbol{\Omega}, g)\delta(\mathbf{r}_i \cdot \mathbf{n}_{is}), \quad (1.7)$$

where $\mathbf{j} = \boldsymbol{\Omega}\psi$ is the angular current vector, and j is the magnitude of the angular current. For brevity, j is referred to as the angular current. This formulation allows the global problem to be formed by sweeping across each local node.

The general solution of Eq. (1.7) for arbitrary incident conditions can be expressed as the convolution of the external source term with an appropriate kernel, or

$$\psi(\mathbf{r}_i, \boldsymbol{\Omega}, g) = \sum_{g'=1}^G \sum_{s'=1}^S \int_{\mathbf{n}_{is'} \cdot \boldsymbol{\Omega}' < 0} R_f(\mathbf{r}'_{is'}, \boldsymbol{\Omega}', g' \rightarrow \mathbf{r}_i, \boldsymbol{\Omega}, g) j(\mathbf{r}'_{is'}, \boldsymbol{\Omega}', g') d\Omega, \quad (1.8)$$

where G is the number of groups, and the “local” and “global” superscripts have been omitted. Similarly, exiting angular currents can be expressed as

$$j^+(\mathbf{r}_{is}, \boldsymbol{\Omega}, g) = \sum_{g'=1}^G \sum_{s'=1}^S \int_{\mathbf{n}_{is'} \cdot \boldsymbol{\Omega}' < 0} R_c(\mathbf{r}'_{is'}, \boldsymbol{\Omega}', g' \rightarrow \mathbf{r}_{is}, \boldsymbol{\Omega}, g) j(\mathbf{r}'_{is'}, \boldsymbol{\Omega}', g') d\Omega, \quad (1.9)$$

where $\mathbf{n}_{is} \cdot \boldsymbol{\Omega} > 0$. The integration kernels R_f and R_c are flux- and current-response functions, which represent the angular flux and outgoing angular current at one point in phase-space due to a unit, incident current at another point in phase-space. However, some effort is required to convert Eqs. (1.8) and (1.9) into

a practical form. The goal is to reduce the effort needed to compute the global solution; thus, some approximations should be made for each phase-space variable.

1.2.3 Projection onto a Space and Angle Subspace

The first approximations are the treatment of the angular and spatial dependence. Projection of local angular currents and fluxes onto a finite subspace represented by an orthogonal basis allows the use of a truncated basis to represent space and angle. Using a truncated basis reduces the problem space, and, hence, the computation cost of ERMM at the expense of reduced accuracy. Let a finite basis be constructed with a set of functions $P^m(\mathbf{r}, \boldsymbol{\Omega})$, $m = 0, 1, \dots, M$, that are orthonormal over some domain of interest (i.e., a volume or a surface). Then the angular flux can be approximated as

$$\psi(\mathbf{r}_i, \boldsymbol{\Omega}, g) \approx \sum_{m=0}^M \psi_i^m(g) P_f^m(\mathbf{r}_i, \boldsymbol{\Omega}), \quad (1.10)$$

where

$$\psi_i^m(g) = \int_{V_i} \int_{4\pi} \psi(\mathbf{r}_i, \boldsymbol{\Omega}, g) P_f^m(\mathbf{r}_i, \boldsymbol{\Omega}) d\Omega d^3r_i, \quad (1.11)$$

and the f subscript denotes a basis suitable for the angular flux. Angular currents can similarly be approximated as

$$j^\pm(\mathbf{r}_{is}, \boldsymbol{\Omega}, g) \approx \sum_{l=0}^L j_{is}^{\pm l}(g) P_c^l(\mathbf{r}_{is}, \boldsymbol{\Omega}), \quad \mathbf{n}_{is} \cdot \boldsymbol{\Omega} \gtrless 0, \quad (1.12)$$

where

$$j_{is}^{\pm l}(g) = \int_{\partial V_{is}} \int_{\mathbf{n}_{is} \cdot \boldsymbol{\Omega} \gtrless 0} j(\mathbf{r}_{is}, \boldsymbol{\Omega}, g) P_c^l(\mathbf{r}_{is}, \boldsymbol{\Omega}) d\Omega d^2r_{is}, \quad (1.13)$$

and the c subscript denotes a basis suitable for the angular current defined on a boundary surface. Then, substitution of Eqs. (1.10) and (1.12) into Eq. (1.8) yields

$$\sum_{m=0}^M \psi_i^m(g) P_f^m(\mathbf{r}_i, \boldsymbol{\Omega}) \approx \sum_{g'=1}^G \sum_{s'=1}^S \sum_{l'=0}^L j_{is'}^{-l'}(g') \langle R_f, P_c^{l'} \rangle, \quad (1.14)$$

where variables have been suppressed and $\langle \cdot \rangle$ indicates the appropriate space and angle integration. Multiplication of Eq. (1.14) by $P_f^m(\mathbf{r}_i, \boldsymbol{\Omega})$ and integration of the result over space and angle leads to a set of flux moments defined by

$$\psi_i^m(g) \approx \sum_{g'=1}^G \sum_{s'=1}^S \sum_{l'=0}^L j_{is'}^{-l'}(g') R_{fi}^{s'l' \rightarrow m}(g' \rightarrow g), \quad (1.15)$$

where

$$R_{fi}^{s'l' \rightarrow m}(g' \rightarrow g) \equiv \langle \langle R_f, P_f^{l'} \rangle, P_f^m \rangle. \quad (1.16)$$

Similarly, the outgoing angular currents can also be projected to yield the moments

$$j_{is}^{+l}(g) \approx \sum_{g'=1}^G \sum_{s'=1}^S \sum_{l'=0}^L j_{is'}^{-l'}(g') R_{ci}^{s'l' \rightarrow sl}(g' \rightarrow g). \quad (1.17)$$

By choosing appropriate basis sets for expanding the spatial and angular variables, few terms are required for sufficient accuracy. For this work, Jacobi polynomials were used for the angular expansion, and Discrete Legendre Polynomials were used for the spatial expansion. The work did not focus on optimization of the spatial and angular basis functions, but rather the best cases from the work of Roberts and Forget [3] were used.

1.2.4 Projection onto an Energy Group Subspace

This work focused on the basis sets for energy expansion, and the projection for energy is formulated similarly to the space-angle projection. This treatment will eliminate an explicit dependence on g . Because g is a discrete variable, bases used to represent dependence on g consist of discrete functions $P^h(g)$, $h = 0, 1, \dots, H$ that satisfy

$$\sum_g^G P^h(g) P^{h'}(g) = \delta_{hh'}, \quad (1.18)$$

where $\delta_{hh'}$ is the Kronecker- δ . By using such a discrete basis, group-dependent flux moments defined by Eq. (1.15) can be approximated as

$$\psi_i^m(g) \approx \sum_{h=0}^H \psi_i^{mh} P^h(g), \quad (1.19)$$

where

$$\psi_i^{mh} = \sum_{g=1}^G \psi_i^m(g) P^h(g). \quad (1.20)$$

Likewise, current moments defined by Eq. (1.17) can be approximated as

$$j_{is}^l(g) \approx \sum_{h=0}^H j_{is}^{lh} P^h(g), \quad (1.21)$$

where

$$j_{is}^{lh} = \sum_{g=1}^G j_{is}^l(g) P^h(g). \quad (1.22)$$

Substitution of Eqs. (1.19) and (1.21) into Eq. (1.15) yields

$$\sum_{h=0}^H \psi_i^{mh} P^h(g) \approx \sum_{g'=1}^G \sum_{s'=1}^S \sum_{l'=0}^L \sum_{h'=0}^H j_{is'}^{-l'h'}(g') \langle \langle R_f, P_f^{l'} \rangle, P_f^m \rangle. \quad (1.23)$$

Next, multiplication of the result by $P^h(g)$, and summation over energy leads to

$$\psi_i^{mh} \approx \sum_{s'=1}^S \sum_{l'=0}^L \sum_{h'=0}^H j_{is'}^{-l'h'} R_{fi}^{s'l'h' \rightarrow mh}, \quad (1.24)$$

where

$$R_{fi}^{s'l'h' \rightarrow mh} \equiv \langle \langle \langle R_f, P_f^{l'} \rangle, P_f^m \rangle, P^h \rangle, \quad (1.25)$$

and the outer brackets represent summation, not integration over g . Similarly, current moments are defined as

$$j_{is}^{+lh} \approx \sum_{s'=1}^S \sum_{l'=0}^L \sum_{h'=0}^H j_{is'}^{-l'h'} R_{ci}^{s'l'h' \rightarrow slh}, \quad (1.26)$$

where

$$R_{ci}^{s'l'h' \rightarrow slh} \equiv \langle \langle \langle R_c, P_c^{l'} \rangle, P_c^l \rangle, P^h \rangle. \quad (1.27)$$

Computation of response function moments $R_{fi}^{s'l'h' \rightarrow mh}$ and $R_{ci}^{s'l'h' \rightarrow slh}$ requires evaluation of Eqs. (1.8) and (1.9) in which the incident current is equal to $P_c^{l'}(\mathbf{r}_{is'}, \boldsymbol{\Omega})P^{h'}(g)$ for each allowed combination of i , s , l , and h .

1.2.5 Response Matrix Formalism

Equations (1.24) and (1.26) can be represented as nodal response matrix equations, i.e.,

$$\begin{aligned}\boldsymbol{\psi}_i &= \mathbf{R}_{fi}\mathbf{j}_i^- \\ \mathbf{j}_i^+ &= \mathbf{R}_{ci}\mathbf{j}_i^-, \end{aligned}\tag{1.28}$$

where $\boldsymbol{\psi}_i$ and \mathbf{j}_i^\pm are vectors of nodal moments and \mathbf{R}_i 's are matrices of nodal response function moments. Response matrix equations for the entire spatial domain can then be written as

$$\begin{aligned}\boldsymbol{\psi} &= \mathbf{R}_f\mathbf{j}^- \\ \mathbf{j}^+ &= \mathbf{R}_c\mathbf{j}^- . \end{aligned}\tag{1.29}$$

By redirecting outgoing currents from one node as incident currents to another via $\mathbf{j}^- = \mathbf{M}\mathbf{j}^+$, where the matrix \mathbf{M} represents geometry and boundary conditions, the global equations become

$$\begin{aligned}\boldsymbol{\psi} &= \mathbf{R}_f\mathbf{j}^- \\ \mathbf{j}^- &= \mathbf{M}\mathbf{R}_c\mathbf{j}^- . \end{aligned}\tag{1.30}$$

In this formulation, the flux is dependent only on incident currents, and the response matrices \mathbf{R}_f and \mathbf{R}_c are functions of the k -eigenvalue. When k is not converged, the balance defined by Eq. (1.30) generally does not have a solution. As an alternative, the current equation can be rewritten as the nonlinear eigenvalue equation

$$\mathbf{M}\mathbf{R}_c(k)\mathbf{j}^- = \lambda\mathbf{j}^- .\tag{1.31}$$

After determining \mathbf{j}^- by solving Eq. (1.31) for an assumed value of k , the flux and other volume moments can be determined. Subsequently, a new value for k is generated similarly to the traditional power method by using the standard balance relation of gains-to-losses. The new k is then used to find \mathbf{j}^- , etc., until the solution has converged. A more detailed presentation of algorithms for solving the response matrix equations is given by Roberts and Forget [3].

1.3 Objectives

The primary focus of this work is to reduce by an order of magnitude the number of energy degrees of freedom needed to achieve sub-0.1% error or better in the relative fission density error or pin powers. The fission density is directly related to the power levels throughout the problem. To evaluate the method, several test problems were devised to test the energy expansion including two 1-D problems and one 2-D problem. The project tested several facets of the orthogonal basis sets produced by the Karhunen-Loève Transform (KLT), which was used for expansion in energy.

The KLT uses snapshots (discussed in detail in Chapter 4) to generate the basis sets. Thus, the work has considered several different models from which to generate snapshots. Each unique set of basis functions is used to expand the energy dependence for the test problem, and results are generated as the relative error in the fission density for the expansion as a function of order. The results of the 1-D test problems are presented in Chapter 5, while the results of the 2-D test problem are presented in Chapter 6. During the course of the project, several parametric studies were developed to test the impact of different aspects of the models. These parametric studies are presented in Appendix A.

Chapter 2

Summary of Previous Work

Response matrix methods have traditionally used a full multigroup representation of the energy variable, which is to say that there is no basis truncation for the energy basis. Response matrix methods have been used successfully with a variety of energy group structures, ranging from three to 190 groups [5–7].

One of the first studies of a truncated energy expansion for response matrix methods was the work of Roberts [8], who used Discrete Legendre Polynomials (DLP) and a modified version of DLP (mDLP) to expand in energy. His approach compared the basis expansions to a full multigroup approximation of the energy variable.

This chapter aims to explore the previous work in expanding the energy variable, beginning with a brief overview of how the full multigroup approach can be represented in the formalism of Chapter 1, followed by a discussion of DLPs and problem-specific, modified DLPs.

2.1 Expanding in Energy

In order to represent the multigroup method exactly in ERMM, the energy variable is represented by a set of Kronecker- δ functions defined by

$$P_{\delta}^h(g) = \delta_{h,g-1}, g = 1, 2, \dots, G, \quad (2.1)$$

where

$$\delta_{h,g-1} = \begin{cases} 1, & \text{if } h = g - 1, \\ 0, & \text{if } h \neq g - 1. \end{cases} \quad (2.2)$$

When a complete set of these vectors (i.e., $H = G - 1$) is used, a generic response function moment $R^{s'l'h' \rightarrow slh}$ can be rewritten as $R^{s'l'g' \rightarrow slg}$. Thus, the group-to-group transfer process is represented in the traditional form. However, this approach requires that all energy groups be included, thus the number of energy degrees of freedom is equal to the number of groups. This method works well when the number of groups is low, but large models become prohibitively expensive when a detailed energy treatment (i.e., more than approximately 10 groups) is used.

If a Kronecker set is truncated, a significant amount of the physics is lost, and the expansion has significantly reduced accuracy. Hence, the Kronecker set should not be used with energy order reduction. In order to improve ERMM performance, a basis set for energy is sought that will capture many-group fidelity while requiring many fewer degrees of freedom (H in Eqs. (1.24) and (1.26)) than a full multigroup analysis (a solution utilizing the complete set of Kronecker vectors).

A way to reduce the energy degrees of freedom is to use an orthogonal basis to approximate the functionality of every group, thus converting the energy degrees of freedom for the problem into the coefficients of expansion. In this case, the energy degrees of freedom are equal in number to G , because $H + 1$ basis functions are required to expand exactly a vector of length G . The problem is then simplified by using a lower expansion order than G . If the low-order vectors in a basis set can better approximate the energy dependence, then the truncation should minimize effect of the error introduced such an expansion. The previous work by Roberts [8] suggested that incorporating physics directly into the basis functions may be more efficient than standard basis sets, and can lead to accuracy close to that of a full multigroup treatment using G groups without needing $H + 1$ energy degrees of freedom.

Each of the basis sets presented here are computed in orthonormal form. This formulation leads to quick determination of the expansion coefficients similarly to Eqs. (1.20) and (1.22), rewritten as

$$a_i = \sum_{g=0}^G w(g) P^i(g) f(g), \quad (2.3)$$

where a_i is the i th coefficient of expansion, $w(g)$ is the weight associated with the basis function, $P^i(g)$ is the i th basis function, and $f(g)$ is the function to expand. The reconstructed function is given by

$$\tilde{f}(g) \approx \sum_{i=0}^I a_i P^i(g) \quad (2.4)$$

where I is the expansion order.

Equation (2.4) is defined akin to Eqs. (1.19) and (1.21). Because a reduced number of energy degrees of freedom is desired, a basis that incorporates some physics of the problem is likely to provide the best expansion. The basis that includes physics should then be compared to more standard basis sets, such as the Discrete Legendre Polynomials (DLPs).

2.2 Discrete Legendre Polynomials

The Legendre polynomials form a standard and proven basis that is used throughout computational physics. The discrete versions of the Legendre polynomials are shown in Fig. 2.1. When using the DLPs, the functions are represented by vectors of length equal to the number of discretized points. The vectors shown in Fig. 2.1 have been orthonormalized over the range of the vectors. Recent work investigated the use of discrete Legendre polynomials (DLPs) for expansion in energy [8–10]. The set of polynomials is generated using the Gram-Schmidt process to orthogonalize the discrete monomials $M^h(g) = g^h$, $g = 0, 1, \dots, G - 1$. To illustrate, let $G = 5$,




FIGURE 2.1: Discrete Legendre Polynomials shown through 3rd order using 40 points

for which the zeroth-order DLP vector is defined as

$$\begin{aligned}
 P_{\text{DLP}}^0(\cdot) &= \frac{M^0(\cdot)}{\sqrt{\sum_{g=1}^G M^0(g)}} \\
 &= \frac{\sqrt{5}}{5} [1, 1, 1, 1, 1]^\top.
 \end{aligned} \tag{2.5}$$

To define the first-order DLP vector, let

$$\tilde{P}_{\text{DLP}}^1(\cdot) = M^1(\cdot) - \left(\sum_{g=1}^G P_{\text{DLP}}^0(g) M^1(g) \right) P_{\text{DLP}}^0(\cdot), \tag{2.6}$$

leading to

$$\begin{aligned}
 P_{\text{DLP}}^1(\cdot) &= \frac{\tilde{P}_{\text{DLP}}^1(\cdot)}{\sqrt{\sum_{g=1}^G \tilde{P}_{\text{DLP}}^1(g) \tilde{P}_{\text{DLP}}^1(g)}} \\
 &= \frac{\sqrt{10}}{5} [-2, -1, 0, 1, 2]^\top.
 \end{aligned} \tag{2.7}$$

This procedure can be repeated for DLP vectors of arbitrary order, up through one less than number of energy groups used. For the provided example, the second-,

third-, and fourth-order DLP vectors are defined as

$$\begin{aligned}
 P_{\text{DLP}}^2(\cdot) &= \frac{\sqrt{146}}{146} [8, -1, -4, -1, 8]^\top \\
 P_{\text{DLP}}^3(\cdot) &= \frac{\sqrt{610}}{610} [-16, 7, 0, -7, 16]^\top \\
 P_{\text{DLP}}^4(\cdot) &= \frac{\sqrt{1090}}{5450} [80, -85, 0, -85, 80]^\top .
 \end{aligned}
 \tag{2.8}$$

The DLPs can be used to truncate an expansion at any desired order with varying levels of accuracy, e.g., use Eq. (2.4) except let I take a value less than the number of values in the expanded function, \tilde{f}_i .

To illustrate, let $f(x)$ be defined as

$$f(x) = \cos(x) \quad 0 \leq x \leq 2\pi . \tag{2.9}$$

To expand $f(x)$ with the set of DLPs, F is first discretized by evaluating it at a number of points, e.g., 40 points. The expansion coefficients are then calculated by Eq. (2.3), and are then used to approximate f to various orders using Eq. (2.4). The results of such expansions are given in Fig. 2.2, where the approximate function is plotted alongside the discretized $f(x)$. Only the even orders are shown in the figure because $\cos(x)$ is an even function, and hence all expansion coefficients for the odd functions of the DLPs are equal to zero. It is evident from Fig. 2.2 that increasing the expansion order reduces the error of the approximation because more terms have been retained. However, adding odd functions to the expansion does not reduce the error of the approximation. In general, an expansion will converge as the basis set becomes more complete (where completeness is determined by the range of the basis set in the vector space of the function to expand), but it is not guaranteed that the expansion will converge to $f(x)$. In the case of the neutron transport equation, the dependence of the solution on the expansions is non-linear in general; thus, the error cannot be guaranteed to decrease with increasing expansion order, but the error is expected to decrease on the average as the order is increased.




FIGURE 2.2: Using 40 discrete points, $\cos(x)$ $x \in [0, 2\pi]$ expanded by Discrete Legendre Polynomials to 4th order

2.3 Modified Discrete Legendre Polynomials

To improve the DLPs as a basis for energy expansion, the polynomials are first modified by superimposing a “shape” vector s on each basis vector, leading to the intermediate vectors

$$\tilde{P}_{\text{mDLP}}^h(g) = P_{\text{DLP}}^h(g)s(g), \quad g = 1, 2, \dots, G. \quad (2.10)$$

The modified Discrete Legendre Polynomial (mDLP) vectors P_{mDLP}^h are subsequently found by orthonormalizing the vectors $\tilde{P}_{\text{mDLP}}^h$. This formulation is referred to as Type 1 mDLP (mDLP-1) [8]. It is obvious that to expand a known function using the mDLP-1 basis, the best shape vector would be the function itself. The power of modifying the DLP vectors is observed when using the same basis set to expand several different but related functions, e.g., the scalar flux as a function of energy for different spatial cells in a test problem.

As an example of mDLPs, consider a 10-pin test problem consisting of 5 pins of UO_2 next to 5 pins of mixed oxide (MOX), i.e., plutonium-bearing, fuel with reflective boundary conditions (this example problem is the first 1-D test problem discussed

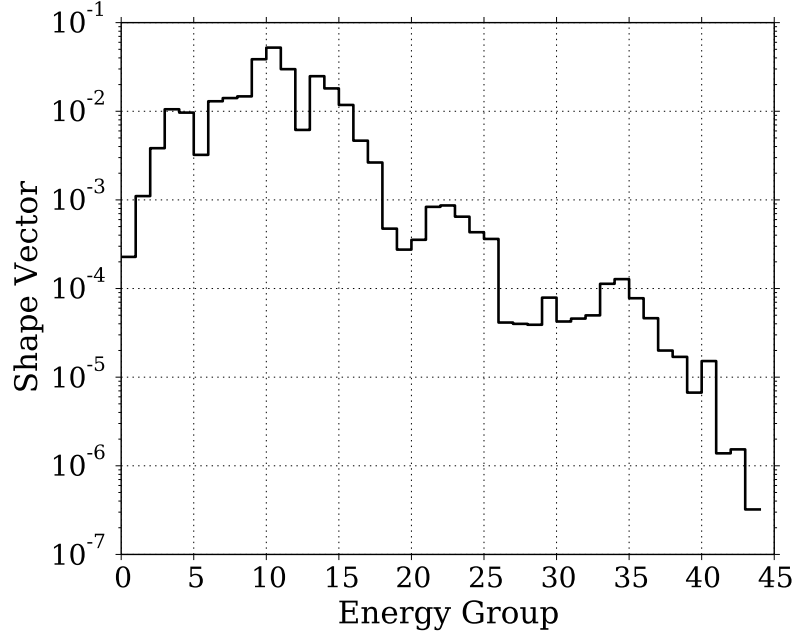


FIGURE 2.3: Example shape vector for creating mDLP basis functions

in Chapter 4 and is shown in Fig. 4.1). This problem can be discretized spatially into several regions, i.e., N spatial cells. Then the scalar flux ϕ can be calculated for each spatial cell, which leads to G group-dependent vectors that represent the scalar flux in each spatial cell. In this case, a shape vector representing the spatially-averaged scalar flux as a function of energy is used to modify the DLP vectors and form the mDLP basis. This example shape vector is shown in Fig. 2.3.

This shape vector will create the basis functions shown in Fig. 2.4 when using mDLP-1. A second type of mDLP (mDLP-2) is formed by multiplying only the first DLP vector (i.e., the flat vector) with the shape function, then orthonormalizing the set of vectors [8]. A sample of mDLP-2 basis vectors is shown in Fig. 2.5.

The efficiency of the mDLP vectors are shown by computing the error due to expansion as a function of order. The error is computed as

$$\epsilon = \sqrt{\frac{\sum_{h=1}^H (f(h) - \tilde{f}(h))^2}{\sum_{h=1}^H f(h)^2}}, \quad (2.11)$$

where $\tilde{f}(h)$ is the h th element of the approximate function \tilde{f} as defined in Eq. (2.4). Equation (2.11) is the discrete definition of the L_2 norm, which measures the deviation between two vectors in a least squares sense. Figure 2.6 shows the




FIGURE 2.4: Example basis functions for mDLP-1 shown through 3^{rd} order




FIGURE 2.5: Example basis functions for mDLP-2 shown through 3^{rd} order

performance of both types of mDLP compared to DLP when used to expand the set of scalar flux vectors described for the example 10-pin case. Note that at the full order (order = 43 for this example), all of the basis sets converge to within machine precision. It is clear that mDLP-1 outperforms mDLP-2, which suggests that more of the physics of the problem is captured by the lower-order functions of mDLP-1. Thus, for this problem, the functions of mDLP-1 can better approximate




FIGURE 2.6: Relative error in the L_2 norm as a function of expansion order for various basis sets

the solution to the problem than mDLP-2.

It must be reiterated that the solution for this problem was already known, and the shape vector was taken to be the average of the group-dependent vectors from each spatial cell. In practice, the scalar flux will not be known *a priori*, thus a representative shape vector must be chosen. For the 10-pin example in this chapter, a suitable shape vector may be constructed by taking the average of the spectrum produced by solving each type of fuel pin individually. Although this approach is expected to produce a viable shape vector, the constructed basis set will not lead to the same accuracy as achieved when using the average of the actual spectra as the shape vector.

2.4 Summary

Later in Chapters 5 and 6, DLP and mDLP are compared to the KLT (Karhunen-Loéve Transform) approach (which is discussed in Chapter 3). The best case of mDLP (i.e., mDLP-1) is compared to each of the various formulations of the KLT. However, the best-case mDLP requires the correct solution *a priori*. A more

practical use of mDLP is not expected to perform as well. The best case mDLP was chosen as the comparison to provide a view into the performance of the KLT as used for basis generation.

Chapter 3

Karhunen-Loève Transform

Although use of mDLPs is one way to build physics into an energy basis, a more powerful approach is to use the Karhunen-Loève transform (KLT) [11]. This method goes by many names such as Proper Orthogonal Decomposition (POD) [12] and Principal Components Analysis (PCA) [11]. KLT has also been used for a variety of purposes, including model reduction in both fluid dynamics [13] and reactor eigenvalue problems [12]. Additionally, KLT has been used in image compression [11].

The central goal of KLT is to approximate a discrete or continuous function $f(x)$ as a truncated expansion in an orthogonal basis whose functions yield the best possible n th-order approximation in terms of least-squares error for all values of n . In some applications, such as image compression, the function f is predetermined, e.g., a set of pixel values. For other applications, such as reduced-order modeling, the function f is not known. While details of its use in these applications may differ, KLT is fundamentally related to the singular value decomposition (SVD).

Suppose the global response matrix equations defined by Eq. (1.30) were solved using a full multigroup treatment and group-dependent boundary currents were reconstructed for each node from the resulting solution. Let N denote the number of these currents, each of length G , the number of energy groups. Furthermore, let the n th current vector be denoted by \mathbf{d}_n . These vectors, referred to as snapshots

later in this work, are combined to form the matrix $\mathbf{D} \in \mathbb{R}^{G \times N}$, which is defined as $\mathbf{D} = [\mathbf{d}_1, \mathbf{d}_2, \dots, \mathbf{d}_N]$.

The present goal is to find a set of vectors that is the best representation of the columns of \mathbf{D} , i.e., the set of snapshots. The full SVD of \mathbf{D} is defined as

$$\mathbf{D} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top, \quad (3.1)$$

where $\mathbf{U} \in \mathbb{R}^{G \times G}$ is an orthogonal matrix of the right-singular vectors, $\mathbf{V} \in \mathbb{R}^{N \times N}$ is an orthogonal matrix of left-singular vectors, and $\mathbf{\Sigma} \in \mathbb{R}^{G \times N}$ is a matrix with diagonal elements $\sigma_{kk} \geq 0$, $k = 1, 2, \dots, K = \min(G, N)$ arranged in decreasing order and zeros everywhere else. The σ_{kk} 's are referred to as singular values of the matrix \mathbf{D} . The first $R = \text{rank}(\mathbf{D}) \leq K$ singular values are positive (and the rank of a matrix is equal to the number of linearly independent columns).

If instead of allowing all of the nonzero singular values into the construction of $\mathbf{\Sigma}$, only the largest L values are used to define $\mathbf{\Sigma}$, the problem space will be reduced. Eckart and Young [14] proved that the rank- L matrix $\tilde{\mathbf{D}} \in \mathbb{R}^{G \times N}$ that satisfies

$$\min_{\substack{\tilde{\mathbf{D}} \in \mathbb{R}^{G \times N} \\ \text{rank}(\tilde{\mathbf{D}}) = L}} \|\mathbf{D} - \tilde{\mathbf{D}}\|_F \quad (3.2)$$

is defined as

$$\tilde{\mathbf{D}} = \tilde{\mathbf{U}}\tilde{\mathbf{\Sigma}}\tilde{\mathbf{V}}^\top, \quad (3.3)$$

where $\tilde{\mathbf{U}} \in \mathbb{R}^{G \times L}$ and $\tilde{\mathbf{V}} \in \mathbb{R}^{N \times L}$ contain the first $L \leq K$ columns of \mathbf{U} and \mathbf{V} , the diagonal matrix $\tilde{\mathbf{\Sigma}} \in \mathbb{R}^{L \times L}$ has nonzero elements equal to the first L singular values of \mathbf{D} , and the Frobenius norm of a matrix \mathbf{A} with elements a_{gn} and columns \mathbf{a}_n is defined as

$$\|\mathbf{A}\|_F = \sqrt{\sum_{g=1}^G \sum_{n=1}^N a_{gn}^2} = \sqrt{\sum_{n=1}^N \|\mathbf{a}_n\|_2^2}. \quad (3.4)$$

Therefore, the matrix $\tilde{\mathbf{D}}$ that satisfies Eq. (3.2) provides a least-squares representation of the entire set of data contained in \mathbf{D} . Equivalently, the columns of $\tilde{\mathbf{D}}$ are

approximations of the columns of \mathbf{D} with the minimum root mean square (RMS) error.

The KLT is designed to generate a set of orthogonal basis vectors that capture the most important information of the defining matrix. In this case, that matrix is \mathbf{D} , so the singular vectors of a matrix capture the most important information about the snapshots. Thus, part of the KLT is to identify the singular values of the matrix \mathbf{D} . By extension of Eq. (3.1)

$$\mathbf{D}^T\mathbf{D} = \mathbf{V}\mathbf{\Sigma}\mathbf{U}^T\mathbf{U}\mathbf{\Sigma}\mathbf{V}^T = \mathbf{V}\mathbf{\Sigma}^2\mathbf{V}^T. \quad (3.5)$$

Thus the SVD of the quantity $\mathbf{D}^T\mathbf{D}$ yields \mathbf{V} , which are the right-singular vectors of \mathbf{D} . When using the SVD of the matrix $\mathbf{D}^T\mathbf{D}$, the result is equivalent to the eigenvalue decomposition of $\mathbf{D}^T\mathbf{D}$. The eigenvalue decomposition is of a similar form to Eq. (3.5) and is given by

$$\mathbf{D}^T\mathbf{D} = \mathbf{B} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^{-1}, \quad (3.6)$$

where the columns of $\mathbf{Q} \in \mathbb{R}^{N \times N}$ are the eigenvectors of $\mathbf{B} \in \mathbb{R}^{N \times N}$, and $\mathbf{\Lambda} \in \mathbb{R}^{N \times N}$ contains the eigenvalues of \mathbf{B} in decreasing order.

From Eq. (3.6), the eigenvectors contained in \mathbf{Q} of the eigenvalue decomposition of $\mathbf{D}^T\mathbf{D}$ are equivalent to the right singular vectors of \mathbf{D} with corresponding singular values contained in $\mathbf{\Sigma}$. Also important to note is that $\mathbf{D}^T\mathbf{D}$ is a symmetric, square, real-valued matrix, which simplifies many linear algebraic computations and guarantees that the eigenvalue decomposition exists. Furthermore, many algorithms exist to find eigenvalues and corresponding eigenvectors of symmetric matrices, which are quicker than computing the SVD of \mathbf{B} .

In short, we desire the information contained in the singular vectors of \mathbf{D} . The singular vectors of \mathbf{D} are equivalent to the eigenvectors of \mathbf{B} . By taking only the N largest eigenvalues and corresponding eigenvectors of \mathbf{B} , the eigenvectors provide the best N th order least squares representation of \mathbf{B} , and correspondingly the least

squares representation of \mathbf{D} . By working with \mathbf{B} instead of \mathbf{D} , many computations are quicker.

This approach leads to the simple implementation algorithm to find the basis functions, $P_{KLT}^i(\cdot)$, first described by Meyer and Matthies [15], as follows. Form the data matrix, \mathbf{D} as above. Then the matrix $\mathbf{B} \in \mathbb{R}^{N \times N}$ is defined as

$$\mathbf{B} = \mathbf{D}^T \mathbf{D}. \quad (3.7)$$

Now find the eigenvectors corresponding to the largest eigenvalues of the matrix \mathbf{B} ,

$$\mathbf{B} = \mathbf{Q} \mathbf{\Lambda} \mathbf{Q}^{-1}, \quad (3.8)$$

where \mathbf{Q} from Eq. (3.8) is equal to \mathbf{V} from the SVD of \mathbf{D} in Eq. (3.1), and is thus the set of eigenvectors of \mathbf{D} . The eigenvectors \mathbf{q}_j (columns of \mathbf{Q} from Eq. (3.8)), are then multiplied by the data matrix to form the basis vectors $P_{KLT}^j(\cdot)$, i.e.,

$$P_{KLT}^j(\cdot) = \mathbf{D} \mathbf{q}_j, \quad (3.9)$$

which are subsequently orthonormalized. Then an approximate representation of an arbitrary N -vector \mathbf{f} in the basis is

$$\mathbf{f} \approx \sum_j a_j P_{KLT}^j(\cdot) \quad \text{where} \quad a_j = \mathbf{f}^T P_{KLT}^j(\cdot). \quad (3.10)$$

The form of Eq. (3.10) is comparable to that of Eqs. (1.19)–(1.22). The KLT is used to create a set of N basis vectors that are best N th order approximation of the data to which the KLT applied. In other words, the k th basis function is stretched such that the basis vector can best approximate a snapshot \mathbf{d} when used in conjunction with all $i < k$ vectors of the basis set. Thus, the KLT maximizes the projection of a basis vector onto the data to which the KLT was applied. This means that the KLT basis is constructed to obtain the most efficient basis functions for order reduction. For example, the most efficient storage for a zeroth expansion of a function is the function itself, which is an uninteresting use of KLT. However,

when instead applied to a set functions, KLT efficiently detects the prominent modes in the set. In this case, the zeroth-order basis function typically resembles the average of the functions to which KLT is applied.

Defining the KLT basis vectors requires that values of the function to be expanded be known *a priori*. For an application such as image compression, those values are completely known. However, for applications in reduced order modeling, knowledge of the true function to be expanded, e.g., the group-dependent angular current, would require the full system to be solved, which may be impossible for sufficiently large models. Additionally, if the complete solution is already known, there remains little need to approximate the solution. Therefore, an approximation must be made in order to sample values of the function to be expanded.

For time-dependent fluid dynamics, Sirovich proposed a method [13] that is now known as the *method of snapshots* [12]. The method was used originally to produce reduced-order models for spatially-discretized, continuous-in-time models by computing the space-dependent solution only at discrete times. The resulting data, or matrix \mathbf{D} , is reduced significantly from the continuous case.

Again, the solution of the global response matrix equations should be avoided because it could be prohibitively large and would preclude the need for a basis expansion. Thus, an effective basis set that can reconstruct the global solution without solving the global solution directly is sought. The true set of vectors \mathbf{d}_n is assumed to be unavailable, so an alternative method to generate snapshot data is required. The implemented approach is similar to traditional lattice physics methods in which a small, representative portion of the global domain (e.g., a many-group, two-dimensional assembly model) is used to generate an approximate global model (e.g., few-group, three-dimensional reactor model). For KLT, small, “snapshot model” problems are solved from which group-dependent fluxes at various space and angle points are extracted as snapshot vectors. These vectors become columns of an approximate \mathbf{D} , and KLT is constructed and applied as described. Since snapshot vectors do not represent the true solution space, optimality of KLT is not guaranteed. However, if the set of snapshots is sufficiently similar to the




FIGURE 3.1: Basis functions for the KLT applied to an example problem

solution, then KLT can provide excellent results when applied to ERMM as an energy basis.

To illustrate application of KLT, reconsider the 10-pin example from Chapter 2. The KLT approximation is applied to expand the known solution for the scalar flux, and its performance is compared to that of DLP and mDLP. The data is arranged as the matrix $\mathbf{D} \in \mathbb{R}^{G \times N}$, where $G = 44$ and $N = 280$. Following the algorithm described starting with Eq. (3.7), the basis function are constructed. These basis functions can be observed in Fig. 3.1. Note that the zeroth-order function is approximately identical to the zeroth order from either type of mDLP. It is in the remaining basis functions that KLT differs from mDLP.

The performance of the KLT basis set is shown in Fig. 3.2. Clearly, KLT outperforms DLP and mDLP for this example problem as KLT retains considerably more physics in the basis functions than mDLP can achieve. Also note that it takes considerably longer to compute the KLT basis function as compared to DLP or mDLP; however, the computation time for the basis generation of KLT is still orders of magnitude shorter than the time required to solve for the snapshot models or the test-problem itself.

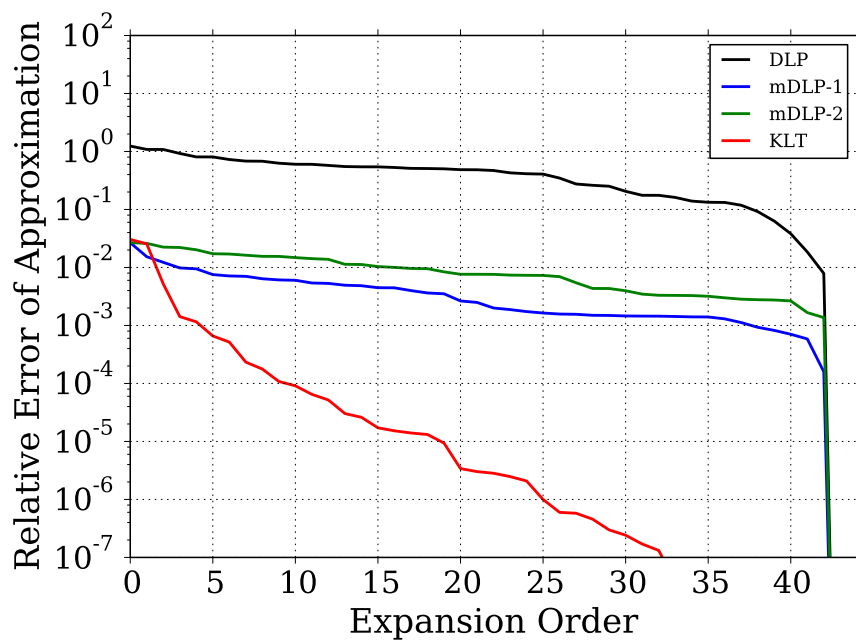


FIGURE 3.2: Relative error in the L_2 norm as a function of expansion order for KLT compared to other basis sets

Chapter 4

Test Problems and Models

For the purpose of testing the KLT as a basis generator for the ERMM, a number of test problems and corresponding snapshot models were developed. A total of three test problems were constructed, two that are 1-D and one that is 2-D. These test problems serve as the way to benchmark the implementation of KLT into the ERMM as a way to approximate the energy variable.

4.1 Test Problems

The test problems considered are (1) a 1-D, 10-pin, UO₂-MOX assembly model, (2) a 1-D, 70-pin BWR core model, and (3) the 2-D C5G7 benchmark [16]. The transport code DETRAN was used with a 16-angle, double Gauss-Legendre quadrature and a step-characteristic spatial discretization for problems (1) and (2). Problem (3) was solved with a Gauss-Chebyshev quadrature for the polar variable and Abu-Shumays' Quadruple Range quadrature for the azimuthal variable [8]. Cross-section libraries for the test problems were generated with SCALE 6.1 in the SCALE 44-group and 238-group formats [17].

For the 1-D problems, a current-conserving, first-order, angular expansion based on Jacobi polynomials [8] was used. For the 2-D problem, a set of Chebyshev polynomials was used for the angular expansion. For the 1-D problems, nodal

boundary currents have no spatial dependence thereby rendering spatial expansions unnecessary, but for the 2-D problem, a spacial expansion based on DLP was used. For energy, DLP, mDLP, and KLT bases were explored as a function of the energy expansion order.

All calculations were performed with the SERMENT parallel response matrix code [3], which links to the DETRAN deterministic transport code [18]. To implement each of the test problems, each was first solved in DETRAN to provide the snapshots (which are discussed later in this chapter). Next, each model was solved again using SERMENT for the full multi-group reference case. The same spatial and angular expansion order was used throughout the calculations for each test problem and corresponding snapshot models. Finally, SERMENT was used to calculate the test problem for several expansion orders with the chosen basis set. These calculations were compared to the reference case, and the results are presented in Chapter 5 for the 1-D problems and Chapter 6 for the 2-D problem.

4.1.1 1-D Test Problems

Two 1-D test problems were developed. The first, the 10-pin problem, serves as the proof of concept. It was a relatively simple problem to show how KLT can be used to approximate the junction between two fuel types with very different spectral properties. The second problem is designed to be more complex and, thus, more difficult to model. It is a 1-D, full-core model with several fuel assembly types, and, due to its complexity, it should provide a better test of KLT.

10-Pin Problem

The first test problem was a 1-D approximation to the junction between a UO_2 and MOX assembly. In the 10-pin test problem, fuel for the left five pins was 4% enriched UO_2 , while the right five pins were composed of 4.3% enriched MOX (The fuel contained 4.3% plutonium with the remainder natural uranium), as shown in Fig. 4.1. Boundary conditions on either side of the model were reflective.




FIGURE 4.1: Configuration for 10-pin Test Problem

Fuel pins were 1.08 cm thick with 0.09 cm of moderator on each side. The baseline pincell discretization consisted of 22 mesh cells of fuel enclosed by three mesh cells of moderator on either side; therefore, each pincell provided 28 energy-dependent snapshots.

Because the test problem reference solution is a full transport approximation, boundary currents generally exhibit coupled angle-energy dependence. By including snapshots that incorporate angular information, the resulting KLT basis set may outperform snapshots based only on the energy-dependent scalar flux ϕ . To include angular information, snapshots were taken of the energy-dependent partial current J_{left} . Snapshots of the net current were previously considered, but those snapshots performed as well as or worse than the partial current in all cases. Further, because the snapshot models were symmetric, the direction of the partial current is selected arbitrarily. Finally, because the spatial discretization produces only one spatial unknown per cell, the snapshot generation approach used provided one (ϕ or J_{left}) or two (ϕ and J_{left}) snapshots per spatial cell.

In addition to inclusion of snapshots of the partial current for each test problem, the addition of snapshots of higher-order angular moments was studied. These moments were generated by expanding the angular flux ψ through a Jacobi expansion in angle. In this basis, the zeroth moment is equivalent to the partial current, while higher moments have less well-defined physical corollaries. Inclusion of these moments (up through second order) was studied for both test problems and are presented at the end of the results section.

Several combinations for the snapshot types were considered and presented in the following chapters. First, snapshots of just ϕ were used. Second, snapshots of just J_{left} were used. Third, KLT was used with snapshots of both ϕ and J_{left} together.

Fourth, snapshots of each of the higher order moments (up through order two) were used individually for the test problems (e.g., snapshots from just angular moment two). Finally, The snapshots for ϕ , J_{left} , and the higher moments (incrementing upward in order) were used for basis generation (e.g., ϕ , J_{left} , order 1 or ϕ , J_{left} , order 1, order 2 etc.).

BWR Test Problem

The second test case, representative of a boiling water reactor (BWR), was comprised of seven assemblies, each with 10 pins. Three core configurations were used, and each core configuration had two unique assemblies. Three fuel types were used, including 4.5% enriched UO_2 , 2.5% enriched UO_2 , and 4.5% enriched UO_2 with 5 wt% Gd_2O_3 . Core and assembly configurations are shown in Fig. 4.2. Boundary conditions for this case were vacuum. This test problem was adapted from the work of Nichita and Rahnema [19] and Ilas and Rahnema [20].

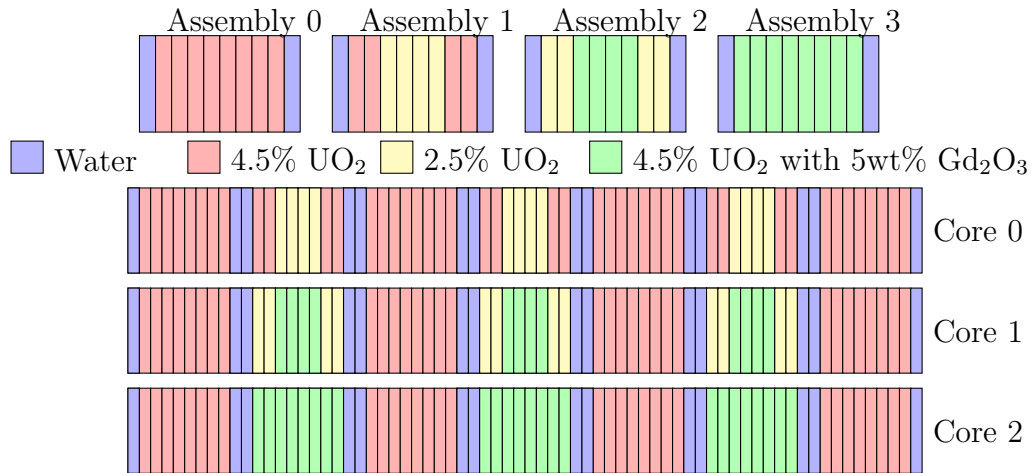


FIGURE 4.2: Configuration for BWR Test Problem

Fuel pins for the second problem were 0.72 cm thick with 0.27 cm of moderator on each side. The baseline pincell discretization consisted of 16 mesh cells of fuel enclosed by six mesh cells of moderator; therefore, each pincell provided 28 energy-dependent snapshots. With 70 pincells, the total number of snapshots was 1960, but the right set of snapshots was identical to the left with respect to scalar

flux due to symmetry. For these simple 1-D problems, finer spatial and angular refinement was unnecessary.

For the BWR test problem, the combinations of potential snapshots were the same as used for the 10-pin test problem; however the snapshot models used to generate snapshots were different, as will be discussed later in this chapter.

4.1.2 2-D Test Problems

The C5G7 benchmark is a well-studied problem for assessing numerical methods in neutron transport and reactor physics. This benchmark consists of modeling a quarter core that has four, 17×17 -pin assemblies [16]. The configuration is detailed in Fig. 4.3. Each of the blocks in Fig. 4.3 represent either a fuel assembly or an equivalent area of moderator. The configuration of the UO_2 assembly is shown in Fig. 4.4 and the configuration of the MOX bundle is shown in Fig. 4.5. Each individual pincell is modeled as shown in Fig. 4.6. Each pincell is broken up into a 7×7 Cartesian mesh. Thus each pincell contained 49 spatial cells and provided 49 energy dependent snapshots of each type. The cladding for the pincells was homogenized into the fuel part of the pincell.

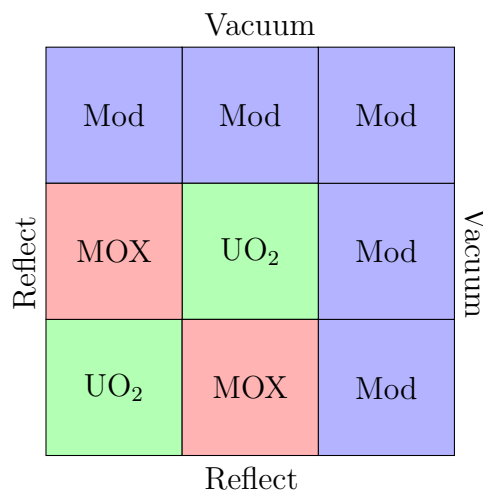


FIGURE 4.3: Configuration for Full-Core. Each square represents the area of a 17×17 pin assembly

Although the original C5G7 benchmark was defined with 7-group data, energy order reduction is unnecessary with so few groups. Consequently, cross-section libraries




FIGURE 4.4: Configuration for UO_2 fuel bundle. The green represents a UO_2 pincell, while the blue represents a guide tube modeled as a pincell filled with moderator

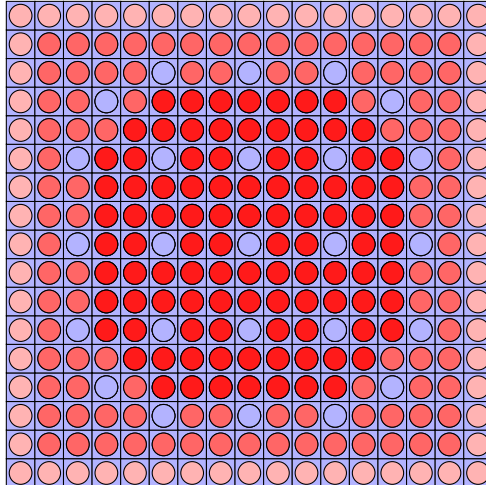


FIGURE 4.5: Configuration for MOX bundle. The light red represents 4.3% MOX fuel, the medium red represents 7.0 % MOX fuel, and the dark red represents 8.7% MOX fuel. The blue represents moderator (i.e., light water)

were computed using Scale 6.1 for the C5G7 problem to adapt the benchmark for use with 44 energy groups like the other test problems, but, all other aspects of the problem remained unchanged.

For this test problem, the spatial and angular orders for the SERMENT solution were set to second order. This was in the interest of time, as utilizing higher orders will take much longer to solve (but be more accurate). Previous work by Roberts [8] suggested that using spatial and angular orders of at least fourth order are required to achieve sub-0.1% errors in pin powers. This error will be present when




FIGURE 4.6: Configuration for pincell. The circular fuel element had a radius of 0.54 cm and was homogenized with cladding for this model.

comparing the multigroup SERMENT solution to the benchmark C5G7 results. However, it is assumed that since the same spatial and angular orders were used for computing the reference case and the snapshot cases, the relative error of the snapshot cases is a function of the energy expansion alone. It is impossible to be sure if this assumption is valid without computing the problem at better resolution.

4.2 Snapshot Generating Models

Ideally, with any form of model reduction, the computational effort required to solve a given problem is reduced, so a quickly-computed basis set is sought. Therefore, the generation of snapshots for KLT was a primary focus for this work. For each test problem, a number of small “snapshot models” were developed. These models are quick to solve and provide group-dependent snapshots that represent the spectrum of the test problem of interest. In this section, the models chosen for each test problem are discussed.

4.2.1 Snapshots for 10-Pin Test Problem

Since the first test problem was a 1-D approximation of the junction between a UO_2 and MOX assembly, a number of small, representative subproblems are clear choices for snapshot generation. Models studied for this problem are summarized

in Table 4.1. The simplest approach by which to generate snapshots is to model each pincell individually subject to reflecting boundary conditions on all surfaces and to extract snapshots (i.e., energy-dependent vectors in distinct spatial cells) from an individual pincell. Additionally, snapshots from several pincell models may be combined together in hopes of improving the range accessible by the KLT basis functions.

TABLE 4.1: Summary snapshot models for 10-pin Test Problem

Abbreviation	Model to generate snapshots
Full-Assembly	Repeating array of 10 UO ₂ and 10 MOX pins
N-pin	Repeating array of N UO ₂ and N MOX pins
Combined-Pins	Combined snapshots from UO ₂ and MOX models, and two-pin, UO ₂ -MOX model
UO ₂ -Pin	UO ₂ pin only
MOX-Pin	MOX pin only

A larger problem space is obtained if a single model includes more than one pin (fuel or moderator) in various arrangements. For this work, an effectively 10×10 model (infinite lattice of ten UO₂ and ten MOX pins; denoted as the full-assembly) was studied. The full-assembly model was equivalent to the test problem of interest and should be expected to yield snapshots that capture the true multi-group solution with the lowest-order energy expansion. The remaining models were effectively $N \times N$, where N was taken to be 1, 2, or 3.

4.2.2 Snapshots for BWR Test Problem

The second test problem was a 1-D approximation of a BWR core. Models for this case are summarized in Table 4.2. Three cases naturally arise from core construction. The first was to take snapshots from the entire core. This model was expected to perform the best, as the model was equivalent to the test problem. The second was to model assemblies with reflective boundary conditions for the given configuration and combine snapshots from the two unique assemblies used for each core configuration. The final approach was to model each pincell in the

test problem with reflective boundary conditions and then combine the snapshots from each unique fuel pin used in the core configuration.

TABLE 4.2: Summary of snapshot models for BWR Test Problem

Abbreviation	Model to generate snapshots
Full-Core	Snapshots from whole core model (i.e., the test problem)
Combined-Assemblies	Snapshots from unique assemblies used in core configuration
Combined-Pins	Snapshots from unique pins used in core configuration

4.2.3 Snapshots for C5G7 Test Problem

There are several models that are readily apparent choices to simplify the C5G7 benchmark for use in generating snapshots. All of the snapshots models for this test problem are summarized in Table 4.3. The first model was to use snapshots from the test problem itself to gain an understanding of the best possible snapshots. However, the number of snapshots obtained for this model is prohibitively large, even after removing all duplicate snapshots from model symmetry, and thus is unusable in the raw state. However, by spatially averaging all the snapshots of a given type from a pincell, the number of snapshots is reduced by a factor of 49, and thus becomes a manageable set with which to create basis functions. The Combined-Assemblies model was formed by taking snapshots from each assembly and combining the snapshots together. For this snapshot model, each assembly was modeled with reflective conditions on all sides. The UO₂ assembly used the same configuration as in Fig. 4.4, and the MOX assembly used the same configuration as in Fig. 4.5.

The next model was the Combined-Pins model, which derives snapshots from individual pincells for each fuel type. These pincells have the same configuration as shown in Fig. 4.6. The pincell snapshots are combined with snapshots from junctions of pincells. These junctions are modeled as 2-pin by 2-pin assemblies with reflective conditions as shown in Fig. 4.7. There were a total of 4 different fuel types, thus six unique junctions to model.

TABLE 4.3: Summary of snapshot models for C5G7 test Problem

Abbreviation	Model to generate snapshots
Reduced Full-Core	Spatially-averaged snapshots from whole core model (i.e., the test problem)
Combined-Assemblies	Snapshots from assemblies used in core configuration
Combined-Pins	Snapshots from pins used in core configuration combined with the pin junctions
Small-Core	Snapshots from the small core model
Small-Assemblies	Snapshots from the small assemblies used in the small core configuration
Reduced Small-Core	Spatially-averaged snapshots from the small core model
1-D Approximation	Snapshots from the 1-D approximation to the C5G7 benchmark




FIGURE 4.7: Configuration for pincell junction; all Reflect BC; Every combination is used

The next model was to create a small version of the C5G7 core, which has the same assembly configuration shown in Fig. 4.3, but instead uses 8-pin by 8-pin assemblies. The small UO_2 assembly is shown in Fig. 4.8, and the small MOX assembly is shown in Fig. 4.9.




FIGURE 4.8: Configuration for small UO_2 fuel bundle. The green represents a UO_2 pincell, while the blue represents a guide tube modeled as a pincell filled with moderator

An additional model arises from the small core model, which is to combine snapshots from each of the small assemblies. This is not expected to perform as well as the




FIGURE 4.9: Configuration for small MOX bundle. The light red represents 4.3% MOX fuel, the medium red represents 7.0 % MOX fuel, and the dark red represents 8.7% MOX fuel. The blue represents moderator, light water

full-sized, combined-assembly model, but it will be quicker to create the snapshots, and thus the basis functions. A third model from the small core was to spatially average the snapshots for each pincell similarly to the spatially-reduced, full-core solution. The reduced, small-core model was used to approximate the non-spatially averaged, full-core model by comparing the reduced and non-reduced small-core models.

The final snapshot model for the C5G7 test problem is to create a similar 1-D model. This is shown in Fig. 4.10. The model is comprised of 51 pins that are each of width 1.26 cm. The fuel section for each pin was modeled at 1.08 cm, leaving 0.09 cm of moderator on either side. This problem was created with a reflective boundary condition on the UO_2 side of the model and a vacuum boundary condition on the moderator side of the model.




FIGURE 4.10: Configuration for 1D approximation to the C5G7 benchmark. The green represents UO_2 , the light red represents 4.3% MOX fuel, the medium red represents 7.0 % MOX fuel, and the dark red represents 8.7% MOX fuel, and the blue represents moderator

4.3 SERMENT and DETRAN

All of the solvers and methods, save for new basis generation, were developed previously and implemented in two transport codes, DETRAN and SERMENT. The SERMENT parallel response matrix code [3] links to the DETRAN deterministic transport code [18] to implement ERMM and allows the choice of several orthogonal basis sets. DETRAN implements discrete ordinates, method of characteristics, and diffusion approximations, along with several advanced solvers developed specifically for use in response function generation. For this work, only its discrete ordinates capabilities were used.

DETRAN was used to generate the data to be used for each of the snapshot models. It requires no basis expansions, and is taken as the true solution for each of the test problems. SERMENT was used to generate the reference case for each test problem against which all of the KLT solutions for the given test problem were compared. The reference case was taken as a full order multi-group approximation with the same spatial and angular expansions as the KLT cases. This was done to isolate the error caused by the expansion in energy to the greatest extent possible.

The new functionality for SERMENT that was developed in this work is the ability to incorporate user-defined basis sets. This allows an arbitrary order and size of basis functions to be passed to SERMENT enabling expansion via that basis.

Chapter 5

Results for 1-D Studies

The goal of this work was to reduce the number of energy degrees of freedom in a given problem to a manageable size without compromising the accuracy of the model. However, the best expansions of a function can only be performed when the function is completely known, which, in this case, means the full solution is required before the expansion, which precludes the need for an expansion. As such, this work focused on two classes of basis expansions. The first class is the expansions based on the test problem of interest, which provide insight to the best that a basis set can do. This class is compared to the sets of basis functions that are based on small, yet similar models to the test problem. These small models are quick to solve, and hopefully similar enough to the test problem that an expansion based on the model would be accurate enough for the larger test problem. Thus, another focus of this work was to identify how similar a model had to be to a given test problem to provide effective basis functions for expansion.

Typically, reactor analysis attempts to compute pin fission densities (or powers) with sub-1% errors. This work added an additional buffer, and focused on the minimization of the energy expansion order required to achieve sub-0.1% maximum relative errors in the fission density. The reference solution in all cases was a full multi-group response matrix solution for the given test problem with consistent angular expansion used throughout. This consistency ensured the observed changes in the solution were functions of only the energy basis used. In practice, it is

not possible to truly separate the effects of space, angle, and energy because the responses in each phase-space variable are coupled, but it is assumed that this treatment will provide adequate insight into the effect of the chosen energy expansion.

With the exception of ‘DLP’ and ‘mDLP’, each curve was generated using a KLT basis with distinct snapshot data, as described previously in Chapter 4. The mDLP results represent the best case previously observed by Roberts [8], which was to use mDLP-1 with the average flux profile from the complete test problem of interest as the shape function. Therefore, the mDLP results are not practical because the solution for the test problem must be known *a priori*.

5.1 mDLP Comparison

Figure 5.1 shows a comparison between practical applications of both versions of mDLP when applied to the 10-pin problem using the 44-group cross-section library. Figure 5.2 shows the same comparison, but instead for the 238-group cross-section library. In Fig. 5.2, both plots show the same data, but the left plot is truncated to order 43 to ease comparison to the 44-group data. For these figures, the denotation “full” means the shape vector was the flux spectrum averaged over space across all 10 pins. The denotation “UO₂” used the flux spectrum from a UO₂ pin as the shape vector. Likewise, the denotation “MOX” used the shape vector as the flux spectrum from a MOX pin.

All cases of mDLP-2 performed worse than to mDLP-1. The “full” denotation represents the impractical expansion in that the complete problem is required prior to the expansion. The “UO₂” and “MOX” denotations represent more practical expansions as they each require only solving a pincell problem to generate the expansion; however, these two models do not have the physics for one of the fuel types, and thus do not perform as well. As expected, the “full” expansions of mDLP-1 perform the best in general and were chosen as the comparison point for the rest of this work. Lastly, the increased error at high-order of Fig. 5.2 is due to

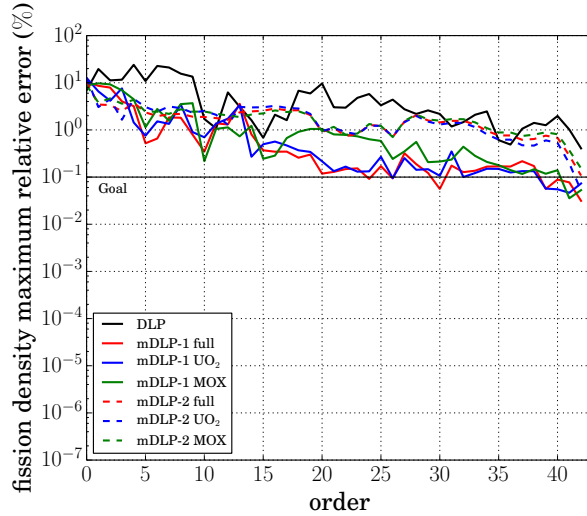


FIGURE 5.1: Comparison of mDLP applied to the 10-pin problem using 44-group cross-section library.




FIGURE 5.2: Comparison of mDLP applied to the 10-pin problem using the 238-group cross-section library.

non-orthogonality of the high-order basis functions due to accumulating roundoff error during orthogonalization.

5.2 Energy Spectra for the Test Problems

The 10-pin test problem was first solved to generate flux profiles to be used for mDLP and KLT expansions. In Fig. 5.3, the spatially averaged flux spectra are presented for the 10-pin problem with the 44-group cross-section library. These averaged spectra were used as the shape vectors for mDLP, while the non spatially-averaged data were used for KLT expansion. These spectra can be compared to




FIGURE 5.3: Flux spectrum for the 10-pin problem using 44-group cross-section library

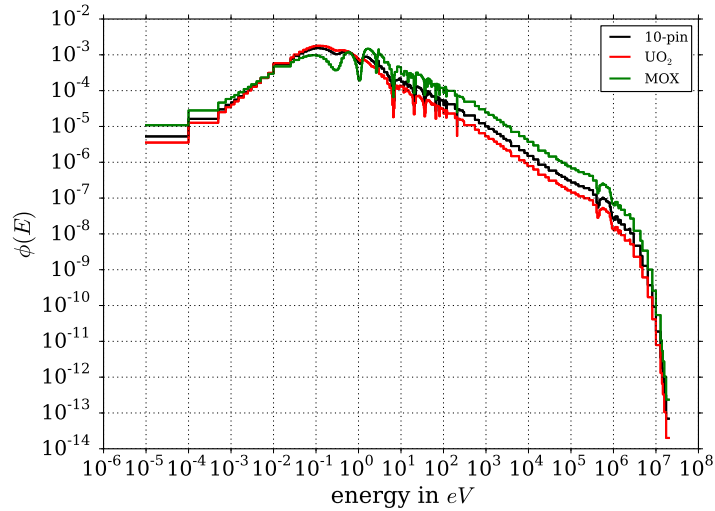


FIGURE 5.4: Flux spectrum for the 10-pin problem using 238-group cross-section library

those in Fig. 5.4, which shows the resulting spectra for the 10-pin problem using the 238-group cross-section library. As can be observed in each of these figures, the basic shape of the spectrum are nearly the same, with the 10-pin solution lying as average between the two individual pin spectra, UO_2 and MOX.

Figure 5.5 shows the spatially averaged spectrum for Core 0 of the BWR test problem, where the spectra denoted “assay1” or “assay2” are the spatially averaged energy spectra for each assembly in Core-0 of the BWR test problem. It is apparent that these spectra do not differ as much as the spectra associated with the 10-pin

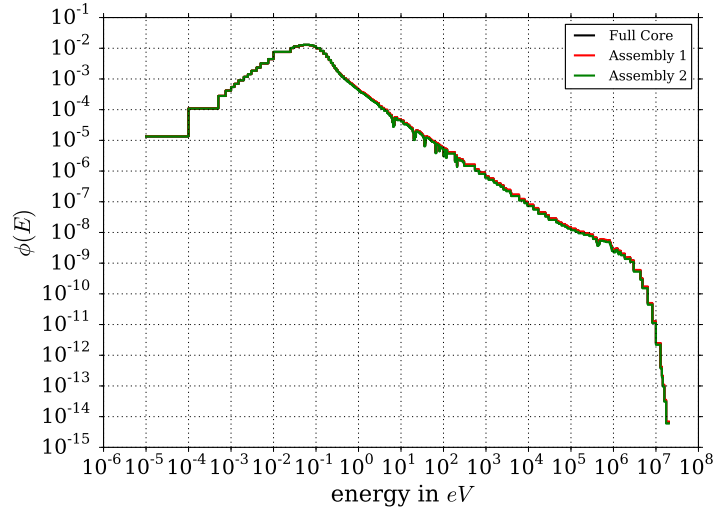


FIGURE 5.5: Flux spectrum for the BWR-Core 0 problem using 238-group cross-section library

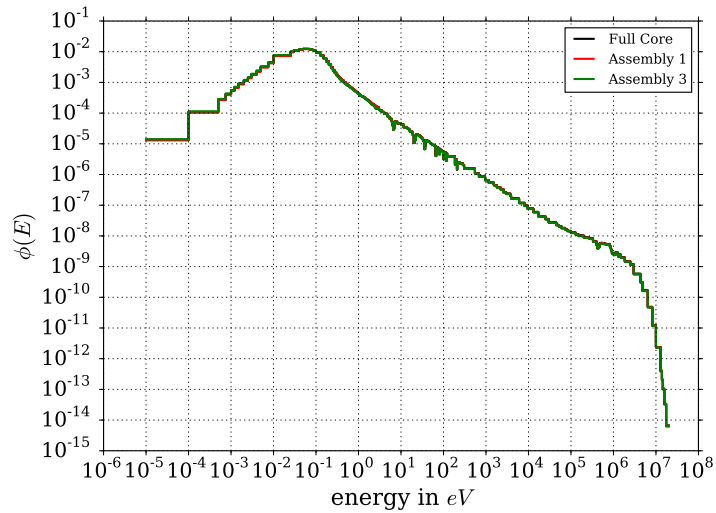


FIGURE 5.6: Flux spectrum for the BWR-Core 1 problem using 238-group cross-section library

problem. As will be shown later, this leads to improved performance of mDLP because there is not as much variation between the various parts in the model. The spectra for Core-1, shown in Fig. 5.6, is similar to that of Core-0. However, the spectra difference becomes more pronounced when observing the spectra for Core-2, shown in Fig. 5.7.




FIGURE 5.7: Flux spectrum for the BWR-Core 2 problem using 238-group cross-section library

5.3 10-Pin Test Problem

As previously discussed, the KLT was applied to the 10-pin test problem using several snapshot models. Additionally, there were several kinds of snapshots taken from each snapshot model i.e., scalar flux (ϕ), leftward partial current (J_{left}), and higher-order angular moments. These snapshots were combined in various ways for the 10-pin problem and are presented in this section. Each of the plots in this section are shown up to order $G - 1$, where G is the number of groups in the cross-section library. The last order is omitted because orthogonal basis functions are exact to machine precision when using a complete expansion. Some of the DLP and mDLP results will not appear to converge nearing the final order, which is due to non-orthogonality of the high orders from accumulating roundoff error in the orthogonalization. Furthermore, since the problem is non-linear, the error is not guaranteed to decrease with increasing expansion order, however, the error should decrease on the average with increasing expansion order.

5.3.1 44-Group Results

Figure 5.8 shows the performance of various formulations of the KLT as applied to the 10-pin problem. The snapshots used for the figure came from only the

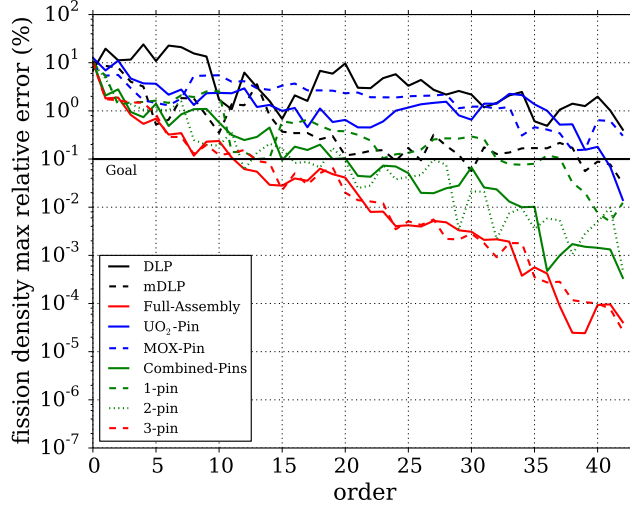


FIGURE 5.8: Performance of the KLT when applied to the 10-pin test problem with snapshots of only ϕ .

44-group scalar flux. When using snapshots from the full assembly model (10-pin), the relative error fell below the 0.1% threshold at an energy order of approximately 12, while more practical models, e.g., Combined-Pin, require approximately order 20 to reach the goal. The results of the Combined-Pins model are far from the goal of using approximately five energy degrees of freedom to reach the error goal. Essentially all of the KLT formulations outperformed mDLP-1 with the exception of the individual pins (UO_2 and MOX), which is expected as those models lack physics for half the problem space.

Figure 5.9 shows the performance of the KLT using snapshots of only J_{left} from the 44-group snapshot models. As shown, the best case KLT can reach the goal by approximately order 6, while more practical models, e.g., Combined-Pin, require approximately order 20 to reach the error goal.

Figure 5.10 shows the results of using the combined snapshots of ϕ and J_{left} each from the 44-group models. With this approach, the best case still reaches the goal at about order 6, but the practical case of Combined-Pins was improved to achieving the goal by order 15. In general, including both sets of snapshots improves the performance because KLT can extract the most important information from both sets and create highly effective basis functions.

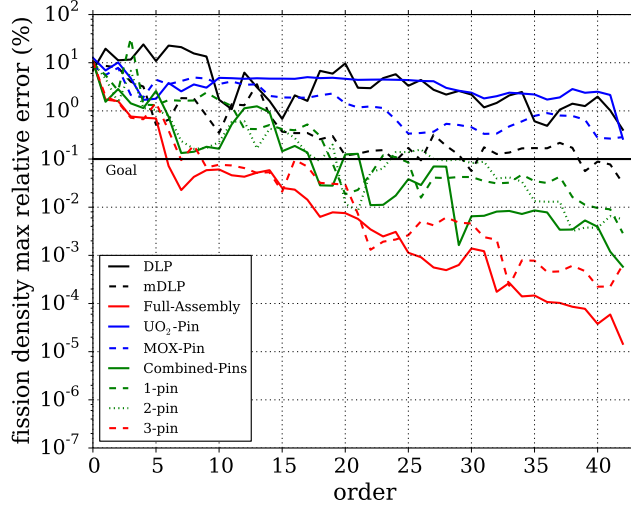


FIGURE 5.9: Performance of the KLT when applied to the 10-pin test problem with snapshots of only J_{left} .




FIGURE 5.10: Performance of the KLT when applied to the 10-pin test problem with snapshots of both ϕ and leftward partial current.

5.3.2 238-Group Results

Similar to the 44-group section, the figures in this section will be presented to showcase how the choice of snapshots impacts the performance of the KLT applied to the 10-pin problem using the 238-group cross-section library. All of the figures in this section include side-by-side comparisons of the same data, one shown to 43th order, while the other presents the full spectrum up to 237th order.

Figure 5.11 presents the results from using snapshots of only ϕ in the basis generation. The best expansion requires approximately order 14, while the Combined-Pins

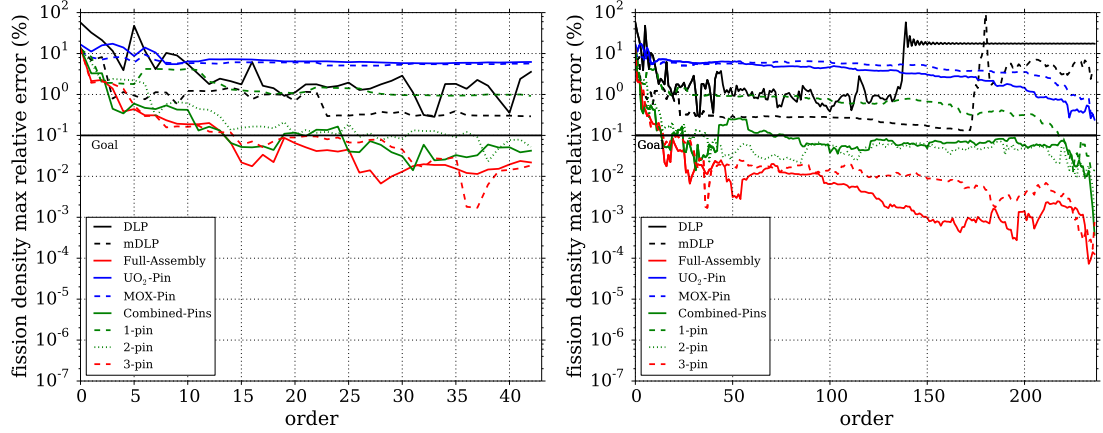


FIGURE 5.11: Relative error for 238-group, 10-pin test problem using snapshots of only ϕ

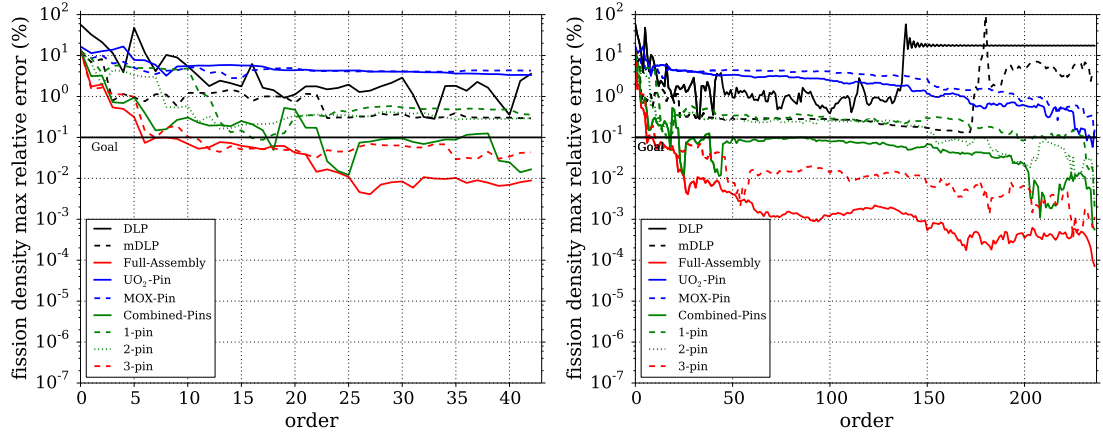


FIGURE 5.12: Relative error for 238-group, 10-pin test problem using snapshots of only J_{left}

case requires approximately order 24 to remain under the error goal. The individual pin models do not perform well, as is expected, as each lacks the information from an entire type of fuel pin.

Presented in Fig. 5.12 are the results from using snapshots of only J_{left} for basis generation. The impractical case of the 10-pin model requires approximately order 6 to reach the error goal, while the practical Combined-Pins model requires approximately order 22 to reach the error goal. These results are an improvement for the 10-pin model over using ϕ snapshots. However, the Combined-Pins model results are relatively unchanged as compared to using ϕ snapshots.

Figure 5.13 shows the results from combining together the snapshots of both ϕ and J_{left} . In this case, the 10-pin model requires approximately order 10 to reach




FIGURE 5.13: Relative error for 238-group, 10-pin test problem using snapshots of both ϕ and J_{left}

the goal, while the Combined-Pins model requires approximately order 14 to reach the goal. As compared to the 44-group data, these data from 238-group do not improve to the same degree when ϕ and J_{left} snapshots are combined.

Note that when comparing the cases between the 44-group and 238-group results, basis functions from a given type of model (e.g., Combined Pin) perform similarly despite the number of groups. It appears that for simple problems, KLT basis sets may contain a relatively constant amount of information despite increasing group size, but more effort is required to confirm this trend.

5.4 BWR Test Problem

The BWR test problem refers to the grouping of three different core configurations as discussed in Chapter 4. Each of the configurations was used with both 44-group and 238-group cross-section libraries. The configurations were designed with increasing inhomogeneity leading to more difficult models. The difficulty results in a larger error for the same energy order while progressing through configurations. Thus, the performance of each expansion is expected to worsen with increasing configuration number.




FIGURE 5.14: Relative error for the 44-group, BWR-Core 0 test problem using snapshots of only ϕ

For each configuration, the Full-Core model is expected to perform the best, as it uses all available unique snapshots for basis generation, while the Combined-Assemblies and the Combined-Pins models represent practical cases for basis generation. Note that for this test problem, mDLP and DLP perform quite well, which is due to the relatively small difference between the spatially averaged flux profile as discussed previously. Again, the mDLP results presented here used the spatially averaged flux profile from the Full-Core model as the shape vector, and, thus, do not represent a practical performance of mDLP, but rather should be compared to the Full-Core KLT results.

5.4.1 Configuration 0

5.4.1.1 44-Group Results

As Fig. 5.14 shows, it takes approximately order 6 for the Full-Core model to reach the goal of sub-0.1% relative error in the fission density. Whereas the practical Combined models require approximately order 15 to reach the goal. These results are when using snapshots of only ϕ .

Many of the results are improved if snapshots of J_{left} are used for basis generation with the KLT, as shown in Fig. 5.15. The Full-Core model requires approximately




FIGURE 5.15: Relative error for the 44-group, BWR-Core 0 test problem using snapshots of only J_{left}




FIGURE 5.16: Relative error for the 44-group, BWR-Core 0 test problem using snapshots of both ϕ and J_{left}

order 4, while the Combined-Assemblies model requires approximately order 10 to reach the goal.

The best results are obtained when snapshots of both ϕ and J_{left} are combined together to generate the KLT basis. These results are presented in Fig. 5.16. The Full-Core model requires order 4 to reach the error goal. The Combined-Assemblies reached the goal at approximately order 10. Finally, the Combined-Pins results are greatly improved, and require approximately order 13 to reach the goal when using both types of snapshots.

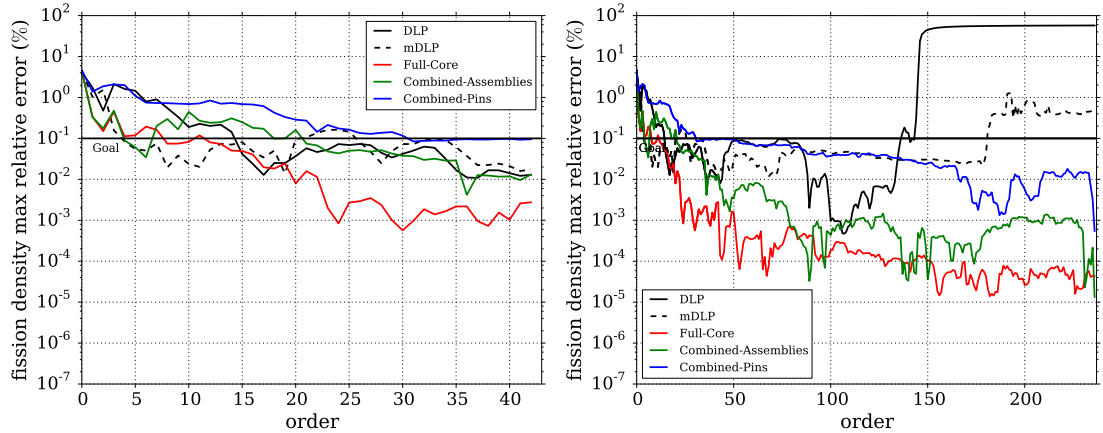


FIGURE 5.17: Relative error for 238-group, BWR-Core 0 test problem using snapshots of only ϕ

5.4.1.2 238-Group Results

As before, the results in this section are presented as side-by-side comparisons of the 238-group data. The left side is shown to only order 43 to ease comparison to the 44-group data, while the right side shows up to 237th order for the 238-group data.

Figure 5.17 shows the results from using snapshots of only the ϕ for the 238-group cross-section library. The impractical model of Full-Core requires approximately order 12 to reach the goal, while the Combined-Assemblies case requires approximately order 21. Each of these models also reach the second goal of an order of magnitude reduction in the required energy degrees of freedom. The Combined-Pins model does not reach the error goal within the order goal.

The required orders for most cases is reduced when using snapshots of only J_{left} to generate the basis as shown in Fig. 5.18. The Full-Core model requires approximately order 3 to reach the goal, while the Combined-Assemblies model requires approximately order 12. However, the performance of the Combined-Pins model is actually worse as compared to using ϕ snapshots.

Once again, the best results are obtained by combining the snapshots of the ϕ and J_{left} together as shown in Fig. 5.19. The required order for the Full-Core model remains at about 3, and the Combined-Assemblies model requires approximately

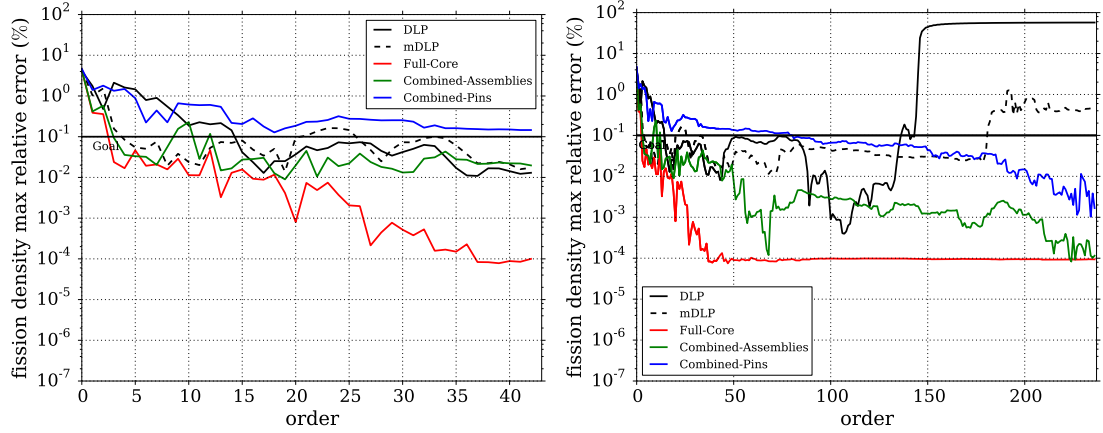


FIGURE 5.18: Relative error for 238-group, BWR-Core 0 test problem using snapshots of only J_{left}




FIGURE 5.19: Relative error for 238-group, BWR-Core 0 test problem using snapshots of both ϕ and J_{left}

order 14 to reach the error goal. The Combined-Pins model is greatly improved and reaches the error goal at about order 20. When using both types of snapshots, the Combined-Pins model reached the required order goal of an order of magnitude reduction.

5.4.2 Configuration 1

5.4.2.1 44-Group Results

Figure 5.20 shows that approximately order 7 is required for the Full-Core model to reach the error threshold. However, the practical, Combined-Assemblies and Combined-Pins models require approximately order 17 and order 26 to reach the




FIGURE 5.20: Relative error for the 44-group, BWR-Core 1 test problem using snapshots of only ϕ




FIGURE 5.21: Relative error for the 44-group, BWR-Core 1 test problem using snapshots of only J_{left}

error threshold. These results are when using snapshots of only ϕ . Both of the Combined models are far from the order goal of approximately 5 order required to meet the error threshold.

Many of the results are improved when snapshots of J_{left} are used for basis generation with the KLT instead of ϕ . These results are shown in Fig. 5.21. The Full-Core model requires approximately order 7, while the Combined-Assemblies model requires approximately order 21 to reach the error goal. The Combined-Pins model require approximately order 27.




FIGURE 5.22: Relative error for the 44-group, BWR-Core 1 test problem using snapshots of both ϕ and J_{left}

The best results are obtained when snapshots of both ϕ and J_{left} are combined together to generate the KLT basis. These results are presented in Fig. 5.22. The Full-Core model requires order 7 to reach the error goal. Both the Combined-Assemblies and Combined-Pins models reached the goal at approximately order 19, when using both types of snapshots.

5.4.2.2 238-Group Results

Figure 5.23 shows the results from using snapshots of only ϕ for the 238-group cross-section library. The impractical model of Full-Core requires approximately order 13 to reach the goal, while the Combined-Assemblies case requires approximately order 26. Each of these models also reach the second goal of an order of magnitude reduction in the required energy degrees of freedom. The Combined-Pins model does not reach the error goal within the order goal.

Again, the required orders for most cases is reduced when using snapshots of only J_{left} as shown in Fig. 5.24. The Full-Core model requires approximately order 8 to reach the goal, while the Combined-Assemblies model requires approximately order 9. The Combined-Pins model reaches the error goal by approximately order 26.

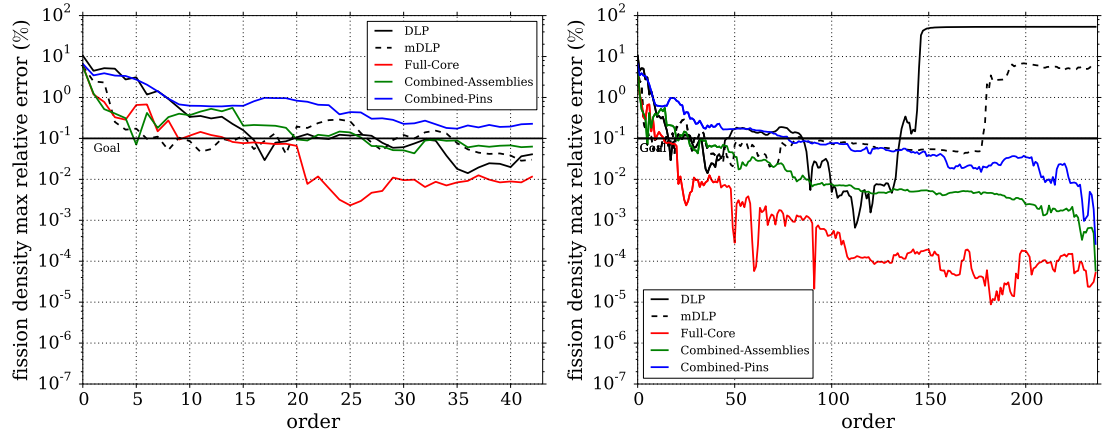


FIGURE 5.23: Relative error for 238-group, BWR-Core 1 test problem using snapshots of only ϕ

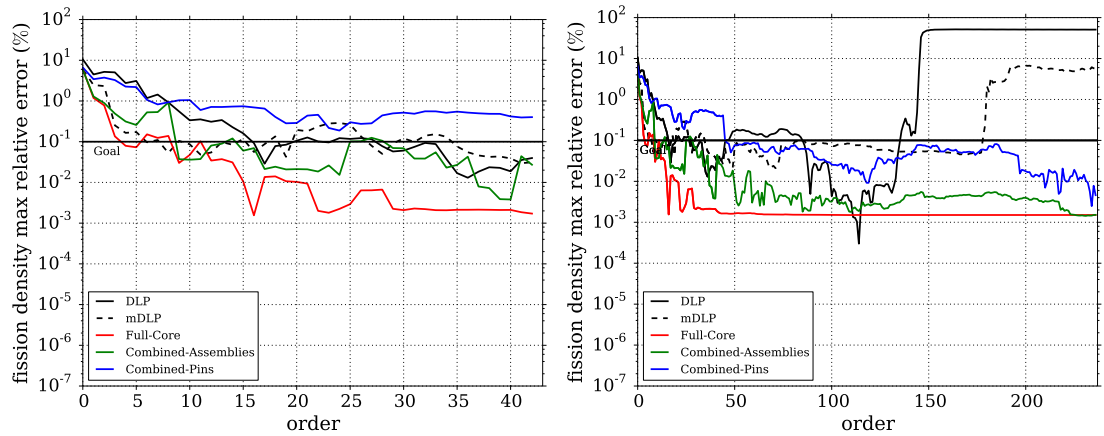


FIGURE 5.24: Relative error for 238-group, BWR-Core 1 test problem using snapshots of only J_{left}

Combining the two types of snapshots together has little effect at low orders for this configuration as shown in Fig. 5.25. The required order for the full core model remains at about 8, and the Combined-Assemblies model requires approximately order 15 to reach the error goal. The Combined-Pins model is greatly improved and reaches the error goal at about order 26. It appears that for this configuration ϕ actually worsens the expansion.

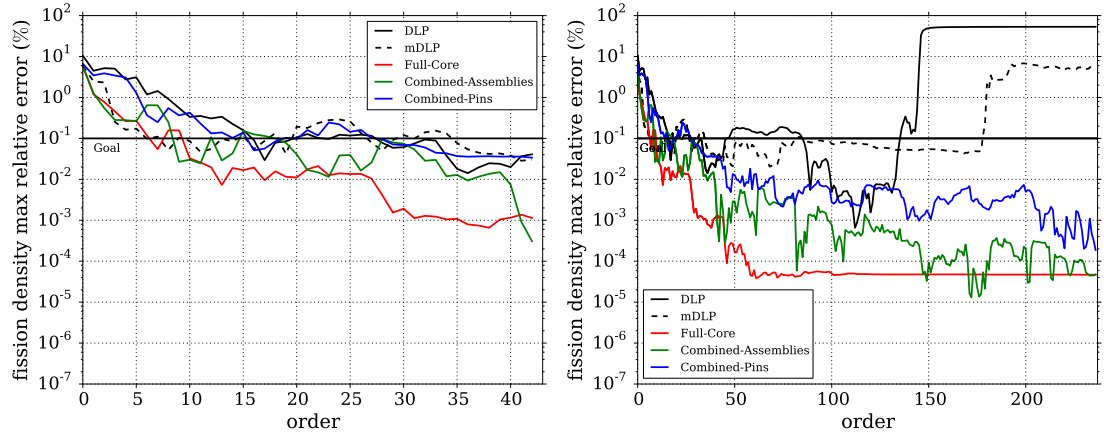


FIGURE 5.25: Relative error for 238-group, BWR-Core 1 test problem using snapshots of both ϕ and J_{left}

5.4.3 Configuration 2

5.4.3.1 44-Group Results

The final configuration of the BWR test problem leads to results very similar to the previous two configurations, but requires slightly higher-order expansion to achieve the error goal, as expected. As Fig. 5.26 shows, it takes approximately order 12 for the Full-Core model to reach the goal of sub-0.1% relative error in the fission density. Whereas the practical Combined models take approximately order 18 to reach the goal. These results are when using snapshots of only ϕ . The Combined-Pins model reaches the goal at approximately order 30. None of these models can achieve the order goal of approximately 5 orders for the 44-group test problems.

Some of the results are slightly improved if instead snapshots of J_{left} are used for basis generation with the KLT, as Fig. 5.27 shows. The Full-Core model requires approximately order 10, while the Combined-Assemblies model requires approximately order 24 to reach the goal, which is worse than the expansion using snapshots of ϕ .

The best expansions for the practical cases are obtained when snapshots of both ϕ and J_{left} are combined together to generate the KLT basis. These results are presented in Fig. 5.28. The Full-Core model requires order 12 to reach the error




FIGURE 5.26: Relative error for the 44-group, BWR-Core 2 test problem using snapshots of only ϕ




FIGURE 5.27: Relative error for the 44-group, BWR-Core 2 test problem using snapshots of only J_{left}

goal. The Combined-Assemblies and the Combined-Pins models reached the goal at approximately order 19.

5.4.3.2 238-Group Results

Figure 5.29 shows the results from using snapshots of only the ϕ for the 238-group cross-section library. The impractical model of Full-Core requires approximately order 22 to reach the goal, while the Combined-Assemblies case requires approximately order 31. Each of these models also reach the second goal of an order of




FIGURE 5.28: Relative error for the 44-group, BWR-Core 2 test problem using snapshots of both ϕ and J_{left}




FIGURE 5.29: Relative error for 238-group, BWR-Core 2 test problem using snapshots of only ϕ

magnitude reduction in the required energy degrees of freedom. The Combined-Pins model does not reach the error threshold within the order goal.

The results for using snapshots of J_{left} for basis generation are shown in Fig. 5.30. The Full-Core model requires approximately order 10 to reach the goal, while the Combined-Assemblies model requires approximately order 36. The performance of both of the Combined models are worse than for the case of ϕ snapshots.

Again, the best results are obtained by combining the snapshots of the ϕ and J_{left} together as Fig. 5.31 shows. The required order for the full core model is about 9, and the Combined-Assemblies model requires approximately order 27 to reach




FIGURE 5.30: Relative error for 238-group, BWR-Core 2 test problem using snapshots of only J_{left}

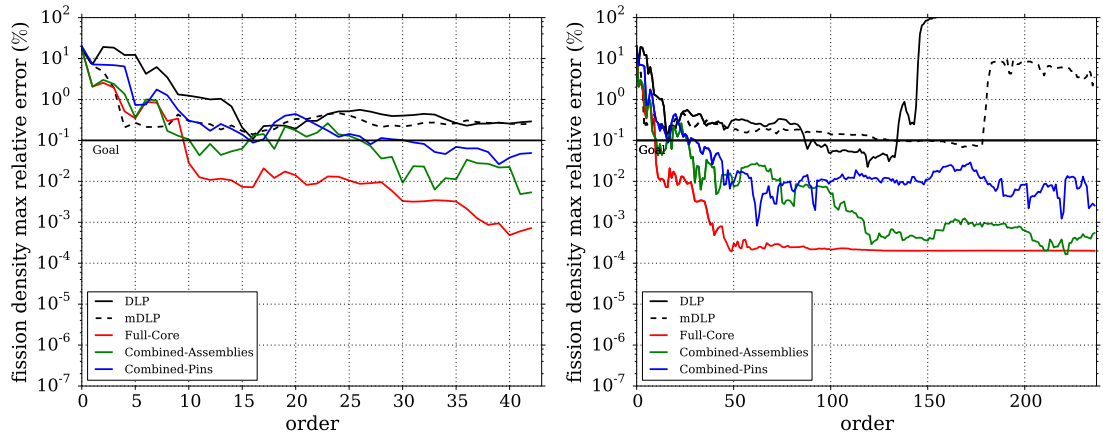


FIGURE 5.31: Relative error for 238-group, BWR-Core 2 test problem using snapshots of both ϕ and J_{left}

the error goal. The Combined-Pins model is greatly improved and reaches the error goal at about order 27. Clearly from this section, it is difficult to achieve an order of magnitude reduction in the required energy degrees of freedom and still reach the error goal for configuration 2. The inhomogeneity simply requires more information to model accurately.

5.5 Higher-Order Angular Moments

After considering the effect of ϕ and J_{left} snapshots, the next step is to compare snapshots of higher-order angular moments. These snapshots are generated by an angular expansion by Jacobi basis functions of the angular flux. In this case, the




FIGURE 5.32: Relative error for 44-group, 10-pin test problem using snapshots from the 10-pin model. Sets of snapshots are used separately for basis generation

zeroth order is exactly J_{left} , while the higher moments do not have direct physical corollaries. These snapshots were used for each test problem for basis generation.

5.5.1 10-pin Problem

In this section, The performance of using additional moments is compared for only the Full 10-pin model, as the resulting conclusions are the same for the excluded figures. Each of the figures in this section are presented with only a single snapshot model that uses various schemes to combine the snapshots as described for each figure.

5.5.1.1 44-Group Results

The individual performance of each set of snapshots for the 44-group 10-pin test problem are presented in Fig. 5.32. The performance of ϕ and J_{left} (0th moment) snapshots are identical to the results presented previously. As shown, the higher-order moment snapshots do not perform well, and require at least 30th order to reach the error threshold.




FIGURE 5.33: Relative error for 44-group, 10-pin test problem using snapshots from the 10-pin model. Sets of snapshots are combined together for basis generation

When used together with the snapshots of ϕ and J_{left} , the results are presented in Fig. 5.33. The performance of the higher-order moments are nearly identical to that of combining only ϕ and J_{left} , which is labeled as 0th moment in the figure.

5.5.1.2 238-Group Results

This section presents the results of using the higher-order moment snapshots for the 238-group 10-pin test problem. All of the figures in this section are shown with side-by-side comparisons of the same data, one shown to 43th order, while the other presents the full spectrum up to 237th order. Figure 5.34 presents the results from using the various sets of snapshots separately. Similarly to the 44-group results, the higher-order moment snapshots do not perform as well as the results of ϕ and J_{left} snapshots.

Figure 5.35 presents the results of combining the various sets of snapshots. The success of the snapshots of higher-order moments do not significantly change the results from combining only the snapshots of ϕ and J_{left} .




FIGURE 5.34: Relative error for 238-group, 10-pin test problem using snapshots from the 10-pin model. Sets of snapshots are used separately for basis generation




FIGURE 5.35: Relative error for 238-group, 10-pin test problem using snapshots from the 10-pin model. Sets of snapshots are combined together for basis generation

5.6 BWR Test Problem

In this section, The performance of using additional moments is compared for only the Full-Core ϕ model. Each of the figures in this section are presented with only a single snapshot model that uses various schemes to combine the snapshots as described for each figure. The performance of the higher-order moment snapshots are compared between each of the three core configurations as discussed previously.




FIGURE 5.36: Relative error for 44-group, BWR-Core 0 test problem using snapshots from the Full-Core model. Sets of snapshots are used separately for basis generation

5.6.1 Configuration 0

5.6.1.1 44-Group Results

The individual performance of each set of snapshots for the 44-group BWR-Core 0 test problem are presented in Fig. 5.36. The performance of ϕ and J_{left} (0th moment) snapshots are identical to the results presented previously. As shown, the higher-order moment snapshots perform similarly to snapshots of ϕ .

When used together with the snapshots of ϕ and J_{left} , the results are presented in Fig. 5.37. The performance of the higher-order moments are nearly identical to that of combining only ϕ and J_{left} , which is labeled as 0th moment in the figure.

5.6.1.2 238-Group Results

This section presents the results of using the higher-order moment snapshots for the 238-group BWR-Core 0 test problem. All of the figures in this section are shown with side-by-side comparisons of the same data, one shown to 43th order, while the other presents the full spectrum up to 237th order. Figure 5.38 presents the results from using the various sets of snapshots separately.

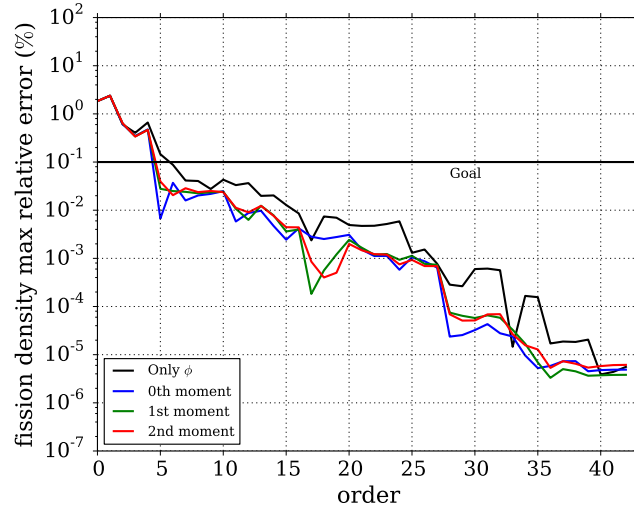


FIGURE 5.37: Relative error for 44-group, BWR-Core 0 test problem using snapshots from the Full-Core model. Sets of snapshots are combined together for basis generation




FIGURE 5.38: Relative error for 238-group, BWR-Core 0 test problem using snapshots from the Full-Core model. Sets of snapshots are used separately for basis generation

Figure 5.39 presents the results of combining the various sets of snapshots. The success of the snapshots of higher-order moments do not significantly change the results from combining only the snapshots of ϕ and J_{left} .




FIGURE 5.39: Relative error for 238-group, BWR-Core 0 test problem using snapshots from the Full-Core model. Sets of snapshots are combined together for basis generation




FIGURE 5.40: Relative error for 44-group, BWR-Core 1 test problem using snapshots from the Full-Core model. Sets of snapshots are used separately for basis generation

5.6.2 Configuration 1

5.6.2.1 44-Group Results

The individual performance of each set of snapshots for the 44-group BWR-Core 0 test problem are presented in Fig. 5.40. The performance of ϕ and J_{left} snapshots are identical to the results presented previously. As shown, the higher-order moment snapshots perform similarly to snapshots of ϕ .




FIGURE 5.41: Relative error for 44-group, BWR-Core 1 test problem using snapshots from the Full-Core model. Sets of snapshots are combined together for basis generation

When used together with the snapshots of ϕ and J_{left} , the results are presented in Fig. 5.41. The performance of the higher-order moments are nearly identical to that of combining only ϕ and J_{left} , which is labeled as 0th moment in the figure.

5.6.2.2 238-Group Results

This section presents the results of using the higher-order moment snapshots for the 238-group BWR-Core 0 test problem. All of the figures in this section are shown with side-by-side comparisons of the same data, one shown to 43th order, while the other presents the full spectrum up to 237th order. Figure 5.42 presents the results from using the various sets of snapshots separately.

Figure 5.43 presents the results of combining the various sets of snapshots. The success of the snapshots of higher-order moments do not significantly change the results from combining only the snapshots of ϕ and J_{left} .




FIGURE 5.42: Relative error for 238-group, BWR-Core 1 test problem using snapshots from the Full-Core model. Sets of snapshots are used separately for basis generation




FIGURE 5.43: Relative error for 238-group, BWR-Core 1 test problem using snapshots from the Full-Core model. Sets of snapshots are combined together for basis generation

5.6.3 Configuration 2

5.6.3.1 44-Group Results

The individual performance of each set of snapshots for the 44-group BWR-Core 0 test problem are presented in Fig. 5.44. The performance of ϕ and J_{left} snapshots are identical to the results presented previously. As shown, the higher-order moment snapshots perform similarly to snapshots of ϕ .

When used together with the snapshots of ϕ and J_{left} , the results are presented in Fig. 5.45. The performance of the higher-order moments are nearly identical to




FIGURE 5.44: Relative error for 44-group, BWR-Core 2 test problem using snapshots from the Full-Core model. Sets of snapshots are used separately for basis generation




FIGURE 5.45: Relative error for 44-group, BWR-Core 2 test problem using snapshots from the Full-Core model. Sets of snapshots are combined together for basis generation

that of combining only ϕ and J_{left} , which is labeled as 0th moment in the figure.

5.6.3.2 238-Group Results

This section presents the results of using the higher-order moment snapshots for the 238-group BWR-Core 0 test problem. All of the figures in this section are shown with side-by-side comparisons of the same data, one shown to 43th order,




FIGURE 5.46: Relative error for 238-group, BWR-Core 2 test problem using snapshots from the Full-Core model. Sets of snapshots are used separately for basis generation




FIGURE 5.47: Relative error for 238-group, BWR-Core 2 test problem using snapshots from the Full-Core model. Sets of snapshots are combined together for basis generation

while the other presents the full spectrum up to 237^{th} order. Figure 5.46 presents the results from using the various sets of snapshots separately.

Figure 5.47 presents the results of combining the various sets of snapshots. The success of the snapshots of higher-order moments do not significantly change the results from combining only the snapshots of ϕ and J_{left} .

5.7 Conclusion

As the 1-D test problems show, KLT basis sets are highly effective. Nearly all the of the KLT basis sets outperform the mDLP and DLP basis sets, which is expected due to the amount of information captured by the KLT. In nearly all cases, the best performing KLT basis functions are those based on snapshots of ϕ combined with snapshots of J_{left} .

It is apparent that including additional information in the snapshots will improve the results only if the new information is distinct from the current set of snapshots. The higher-order moments did not improve the overall performance of the basis set because the information contained in the higher-order moments was already captured in the 0th moments (J_{left}).

Furthermore, successful KLT basis functions contain all available material types. As the results from the 10-pin test problem show, disregarding snapshots from a fuel type leads to poor expansions of the actual solution.

Finally, as expected, more difficult problems require a larger number of basis functions for an expansion of the same relative error. As the difficulty of the BWR models increases, the number of required degrees of freedom to meet the goal were also increased. Thus, the success of the KLT is problem dependent.

Chapter 6

Results for 2-D Studies

The goal for this section was to see how applicable the method of snapshots and the KLT was to 2-D problems. The C5G7 benchmark was adapted to use a 44-group cross-section library, which greatly increased the size of the problem, and thus the computation time. Only the 44-group library was used to generate results due to the long computation times necessary to solve for each case.

Further, in the interest of time, the results for this section have been truncated to 10th order. The overarching goal of the work was to achieve sub-0.1% relative pin power errors while reducing the necessary energy degrees of freedom by an order of magnitude. As such, 10th order is sufficient to determine if KLT can reach the goal in a 2-D problem-space.

As mentioned previously, the response matrix solution in all 2-D calculations was only taken to second order in each angle as well as space, which means that the pin powers and fission densities solved by SERMENT will differ greatly from the original C5G7 results. However, the reference case and the snapshot models utilized the same spatial and angular order, thus the error observed (with respect to the SERMENT reference) is assumed only to be a function of the energy expansion order.




FIGURE 6.1: Flux spectrum for the C5G7 problem using 44-group cross-section library

6.1 Energy Spectra for the Test Problem

Each of the snapshot models and the full 44-group benchmark problem were solved first to provide the spatially averaged flux profiles, which were used for the mDLP results. This spatially averaged spectrum is presented in Fig. 6.1. As compared to the 1-D test problems, the spatially averaged spectra are quite distinct between the various models for the 2-D problem. As such, it is expected that mDLP will not perform as well, and additionally that effective basis functions will be difficult to attain.

6.2 Nodal Fission Density Results

In order to determine the effective constituents for basis function construction, several versions of basis functions are derived and compared as was done for the 1-D cases. In this section, each of the results will compare the maximum relative error of the nodal fission density for a snapshot model. The nodal fission density is almost equivalent to the assembly power, and the reference is taken to be the nodal fission density from a full multigroup solution of the C5G7 problem using the 44-group cross-section library. First, the results of utilizing snapshots of only




FIGURE 6.2: Relative error in fission density for 44-group, C5G7 test problem using snapshot of only ϕ .

the scalar flux ϕ to generate the KLT basis functions are presented in Fig. 6.2. In this figure, the maximum relative error of the nodal fission density is shown as a function of energy degree of freedom.

By nine degrees of freedom, none of the models have consistently reached the goal of sub-0.1% errors; however, several models perform well comparatively. The Small-Core model performs well despite spatially reducing the problem by a factor of four. Surprisingly, one of the best performing models is the Combined-Pins model despite the simplicity of the model. Both of these snapshot models outperform mDLP and can reach nodal fission errors of sub-1% by at least 7th order. It would be interesting to view the relative performance of the snapshot models are higher-order, but cannot be provided here due to time constraints.

As Fig. 6.2 shows, the results of the Small-Core and Reduced Small-Core are quite similar, with the Small-Core results being slightly better in terms of the maximum relative error. Due to this comparison, it may be assumed that the Reduced Full-Core results are approximately equal to the Full-Core results, and thus may be used as a comparison point to the best that a basis set can expand the C5G7 solution.

When the snapshots of only the upward and downward partial current are used for basis generation, the results are shown in Fig. 6.3. Many of the cases have

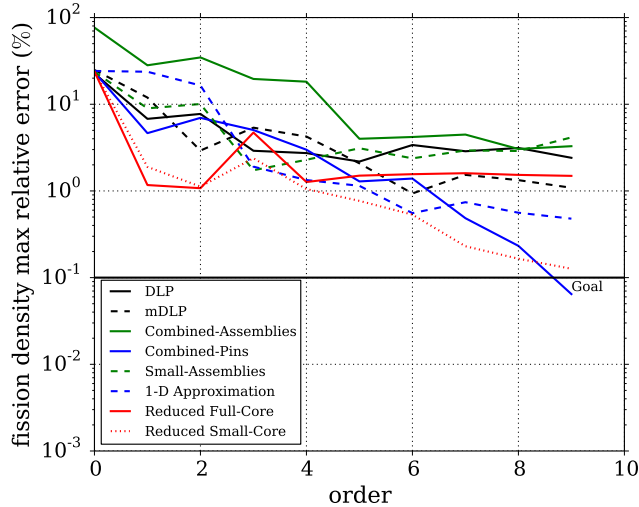


FIGURE 6.3: Relative error in fission density for 44-group, C5G7 test problem using snapshot of J_{up} and J_{down} .

improved as compared to using snapshots of only ϕ including Combined-Pins and Reduced Small-Core. However, some models have reduced effectiveness including Combined-Assemblies and Reduced Full-Core. Unlike the 1-D case, the partial current cannot be assumed symmetric in all directions, thus the partial current in the up and down directions were chosen. The combination of up and down was equivalent due to symmetry to the combination of left and right, thus the choice was made arbitrarily.

Figure 6.4 presents the results from using snapshots from ϕ , J_{up} , and J_{down} . The inclusion of the partial current had varying effects of the success of each snapshot model, e.g., the reduced small-core model was improved in general, while the small-assemblies was worsened. The Combined-Pins results were relatively unchanged, and performed surprisingly well despite the model simplicity. The best performing basis sets in general are those that include the information of the partial current combined together with the scalar flux.

6.3 Pin Power Results

For an additional comparison for the C5G7 problem, the maximum relative error in the pin powers is compared for each of the snapshot models. The reference case




FIGURE 6.4: Relative error in fission density for 44-group, C5G7 test problem using snapshot of ϕ , J_{up} , and J_{down} .

for this section is taken as the pin powers from a full multigroup solution of the C5G7 problem using the 44-group cross-section library. It is expected that the results in this section to be inferior to those in the previous section. Resolving the pin powers is more difficult than resolving the nodal fission density, thus more energy degrees of freedom should be required to achieve the same level of accuracy.

Figure 6.5 shows the results utilizing snapshots of only the scalar flux ϕ . As expected, the results are approximately an order of magnitude worse than the equivalent results solving for the fission density. However, the same snapshot models all appear to perform the same relative to each other. The best performing case is the Reduced Small-Core, while the 1-D Approximation also performs well. Also note that the Reduced Small-Core results are nearly identical to the Small-Core results, which suggests that the Reduced Full-Core results may be nearly equivalent to the Full-Core results. However, more energy degrees of freedom are needed to resolve the pin powers to the desired accuracy.

The performance is somewhat improved when considering only snapshots of the partial current in basis generation as shown in Fig. 6.6. For this figure, the Small-Core results cannot be shown because too many snapshots are available and cannot all be used without running into memory issues. The Reduced Small-Core results must be used as a supplement for the Small-Core results. The Combined-Pins and




FIGURE 6.5: Relative error in pin power for 44-group, C5G7 test problem using snapshot of only ϕ .




FIGURE 6.6: Relative error in pin power for 44-group, C5G7 test problem using snapshot of both J_{up} and J_{down} .

Reduced Small-Core models perform quite well compared to the other snapshot models.

When combining snapshots of ϕ , J_{up} , and J_{down} , the relative error in pin power resolution is slightly improved as Fig. 6.7 shows. The Small-Core model reaches sub-1% relative error by approximately 7th order, provided that the error does not increase past the orders shown in the figure. When using all available snapshots, most cases perform better than using only one type of snapshot alone.

In order to provide some context for the C5G7 pin power results, a pin power heat




FIGURE 6.7: Relative error in pin power for 44-group, C5G7 test problem using snapshot of both ϕ , J_{up} , and J_{down} .




FIGURE 6.8: Pin power heat map for the DETRAN reference solution. The upper left corner is the center of the core

map for the test problem was created for each of the reference cases. The heat map was designed to view the location of the maximum errors in the test problem and related cases. The first heat map is presented in Fig. 6.8. As expected the greatest pin powers are present in the UO_2 assembly at the center of the model.

The error of the SERMENT reference solution to the DETRAN solution is shown




FIGURE 6.9: Error of pin powers in the SERMENT reference solution relative to the DETRAN reference solution. The upper left corner is the center of the core

in Fig. 6.9. In this figure, the center of the model is the upper left corner, and the moderator sections of the model have been omitted. The greatest error for the SERMENT reference solution occurs at the bottom right corner of the figure, which represents a UO_2 pin adjacent to the moderator section.

The error in the pin powers of the best performing KLT case (9th order Reduced Small-Core) using only ϕ snapshots relative to the SERMENT reference solution is presented as a heat map in Fig. 6.10. In this figure, the maximum error is shown at the center of the figure, which corresponds to the center of the reactor quadrant. The KLT solution also has a relatively large error at the assembly boundaries.

The error heat map for using snapshots of only the partial currents J_{up} and J_{down} is presented in Fig. 6.11. As shown, on average, the pin powers are predicted well for the majority of the pins with the exception of the center of the core. The partial current appears to lead to better predictions of the pin powers in the 2-D model.

Finally, the error of the best performing KLT case from using J_{up} and J_{down}




FIGURE 6.10: Error in the pin powers of the best performing KLT case (9th order, Reduced Small-Core, snapshots of ϕ) relative to the SERMENT reference solution. The upper left corner is the center of the core

snapshots and ϕ snapshots (9th order, Reduced Small-Core) is presented in Fig. 6.12. The maximum error for the case of including the partial current is reduced from the case of only using snapshots of ϕ , and the location of the maximum error has moved to the center of the core model.

6.4 Conclusion

The KLT can be effective as compared to the results of mDLP for 2-D models. In general it seems that more degrees of freedom are required in general for the C5G7 test problem as compared to the 1-D test problems, but that is expected due to the increased difficulty of modeling 2-D problems.

In general, it seems that the size of the snapshot model is not a strong predictor of the success of a basis set. The Combined-Pins model performs surprisingly well




FIGURE 6.11: Error in the pin powers of the best performing KLT case (9th order, Reduced Small-Core, snapshots of J_{up} and J_{down}) relative to the SERMENT reference solution. The upper left corner is the center of the core

despite the model simplicity, and in some cases is the best performing in terms of the smallest maximum error at order 9.

The largest errors typically occur at material and assembly junctions, which is expected. The C5G7 benchmark was designed with difficult to capture junctions, to increase the difficulty of the benchmark. It is interesting that the including the partial current snapshots moves the largest relative error to the center of the core instead of the material junctions.

Similar the the 1-D results, the best performing basis functions are attained on average by including all available snapshots into basis generation. However spatially averaging the snapshots does not appear to have a strong effect on the success of the basis set. This is also expected because the snapshots provided by neighboring spatial cells are expected to be exceedingly similar.




FIGURE 6.12: Error in the pin powers of the best performing KLT case (9th order, Reduced Small-Core, snapshots of ϕ , J_{up} , and J_{down}) relative to the SERMENT reference solution. The upper left corner is the center of the core

Chapter 7

Conclusions and Future Work

7.1 Summary

The overarching goal of this work was to find an effective basis with which to expand the energy variable for the eigenvalue response matrix method such that the required energy degrees of freedom are reduced by an order of magnitude while achieving sub-0.1% relative error in the pin powers or fission densities. To this end, the Karhunen Loéve Transform was explored to create the set of vectors for the energy variable. The KLT basis was compared to the discrete Legendre polynomials and the modified DLPs.

To base the comparison, three test problems were constructed including a (1) 1-D 10-pin model representative of a UO₂-MOX junction, (2) a 1-D, 70-pin model representative of a BWR core, and (3) a 2-D model based on the C5G7 benchmark expanded to 44 energy groups. The difficulty of the models increased in the order of presentation, thus the 10-pin test problem was the simplest and the C5G7 benchmark was the most difficult. This difficulty manifested in an increase in the number of required energy degrees of freedom necessary to achieve sub-0.1% maximum errors relative to a full multigroup solution of the particular test problem.

In nearly all cases, the KLT basis outperformed the mDLP and DLP bases for ERMM energy expansions. This result is not surprising due to the method of

constructing the KLT basis. The KLT identifies the most common characteristics of a set of functions (or snapshots in the case of this work) and creates the basis functions from those characteristics. KLT is ordered such that the most important information of a problem-space is contained in the low-order basis vectors, thus in general the KLT can achieve high-fidelity expansions with few orders. The downside to KLT is the requirement of initial data with which to create the basis vectors. Since the actual solution is deemed unavailable (because requiring the solution to approximate the solution is a waste of effort for this application), a number of representative snapshot models were constructed for each of the test problems.

The choice of snapshot model greatly impacts the effectiveness of a KLT basis, and thus the results of an expansion in that basis. In general, snapshot models should be computationally quick to solve, yet similar to the test problem. The results presented in the body of this work indicate that to be effective, a snapshot model does not need to be a reduced version of the test problem, but the snapshot model should contain features consistent with the test problem.

For example, consider the results of the 10-pin snapshot models. The results of the UO₂ and MOX models were not desirable as each of the models was lacking the information of one of the fuel types of the 10-pin problem; however, the results were greatly improved when the snapshots of those two models were included together with the snapshots representing the UO₂-MOX pin junction. This trend remains consistent throughout the remainder of the test problems. The small snapshot models can be as effective as the larger models provided that the unique features of a test problem are captured by the snapshot model.

In regards to the overarching goal, the KLT can be used to construct an effective basis set for expansion in ERMM, and the goal of an order of magnitude reduction in the required energy degrees of freedom can be achieved in some cases. Unfortunately, those cases typically utilized the solution of the test-problem as snapshots for basis generation, which is not a realistic use of the KLT basis for this application. However, the KLT basis functions utilizing the test problem represent the best

approximation of a truncated basis expansion for a given order. In other words, there is not another basis set that can outperform a basis set created by KLT using snapshots of the test problem.

Many of the smaller snapshot models presented throughout this work provided encouraging results due to the simplicity of the snapshot model, particularly the 1-D approximation to the 2-D test problem. Such a method could be used to greatly improve the computation time of response matrix methods. As mentioned previously, response matrix solutions are very slow compared to other solution methods e.g., discrete ordinance unless the phase space is significantly reduced. Using KLT for energy can provide a greater reduction in the the required energy degrees of freedom as compared to the modified DLPs provided that the snapshot model(s) are chosen well.

The results also indicated that including additional information in the form of snapshots will typically improve the results as long as the extra information is relevant to the expansion. For example, the scalar flux expansion was improved by including the leftward partial current; however, the results remained relatively unchanged if additional higher order angular moments were also included in the basis generation. This suggests that adding more snapshots improves the data only if the new snapshots are unique and relevant. Therefore, the inclusion of partial current snapshots improved results because the snapshots differed meaningfully from scalar flux snapshots. Contrarily, the use of a greater number of snapshots from a finer spatial mesh yielded diminishing improvement because the new snapshots became increasingly similar.

The final conclusion to make is when considering test problems using different energy groupings (i.e., 44-group and 238-group), the resulting relative fission density errors were approximately the same between the two libraries, suggesting that KLT efficiency is almost independent of the number of groups in the library. KLT can capture a large amount of physics information in low-order vectors, and, thus, KLT can perform as well or better with the inclusion of additional information from a higher group library compared to a lower group library.

7.2 Future Work

The effort in expanding the energy variable has provided some insight, but ultimately more work is required to finish the exploration of the KLT as a basis for expanding the energy variable. There are a few direct extensions of work presented here that would be of substantial value to this area of research.

The first such project is to further explore the KLT as applied to 2-D problems. Much of the 2-D work presented here was truncated due to a time requirement. The C5G7 results would benefit from computations at higher spatial and angular order to allow for better congruence between the ERMM solution and other accepted solutions for the C5G7 benchmark.

In addition, it would be of substantial value to explore other energy grouping to determine a dependence of KLT on the number of energy groups in an expansion. Although the results here suggested a weak dependence, further research is needed to confirm these findings.

Any ongoing work with KLT will additionally focus on snapshot generation and effectiveness. Most problems of interest contain a sufficiently large number of unique spatial regions for use in snapshot generation; therefore, a key goal of ongoing work is to determine other variables (e.g, increased angle dependence) that can provide greater snapshot variety.

Another interesting test of KLT performance would be to use each of the KLT basis sets to expand the SERMENT solution vectors directly, instead of solving for the pin powers or fission densities. This would be interesting to ensure that the behavior is similar to the previously presented results. It is assumed that the general trend will be similar, but the nonlinear behavior would likely differ as a function of expansion order as compared to the results presented here.

Finally, this work has focused solely on the expansion of the energy variable. However, ERMM requires expansion in all phase-space variables, and it would be interesting to compare the effectiveness of a KLT expansion for either the spatial or

angular variable. Provided that snapshots can be obtained, the KLT is not specific to any particular variable, thus is expected to provide an effective expansion in any phase-space variable to which it is applied.

Bibliography

- [1] A. Shimizu. Response matrix method. *Journal of Nuclear Science and Technology*, 5:1–10, 1963.
- [2] A. Shimizu, K. Monta, and T. Miyamoto. Application of the response matrix method to criticality calculations of one-dimensional reactors. *Journal of Nuclear Science and Technology*, 5:11–17, 1963.
- [3] J.A. Roberts and B. Forget. Solving eigenvalue response matrix equations with nonlinear techniques. *Annals of Nuclear Energy*, 69:97–107, 2014.
- [4] J.A. Roberts. A high-order, time-dependent, response matrix method for reactor kinetics. *Nuclear Science and Engineering*, 179(3), 2015.
- [5] K. Ishii, T. Hino, T. Mitsuyasu, and M. Aoyama. Three-dimensional direct response matrix method using a monte carlo calculation. *Journal of Nuclear Science and Technology*, 46(3):259–267, 2009.
- [6] B. Forget. *A Three Dimensional Heterogeneous Coarse Mesh Transport Method for Reactor Calculations*. PhD thesis, Georgia Institute of Technology, 2006.
- [7] B. Forget, G. Ilas, and F. Rahnema. A heterogeneous coarse mesh transport solution for a 2-d candu-6 benchmark problem. *CNS Sixth International Conference on Simulation Methods in Nuclear Engineering*, 2004.
- [8] J. Roberts. *Advanced Response Matrix Methods for Full Core Analysis*. PhD thesis, Massachusetts Institute of Technology, 2014.
- [9] L. Zhu and B. Forget. A discrete generalized multigroup energy expansion theory. *Nuclear Science and Engineering*, 166:239–253, May 2010.

- [10] L. Zhu and B. Forget. An energy recondensation method using the discrete generalized multigroup energy expansion theory. *Annals of Nuclear Energy*, 38:1718–1727, May 2011.
- [11] R.D. Dony. Karhunen-loève transform. In K.R. Rao and P.C. Yip, editors, *The Transform and Data Compression Handbook*, chapter 1. CRC Press LLC, 2001.
- [12] A.G. Buchan, C.C. Pain, F. Fang, and I.M. Navon. A pod reduced-order model for eigenvalue problems with application to reactor physics. *International Journal for Numerical Methods in Engineering*, 95:1011–1032, July 2013.
- [13] L. Sirovich. Turbulence and the dynamics of coherent structures, part iii; dynamics and scaling. *Quarterly of Applied Mathematics*, 45:583–590, 1987.
- [14] C. Eckart and G. Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936.
- [15] M. Meyer and H. G. Matthies. Efficient model reduction in non-linear dynamics using the karhunen-loève expansion and dual-weighted-residual methods. *Computational Mechanics*, December 2002.
- [16] Nuclear Energy Agency. Benchmark on deterministic transport calculations without spatial homogenization. *Nuclear Science*, NEA/NSC/DOC(16), 2005.
- [17] Scale: A modular code system for performing standardized computer analyses for licensing evaluations. Technical Report ORNL/TM-2005/39, Version 5.1, Vols. I-III, Available from Radiation Safety Information Computational Center at Oak Ridge National Laboratory as CCC-723, November 2006.
- [18] J.A. Roberts and B. Forget. Multigroup diffusion preconditioners for multiplying fixed-source transport problems. *Journal of Computational Physics*, 271:455–472, 2014.
- [19] E. M. Nichita and F. Rahnema. A finite element method for boiling water reactors. *International Conference on the Physics of Nuclear Science and Technology*, 1:351, 1998.

- [20] D. Ilas and F. Rahnema. A heterogeneous coarse mesh transport method. *Transport Theory and Statistical Physics*, 32:445–471, 2003.

Appendix A

Parametric Studies


Over the course of the project, several parametric studies were devised and implemented to ensure reasonable choices for each of the parameters discussed in this Appendix. The parametric studies may be grouped into two sections: snapshot parameters, and database parameters.

A.1 Snapshot Parameters


A.1.1 Sensitivity to Snapshot Selection

The first snapshot study was designed to determine which snapshots created the most effective basis sets. For this study, a variable number of snapshots from the baseline BWR test problem were chosen. These snapshots were then used to generate KLT basis sets, which were used to reconstruct the pin powers for the BWR test-problem. The available snapshots totaled 1960 (28 snapshots from each of the 70 pins), however only half of the snapshots were unique due to the problem symmetry.


The study used several snapshot selection schemes to select from the total snapshots. The first scheme was to use evenly spaced snapshots (i.e., only every x th snapshot was used to generate the basis). The second used a random selection of a given




(A) Using evenly spaced snapshots



(B) Using randomly selected snapshots



(C) Using beginning snapshots



(D) Using middle snapshots

FIGURE A.1: Relative error for 238-group, BWR test problem when using various snapshot selection schemes. The schemes converge as the number of snapshots increases. The legend numbers correspond to the number of snapshots selected for KLT basis generation.

number. The third used the first x snapshots, and the last scheme used x snapshots from the center (because the problem was spatially symmetric).

As shown in Figs. A.1a and A.1b, the resulting impact to fission density errors was negligible for the first two schemes; however, the last two schemes produced a greater variation until snapshots from all types of cells were included as shown in Figs. A.1c and A.1d. In other words, inclusion of additional snapshots did not improve the results unless new snapshots were meaningfully distinct from the current set of snapshots. This study suggests that including all the the available snapshots for a given model will provide the best basis sets, at a slight cost of computation time generating a basis using additional data.




FIGURE A.2: Relative error for 44-group, 10-pin test problem using snapshots from a spatially reduced model. The legend numbers correspond to the total number of available snapshots. The problem was reduced from 280 total snapshots.

A.1.2 Sensitivity to Fineness of Spatial Mesh

To determine the effect of adjusting the spatial mesh to produce various numbers of snapshots per pincell, a second parametric study was created. This study used the 44-group, 10-pin test problem. The spatial discretization was varied to produce various numbers of spatial cells and potential snapshots. As described in Chapter 4, the 10-pin test problem originally had 280 spatial cells (28 per pincell). By proportionally enlarging the mesh regions in the problem, the total number of snapshots may be reduced.

For this study, all available snapshots were used from the spatially reduced results. Results are shown in Fig. A.2, and they indicated that mesh size had a relatively small effect on the efficacy of the generated basis set. Additional snapshots from the same model were exceedingly similar to the other snapshots, which did not add additional information to the expansion.

A.2 Database Parameters

To improve computation time, databases that included responses for every order were created for each problem and model, allowing the relevant information to be read quickly without using on-the-fly computations. Responses are a function of the k eigenvalue; thus, two studies were designed that explored both the number of k values selected as well as the range of the k values. Both of these studies were based on the BWR test problem, configuration 0.

A.2.1 Sensitivity to Number of k Values in Databases

The first study for the databases explored how many values of k were needed for database success. A cubic interpolation was used in conjunction with the database. Thus, a minimum of three values of k on either side (total of 6 k values) of the true value of k for the BWR test problem were used in the database.

As shown in Fig. A.3, adding more values of k to the database as no effect on the relative fission density error. However, adding additional values of k to the database significantly increased the computation time of the databases. However, for more difficult problems, the converged value of k for a given model may be different from the true value, thus a small buffer in the number of k value was chosen by selecting eight values of k in each database, centered about the reference k value.

A.2.2 Sensitivity to Range of k Values in Databases

The second parametric study for the databases was to determine the necessary range for the k values to span. The range was centered on the reference value for k and used a total of eight values of k within the database. The results are shown in Fig. A.4. There was no difference in the fission density as a function of the spanned range. Thus a range of ± 0.15 from the reference k value was chosen




FIGURE A.3: Relative error for 238-group, BWR test problem a database of responses filled with a number of k values as indicated by the legend. The values spanned the range of ± 0.15 of the the true value of k . The number of k values had no effect of the results.




FIGURE A.4: Relative error for 238-group, BWR test problem a database of responses filled with eight k values. The values spanned the range of \pm the value in the legend of the the true value of k . The size of the range no effect of the results.

for all test problems and models. The databases were found to not influence the results with the chosen parameters.