

**NEUTRONIC MODELING AND DESIGN OPTIMIZATION OF
ADDITIVELY MANUFACTURED CONTROL ELEMENTS FOR THE
HIGH FLUX ISOTOPE REACTOR**

A Dissertation
Presented to
The Academic Faculty

by

Joseph Raymond Burns

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy
Nuclear and Radiological Engineering Program
School of Mechanical Engineering

Georgia Institute of Technology
December 2017

Copyright © 2017 by Joseph Raymond Burns

**NEUTRONIC MODELING AND DESIGN OPTIMIZATION OF
ADDITIVELY MANUFACTURED CONTROL ELEMENTS FOR THE
HIGH FLUX ISOTOPE REACTOR**

Approved by:

Dr. Bojan Petrovic, Advisor
Nuclear and Radiological Engineering
Program
George W. Woodruff School of
Mechanical Engineering
Georgia Institute of Technology

Dr. Eva K. Lee
H. Milton Stewart School of Industrial
and Systems Engineering
Georgia Institute of Technology

Dr. C.-K. Chris Wang
Nuclear and Radiological Engineering
Program
George W. Woodruff School of
Mechanical Engineering
Georgia Institute of Technology

Dr. David Chandler
Research Reactors Division
Neutron Sciences Directorate
Oak Ridge National Laboratory

Dr. Nolan Hertel
Nuclear and Radiological Engineering
Program
George W. Woodruff School of
Mechanical Engineering
Georgia Institute of Technology

Dr. Kurt A. Terrani
Fusion and Materials for Nuclear
Systems Division
Nuclear Science and Engineering
Directorate
Oak Ridge National Laboratory

Date Approved: November 6th, 2017

To my fellow Yellow Jackets, past, present, and future: if I can do it, so can you. Never stop believing in yourself.

ACKNOWLEDGEMENTS

I must first express my sincerest thanks to my advisor, Dr. Bojan Petrovic, with whom I have worked since my undergraduate years at Georgia Tech. Dr. Petrovic always believed in me through all my struggles, and encouraged me to never settle for less than my best. I give thanks to the rest of my committee as well: Dr. Wang, Dr. Hertel, and Dr. Lee, who comprise an apt cross section of the multitude of brilliant and supportive professors at Georgia Tech from whom I have had the greatest pleasure to learn; and Dr. Terrani and Dr. Chandler at ORNL, whose support and mentorship have touched me and helped me grow tremendously. I am incredibly grateful to have had the opportunity to meet and work with so many world-class researchers at ORNL, an experience which reignites my passion for learning each morning I step foot on the lab campus. I must also mention how profoundly touched I am by the support and friendship of the postdocs and fellow graduate student researchers in Dr. Terrani's group – they, too, have my greatest thanks.

My time in Knoxville would not have been the same had I not stumbled across the path of one Mariah Williford, whose unending love and encouragement have kept me going through my best and worst days. Finally, I must give my deepest thanks to my parents, Ray Burns and Carol Bruce. Needless to say, I would not be where I am today without them, and I cannot hope to articulate just how much their love and support means to me. I am astounded by how fortunate I am to be surrounded by so much love and support. I thank you all from the bottom of my heart.

The research in this dissertation was carried out with support from Laboratory-Directed Research and Development funding at ORNL as well as the ASTRO program administered by ORAU.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	iv
LIST OF TABLES	vii
LIST OF FIGURES	viii
SUMMARY	xi
CHAPTER 1. INTRODUCTION	1
1.1. Motivation: The Present State of Nuclear Power	1
1.2. Overview of the High Flux Isotope Reactor	5
1.2.1. Control Elements	8
1.3. Organization	11
CHAPTER 2. BACKGROUND	12
2.1. HFIR Control Element Fabrication	12
2.2. Ultrasonic Additive Manufacturing	16
2.2.1. Application to HFIR Control Elements	18
CHAPTER 3. CORE PHYSICS ANALYSIS	20
3.1. Methodology	20
3.1.1. Model Requirements.....	21
3.1.2. Explicit Model	23
3.1.3. Approximate Model.....	24
3.2. Physics Results	26
3.3. Summary and Discussion	37
CHAPTER 4. PERFORMANCE EVALUATION.....	39
4.1. Methodology	39
4.1.1. Depletion Model	41
4.1.2. Long-Term Depletion Approximations	42
4.1.3. Performance Characterization	46
4.2. Depletion and Performance Results	48
4.2.1. Approximating Spectra	48
4.2.2. Isotopics.....	50
4.2.3. Reactivity Worth.....	54
4.2.4. Power Peaking	59
4.2.5. Photon Heating	62

4.3. Summary and Discussion	64
CHAPTER 5. DESIGN OPTIMIZATION	66
5.1. Optimization Problem Formulation	67
5.1.1. Objective.....	67
5.1.2. Constraints.....	68
5.1.3. Problem Statement.....	69
5.2. Methodology	70
5.3. Results	76
5.3.1. Nominal Equivalent Design Sensitivities	76
5.3.2. Response Surface Fitting.....	77
5.3.3. Optimal Design.....	81
5.4. Optimal Design Assessment.....	84
5.5. Summary and Discussion.....	85
CHAPTER 6. CONCLUSION.....	86
APPENDIX A. MISCELLANEOUS NEUTRONICS PLOTS	91
REFERENCES	101

LIST OF TABLES

Table 1.1. Recently Announced Nuclear Plant Closures [6]	3
Table 2.1. HFIR CE Fabrication Cost Magnitudes, FY 2016 Dollars [17, 23]	16
Table 3.1. Explicit Model Absorber Mass Matching and Geometry	24
Table 3.2. Approximate Model Absorber Mass Matching and Geometry	25
Table 3.3. Neutron Multiplication of Homogeneous and Discrete CE Models.....	26
Table 3.4. Thermal Flux Depression in Absorber Compacts: Approximate Model.....	32
Table 3.5. Thermal Flux Depression in Absorber Compacts: Explicit Model	33
Table 4.1. Differential Worth Comparison under Explicit Depletion Approximation Scheme ...	55
Table 4.2. Differential Worth Comparison under Implicit Depletion Approximation Scheme in One Iteration	55
Table 4.3. Differential Worth Comparison under Implicit Depletion Approximation Scheme in Two Iterations	56
Table 4.4. Differential Worth Comparison under Midpoint Depletion Approximation Scheme .	56
Table 4.5. Integral Worth Comparison under Explicit Depletion Approximation Scheme.....	57
Table 4.6. Integral Worth Comparison under Implicit Depletion Approximation Scheme in One Iteration	57
Table 4.7. Integral Worth Comparison under Implicit Depletion Approximation Scheme in Two Iterations	58
Table 4.8. Integral Worth Comparison under Midpoint Depletion Approximation Scheme	58
Table 5.1. Nominal Equivalent Design Variable Settings	70
Table 5.2. Fractional Factorial Experiment Cases	73
Table 5.3. CCD Cases	74
Table 5.4. Sensitivity Cases	75
Table 5.5. Nominal Equivalent Design Sensitivity Results	77
Table 5.6. Compiled Design Perturbation Results.....	78
Table 5.7. Objective Response Surface Regression Coefficients	79
Table 5.8. Eigenvalue Response Surface Regression Coefficients.....	80
Table 5.9. Objective Response Surface Global Minimum	82
Table 5.10. Optimal Design Variable Settings	83
Table 5.11. Optimal Design Sensitivity Results	83
Table 5.12. Optimal Design BOC 400 Reactivity Worth Comparison.....	85

LIST OF FIGURES

Figure 1.1. CE Comparison	4
Figure 1.2. Neutronic Assessment Flowchart	5
Figure 1.3. HFIR Core Schematic [17].....	6
Figure 1.4. HFIR Experimental Facilities [20]	7
Figure 1.5. Ta and Eu Microscopic Absorption Cross Sections [21]	9
Figure 1.6. Tantalum Transmutation Chain [22]	9
Figure 1.7. Europium Transmutation Chain [22].....	9
Figure 1.8. Absorption Cross Sections of Selected Ta and Eu Transmutation Products [21]	10
Figure 1.9. CE Withdrawal during HFIR Operation [22].....	11
Figure 2.1. HFIR CE Fabrication Flowchart [17].....	12
Figure 2.2. Critical Path for HFIR CE Fabrication [17]	15
Figure 2.3. UAM Sonotrode (left) and Mill (right) [18].....	17
Figure 2.4. Discrete Absorber Compact	19
Figure 3.1. Planar View of Absorber Compact Array	22
Figure 3.2. Cross Section of Explicit Model CE Representation	23
Figure 3.3. Side View of Explicit Model CE Representation.....	24
Figure 3.4. Approximate Model CE Representation	25
Figure 3.5. Relative Difference in Thermal Flux: Approximate vs. Homogeneous Model (top left) and Uncertainty (bottom left); Explicit vs. Homogeneous Model (top right) and Uncertainty (bottom right)	28
Figure 3.6. Relative Difference in Fission Rate Density: Approximate vs. Homogeneous Model (top left) and Uncertainty (bottom left); Explicit vs. Homogeneous Model (top right) and Uncertainty (bottom right)	29
Figure 3.7. Relative Difference in Absorption Rate Density: Approximate vs. Homogeneous Model (top left) and Uncertainty (bottom left); Explicit vs. Homogeneous Model (top right) and Uncertainty (bottom right)	31
Figure 3.8. Relative Difference in Effective Thermal Absorption Cross Section: Approximate vs. Homogeneous Model (top left) and Uncertainty (bottom left); Explicit vs. Homogeneous Model (top right) and Uncertainty (bottom right)	35
Figure 3.9. Neutron Energy Spectra in CE Absorbing Regions	37
Figure 4.1. Predictor-Corrector Depletion Calculation Flow	40
Figure 4.2. Cycle 400 Withdrawal Schedule [19].....	42
Figure 4.3. Explicit Euler Method	45

Figure 4.4. Implicit Euler Method in One Iteration	45
Figure 4.5. Implicit Euler Method in Two Iterations.....	45
Figure 4.6. Implicit Midpoint Method	46
Figure 4.7. Approximating Spectra.....	49
Figure 4.8. Relative Difference in Isotopics: Single-Cycle Approximation vs. Explicit Simulation	50
Figure 4.9. Midpoint Scheme Lifetime Isotopics in Homogeneous Black OCE Region Closest to Core Midplane	51
Figure 4.10. Relative Difference in Isotopics: Explicit Scheme vs. Midpoint Scheme.....	52
Figure 4.11. Relative Difference in Isotopics: Implicit Scheme in One Iteration vs. Midpoint Scheme.....	53
Figure 4.12. Relative Difference in Isotopics: Implicit Scheme in Two Iterations vs. Midpoint Scheme.....	53
Figure 4.13. Absolute Difference in RFD between Homogeneous and Discrete CE Cases (left) and Uncertainty (right) at BOC 400.....	60
Figure 4.14. Absolute Difference in RFD between Homogeneous and Discrete CE Cases (left) and Uncertainty (right) at BOC 424.....	61
Figure 4.15. Absolute Difference in RFD between Homogeneous and Discrete CE Cases (left) and Uncertainty (right) at BOC 449.....	61
Figure 4.16. Local Peaking in Outermost OFE Edge (left) and Uncertainty (right) at BOC 400	62
Figure 4.17. Relative Difference in Photon Heating between Homogeneous and Discrete CE Cases (left) and Uncertainty (right) at BOC 400	63
Figure 4.18. Relative Difference in Photon Heating between Homogeneous and Discrete CE Cases (left) and Uncertainty (right) at BOC 424	63
Figure 4.19. Relative Difference in Photon Heating between Homogeneous and Discrete CE Cases (left) and Uncertainty (right) at BOC 449	64
Figure 5.1. Optimization Computational Flow	76
Figure 5.2. Absolute Difference in RFD between Homogeneous and Optimal CE Designs (left) and Uncertainty (right) at BOC 400.....	84
Figure A.1. Thermal Flux (left column) and Uncertainty (right column) in Homogeneous (top row), Approximate (middle row), and Explicit (bottom row) Models	91
Figure A.2. Fission Rate Density (left column) and Uncertainty (right column) in Homogeneous (top row), Approximate (middle row), and Explicit (bottom row) Models.....	92
Figure A.3. Absorption Rate Density (left column) and Uncertainty (right column) in Homogeneous (top row), Approximate (middle row), and Explicit (bottom row) Models.....	93

Figure A.4. Effective Thermal Absorption Cross Section (left column) and Uncertainty (right column) in Homogeneous (top row), Approximate (middle row), and Explicit (bottom row) Models.....	94
Figure A.5. BOC 400 Relative Fission Density (left column) and Uncertainty (right column) in Homogeneous (top row) and Explicit (bottom row) Models.....	95
Figure A.6. BOC 424 Relative Fission Density (left column) and Uncertainty (right column) in Homogeneous (top row) and Explicit (bottom row) Models.....	96
Figure A.7. BOC 449 Relative Fission Density (left column) and Uncertainty (right column) in Homogeneous (top row) and Explicit (bottom row) Models.....	97
Figure A.8. BOC 400 Photon Heating (left column) and Uncertainty (right column) in Homogeneous (top row) and Explicit (bottom row) Models.....	98
Figure A.9. BOC 424 Photon Heating (left column) and Uncertainty (right column) in Homogeneous (top row) and Explicit (bottom row) Models.....	99
Figure A.10. BOC 449 Photon Heating (left column) and Uncertainty (right column) in Homogeneous (top row) and Explicit (bottom row) Models.....	100
Figure A.11. BOC 400 Relative Fission Density (left) and Uncertainty (right) of Optimal CE Design	100

SUMMARY

One of the key challenges to the advancement of nuclear power technology nationally and globally is the substantial upfront capital investment required for the construction of new nuclear power plants. A research effort has been launched at the Oak Ridge National Laboratory to investigate the potential for advanced manufacturing methods to fabricate nuclear reactor components with reduced cost, time, and labor by demonstrating production of control elements (CEs) for the High Flux Isotope Reactor (HFIR) using ultrasonic additive manufacturing (UAM). At this scale, this practice offers particular potential for reducing costs associated with continuous maintenance of reactor components, considering the regular intervals of replacement of the HFIR CEs. UAM yields a unique CE design with lumped neutron absorber regions, in contrast to typical CEs with uniformly distributed absorbers, thereby necessitating an analysis of the neutronic impact of this design change. This dissertation takes on the computational modeling of additively manufactured CEs in HFIR to assess their neutronic and operational feasibility. The impact of the additively manufactured CE design on the HFIR core physics is investigated, and the performance of the new CEs (characterized by reactivity worth and core power shaping) is compared with that of the original homogeneous CEs.

It is found that while some limited changes in the HFIR core physics behavior are introduced by the additively manufactured CEs, they consistently exhibit performance that is comparable to the original CE design with no prohibitive impact to reactor safety. It is therefore concluded that the additively manufactured CEs are feasible for use in the operation of HFIR. Finally, the additively manufactured CE design is optimized by taking advantage of new design variables introduced by UAM fabrication that were not applicable to the original CE design. The additively manufactured CEs are designed to minimize their impact on the HFIR core physics so

as to provide a seamless transition between CE designs and avoid reevaluation of compliance with the well-established HFIR safety margins. This work provides a first look at the neutronic characteristics of in-core nuclear reactor components fabricated with additive manufacturing, and it is hoped that these promising results will encourage future consideration of applications of advanced manufacturing methods elsewhere in the nuclear industry to realize economic improvements.

CHAPTER 1. INTRODUCTION

1.1. Motivation: The Present State of Nuclear Power

Humanity currently faces a historical juncture. Over the past decade particularly, global leaders have accelerated and promoted international cooperative efforts to curb carbon emissions, acknowledging human impact on climate change and the severe hazards it poses to our way of life [1, 2]. The bulk of such efforts has been overtly directed toward renewable energy sources such as solar and wind power, but there is also broad, albeit muted, recognition that nuclear power must play a prominent role in the development of sustainable, environmentally responsible energy infrastructure [3-5]. As solar and wind are intermittent power sources, raising concerns regarding grid stability and energy storage capabilities, nuclear power is the only emissions-free energy source that can match the large capacity baseload generation of fossil plants [5, 6]. However, nuclear power remains a politically contentious topic, due especially to questions about waste, safety, and financing. Technology solutions already exist for these former two issues. While geological storage of spent fuel faces significant political and social hurdles [7], more effective technology solutions are available, including reprocessing technology for recovering fissile and fertile isotopes from spent fuel [8] as well as Generation IV fast reactor designs that can breed and burn fissile material from spent fuel [9, 10]. In regards to safety, technological advances, the fruit of heavy investment in research and development as well as the explosion of computing power over recent decades, have enabled significant improvements in the form of passive safety systems, integral reactor designs, advanced materials, and more. In some cases, such advances eliminate specific modes of accidents taking place, as in integral designs without large primary cooling water piping; in

others, accident scenarios are mitigated via accident-resistant materials and passive safety systems that do not require human interference for actuation or operation [11-15].

Yet, the cost of new nuclear construction remains perhaps the greatest hindrance to advancing the deployment of new nuclear technology. The most recent American projects, the Westinghouse AP1000 reactors under construction (as of this writing) at the Vogtle and V.C. Summer sites, are the first new nuclear projects in the country in decades, and are years behind schedule and billions of dollars over budget¹ [16]. Such challenges are typical of deployment of first-of-a-kind reactors, which are associated with significant uncertainties arising from the use of untested manufacturing and installation procedures, thereby discouraging commercial investment in such financially risky projects. Further, it is not only the new projects that are struggling; recent years have seen imbalanced energy markets as a result of federal and state government programs failing to credit nuclear power as an emissions-free energy source, rendering many existing nuclear plants unprofitable and unable to compete with wind, solar, and natural gas generation [6]. This has led to the announcement of several early unit closures, as listed in Table 1.1 [6]. The lost generation from these plants is in turn replaced most commonly with natural gas plants, which are not emissions-free, thereby reversing the environmental boon of the original nuclear plant [6]. Thus, in order to realize the potential of nuclear technology as a clean, reliable, sustainable energy source, leaders and researchers in the nuclear field must address the quandary of making nuclear power more economically attractive.

¹ These financial struggles ultimately led to Westinghouse declaring for bankruptcy, casting doubt regarding the future of these new build projects.

Table 1.1. Recently Announced Nuclear Plant Closures [6]

Site	Operator	Originally Licensed Final Operating Year	New Planned Closure Year
San Onofre 2 & 3	Southern California Edison	2023/2024	2013
Kewaunee	Dominion	2033	2013
Vermont Yankee	Entergy	2032	2014
Fort Calhoun	Exelon	2033	2016
Palisades	Entergy	2031	2018
Pilgrim	Entergy	2032	2019
Oyster Creek	Exelon	2029	2019
Three Mile Island 1	Exelon	2034	2019
Indian Point 2 & 3	Entergy	2013/2015	2020/2021
Diablo Canyon 1 & 2	Pacific Gas & Electric	2024/2025	2024/2025

To this end, an effort has been launched at the Oak Ridge National Laboratory to apply advanced manufacturing techniques to the fabrication of nuclear reactor components with the intent of reducing cost. As an initial demonstration, the ultrasonic additive manufacturing (UAM) process is employed for fabrication of the control elements (CEs) of the High Flux Isotope Reactor (HFIR). The current fabrication process for the HFIR CEs is labor-intensive, outdated, and expensive [17]; thus, application of UAM is expected to improve the operational economics of HFIR while demonstrating the feasibility of producing nuclear reactor components with advanced manufacturing methods. UAM is a solid-state welding process that takes place at room temperature, as opposed to the original HFIR CE fabrication process which requires hot-rolling and traditional welding steps [17]. Consequently, the CE design yielded by UAM contains an array of discrete compacts of neutron-absorbing material (as shown in Figure 1.1), which is a significant departure from the typical CE design with uniformly distributed absorbers. Therefore, as a complement to investigations of material and mechanical adequacy, it is necessary to assess the neutronic feasibility of employing this unique CE design in HFIR. The use of additively manufactured CEs in the operation of HFIR would be precluded by their failure to maintain acceptable operational and safety margins (or necessitating a drastic reassessment

thereof). This research takes on the neutronics analyses to address the operational feasibility of additively manufactured CEs in HFIR.




Figure 1.1. CE Comparison

The neutronic assessment of the additively manufactured HFIR CEs is carried out in three phases, as diagrammed in Figure 1.2. The first phase consists of the development of computational models of the additively manufactured CEs and their use in identifying the impact of the new CE design on the HFIR reactor physics in detail, so as to characterize the neutronic behavior of the new CEs. Following this exploratory phase, a thorough assessment of the performance of the additively manufactured CEs throughout their useful lifetime is undertaken. This investigation determines whether the new CE design is capable of maintaining adequate safety and performance parameters throughout the CE life, providing a definitive answer regarding the feasibility of using additively manufactured CEs in HFIR. Finally, it is attempted to optimize the additively manufactured CE design by matching the HFIR core behavior under the original CEs as closely as possible, thereby avoiding the need for any reconsideration of the established HFIR safety documentation and operational procedures. Sensitivity of the core behavior to the design parameters is investigated to inform development of a computationally efficient optimization strategy, and an optimized design satisfying the operational constraints is recommended. It is envisioned that the use of UAM in the fabrication of the HFIR CEs can

provide economic benefits to HFIR with no loss in performance while developing impetus for future consideration of applications of advanced manufacturing elsewhere in the nuclear power industry.

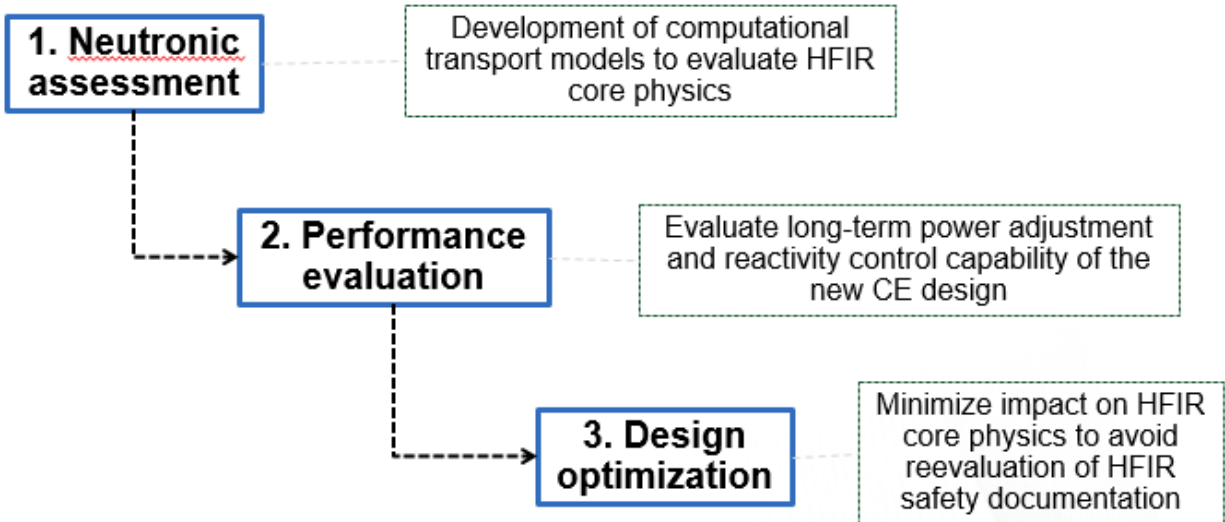


Figure 1.2. Neutronic Assessment Flowchart

1.2. Overview of the High Flux Isotope Reactor

HFIR provides one of the highest steady-state neutron sources in the world [19]; HFIR neutrons are used for a broad array of research activities, including neutron scattering, isotope production, and materials irradiation testing. HFIR is a flux trap type research reactor with a compact, high-power density, light water-cooled core operating at 85 MWt. A typical HFIR cycle runs for about 25 days and is followed by a refueling outage which typically lasts another 25 days. Figure 1.3 [17] shows a schematic of the HFIR core.




Figure 1.3. HFIR Core Schematic [17]

In addition to the fuel, water coolant/moderator, control elements, and reflector components, HFIR contains dozens of experimental facilities available for housing material samples subject to irradiation as well as beam tube penetrations for piping neutrons to external facilities. The HFIR core is designed to provide a large flux magnitude to each of these experimental sites. Figure 1.4 [20] shows the radial layout of experimental sites in the core.




Figure 1.4. HFIR Experimental Facilities [20]

The HFIR core comprises several distinct concentric annular regions. In the center of the core is the flux trap, which includes 37 radially arranged aluminum tubes serving as irradiation experiment facilities; these are housed in an aluminum basket and surrounded by moderating water. The flux trap basket is 12.7 cm in diameter. The thermal flux distribution in the core achieves a peak in the flux trap on the order of 10^{15} n/cm²-s. Surrounding the flux trap are the two fuel elements, an inner fuel element (IFE) and an outer fuel element (OFE), each about 7 cm thick. The fuel itself is a 93% enriched U₃O₈-Al dispersion loaded into individual fuel plates of 1.27 mm thickness shaped into the involute of a circle, thereby maintaining a constant cooling channel width between plates loaded into the fuel element annuli. The fuel plate clad is made of the aluminum 6061 alloy. Fuel in the IFE is loaded with boron carbide (B₄C) burnable absorber for power shaping. The IFE contains 171 individual fuel plates, while the OFE contains 369. The active height of the fuel is about 50 cm. Cooling water flows axially downward through the core

at a rate of roughly $1 \text{ m}^3/\text{s}$. Surrounding the OFE are the two CEs, an inner CE (ICE) and outer CE (OCE), each about 0.5 cm thick; these are detailed separately in the next subsection. Finally, surrounding the CEs is a large beryllium reflector which also provides neutron moderation. The reflector is about 30 cm thick and contains several beam tube penetrations and dozens of experimental facilities. The entire core assembly is submerged in a water reflector and contained in a stainless steel pressure vessel with an inner diameter of 244 cm.

1.2.1. Control Elements

The cylindrical CEs are each axially zoned into three regions according to neutron absorption opacity: the white regions contain only structural aluminum and do not significantly contribute to absorption; the gray regions, which provide fine power regulation, contain tantalum (Ta) as the primary absorber in a concentration of 30% by volume and are about 12 cm long; and the black regions, used for heavy shim and shutdown, contain europium (Eu) as the primary absorber in the form of Eu_2O_3 with a concentration of 33% by volume and are about 55 cm long. The CEs are clad in 6061 Al alloy. Holes are drilled in the gray and white regions in a regular array to mitigate the hydraulic stresses acting on the CEs from the flow of cooling water; no holes are drilled in the black regions to prevent Eu_2O_3 from reacting with water.

The 300 K ENDF/B-VII.0 microscopic absorption cross sections of the naturally occurring isotopes of Ta (^{181}Ta) and Eu (^{151}Eu and ^{153}Eu) are plotted in Figure 1.5 [21]. Neutron capture by Ta and Eu yields transmutation products with absorption cross sections that are comparable to those of their parent species; tungsten (W) and other isotopes of Ta are produced in the gray regions, and the black regions experience a buildup of samarium (Sm), gadolinium (Gd), and other isotopes of Eu. In this way, the CEs maintain absorption strength and reactivity worth over long periods of use. Transmutation chains of Ta and Eu are respectively given in

Figures 1.6 and 1.7 [22], and Figure 1.8 [21] plots the 300 K ENDF/B-VII.0 microscopic absorption cross sections of some of these isotopes.

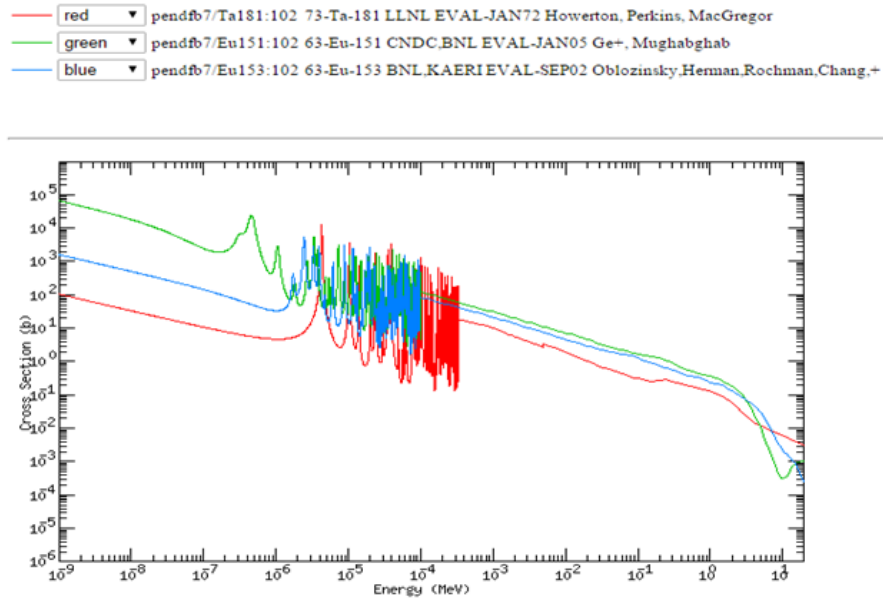


Figure 1.5. Ta and Eu Microscopic Absorption Cross Sections [21]




Figure 1.6. Tantalum Transmutation Chain [22]




Figure 1.7. Europium Transmutation Chain [22]





Figure 1.8. Absorption Cross Sections of Selected Ta and Eu Transmutation Products [21]

In the reactor, the ICE operates as a single cylinder, while the OCE is operated as four separate quadrants; any single OCE quadrant is capable of scrambling the reactor on its own. The arrangement of the absorption zones of the ICE and OCE are inverted with respect to one another, and the CEs are driven symmetrically in opposite vertical directions during operation to maintain criticality over the cycle, compensating for reactivity loss from fuel burnup by gradually driving absorbers out of the core and widening the window between the fuel and beryllium reflector. This operation is depicted in Figure 1.9 [22]. The cycle ends when the absorbing CE regions are fully withdrawn from the core and criticality can no longer be maintained.




Figure 1.9. CE Withdrawal during HFIR Operation [22]

As HFIR is a mature, well-established research reactor now with a half century of operating experience, the additively manufactured CEs should be designed to yield minimal impact to the HFIR core physics and operation so that the established safety analyses and operational procedures do not require reevaluation. This is the motivation for the optimization phase of the analysis, which aims to minimize changes in core physics yielded by the new CE design. The optimized design yielded by this process can therefore be seamlessly transitioned into the reactor following the necessary testing.

1.3. Organization

This dissertation is organized as follows. Chapter 2 compares the current CE fabrication process with UAM and discusses the ongoing research on UAM. Chapter 3 describes the computational models used for the additively manufactured CE assessment and evaluates their transport simulation results to characterize the impact of the new CE design on the HFIR core physics. Chapter 4 covers the CE performance and neutronic feasibility assessment and details the long-term depletion calculation methodology. Chapter 5 formulates the optimization problem and develops an optimal additively manufactured CE design. Chapter 6 concludes the dissertation and provides discussion and design recommendations based on the preceding results.

CHAPTER 2. BACKGROUND

2.1. HFIR Control Element Fabrication

The HFIR CE fabrication process summarized in this section is thoroughly detailed in [17]. The fabrication process currently in place has been used since the initial startup of HFIR in the early 1960's. CE fabrication campaigns take place every 7 to 9 years, each requiring 2 to 3 years and approximately \$3 million (FY 2016 dollars [23]) to complete a set of CEs, including spare components. A high-level flowchart of the CE fabrication steps is shown in Figure 2.1 [17]; each of these steps may include dozens of additional preparatory, inspection, and repair substeps.

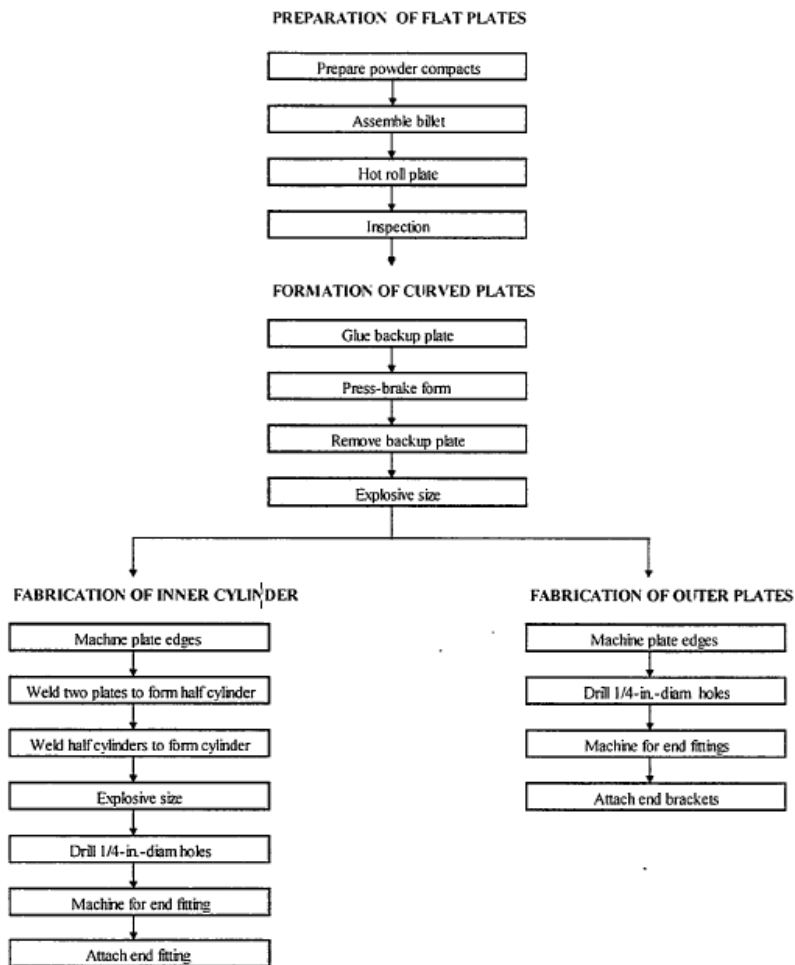


Figure 2.1. HFIR CE Fabrication Flowchart [17]

The process begins with preparation of Al, Ta, and Eu_2O_3 powders. The powders are ordered from commercial vendors and undergo thorough processing and inspection to ensure purity and achieve appropriate particle size distributions necessary for mechanical and material integrity in the final product. The absorber powders are weighed to achieve the specified absorber loading in each CE section, and these quantities are blended with the Al powder. The powder mixtures are then cold pressed into discrete rectangular compacts. An Al billet frame is prepared and heated, and the absorber compacts are loaded into the frame by hand. A cover is welded to the frame, and the billet is evacuated and sealed, at which point it is ready to undergo hot-rolling. The rolling mill is prepared and a furnace is heated to 500 °C. The billet is rolled in about 20 passes punctuated by preheat phases of prescribed times. Following annealing and shaping, the flat plate is sized and inspected in exacting detail for dimensions, defects, and microstructural integrity by visual, ultrasonic, and radiographic means; repairs are carried out as needed. This completes preparation of the flat plates.

To prepare for forming curved plates, the flat plates undergo an additional annealing and flattening phase, and an Al backup plate is prepared. The backup plate is glued to the control plate using a commercial epoxy adhesive, and the resulting bond strength is tested. The curved plates are then formed to a prescribed radius of curvature by a press-brake operation. The backup plate serves to shift the neutral stress axis out of the control plate absorbing core, thus keeping the core in compression and avoiding cracking during the forming operation. The compound plate is then heated to soften the epoxy so that the backup plate can be removed. Following cleaning, the control plate is stress annealed and inspected for proper dimensions. For sizing, a carbon steel explosive-forming die is prepared; a plate is loaded and securely fixed, and the die is sealed. The loaded die is submerged in water and the plate is explosively sized. The assembly is

removed and inspected for damage, and the explosive sizing is repeated until the plate achieves the desired contour. The final plates are then inspected for damage and dimensional conformity.

The OCE is then formed from four separate plates. The hydraulic stress-mitigating holes are drilled in a 2.54 cm (1 in) grid in the gray and white regions, each with a diameter of 0.635 cm (0.25 in). The end fittings and journal bearings are attached to the plates, and following a final cleaning and dimensional inspection, the plates are ready for use in HFIR. For the ICE, four plate quadrants must be joined to form a single cylinder. Two plates are first seam welded to form a half cylinder that is inspected and cleaned, and two half cylinders are then seam welded to form the full cylinder. The full cylinder then undergoes another explosive sizing operation and is inspected for damage and dimensions. The array of hydraulic stress-mitigating holes is drilled into the cylinder, and the drive rod bracket is installed. The completed cylinder undergoes a final cleaning and inspection before installation in HFIR.

Each step in the CE fabrication process is demanding in its own right, requiring many exacting substeps not detailed in the preceding description. Figure 2.2 [17] details the critical path for CE fabrication; with no delays or other difficulties, the quickest time to complete the full process is about 3 years. Table 2.1, adapted from [17], breaks down the magnitude of the costs in FY 2016 dollars associated with each phase of fabrication. Given the regular replacement intervals of the HFIR CEs, it is highly economically desirable in terms of both monetary and labor savings to expedite and simplify the CE fabrication process.




Figure 2.2. Critical Path for HFIR CE Fabrication [17]

Table 2.1. HFIR CE Fabrication Cost Magnitudes, FY 2016 Dollars [17, 23]

Phase	Setup and Qualification Cost	Incremental Cost
Preparation of flat plates	\$600,000	\$45,000
Formation of curved plates	\$450,000	\$45,000
Fabrication of outer plates	\$225,000	\$105,000
Fabrication of inner cylinder	\$450,000	\$525,000

2.2. Ultrasonic Additive Manufacturing

Though UAM is a relatively young technology that has not yet been applied to the production of in-core reactor components², it is a promising candidate for improving the HFIR CE fabrication process. UAM is a form of 3D printing which harnesses the energy from ultrasonic vibrations to sequentially deposit and bond layers of metal tape on top of one another [25-30]. Coupling this additive process with computer numerically-controlled (CNC) milling enables automated printing of intricate 3D geometries that could not feasibly be constructed using traditional welding methods. UAM is particularly attractive for such applications as fabrication of parts with embedded components, including fiber optics and shape memory alloys, which could not withstand the extreme conditions imparted by traditional welding. As a low-temperature solid-state process, UAM is also able to join dissimilar metals, a feat which cannot be accomplished with traditional welding [26].

The key components of a UAM printer are the sonotrode, a base plate, a metal feed tape, and a CNC mill for subtractive steps. A schematic of a typical sonotrode and mill is shown in Figure 2.3 [18]. Metal tape is fed to the horn of the sonotrode, which presses normal to the base plate and runs along its surface as the horn oscillates at ultrasonic frequencies (on the order of 20 kHz) along an axis perpendicular to the sonotrode's direction of travel. The normal force of the

² However, Siemens has successfully applied additive manufacturing to produce a pump impeller for the Krško plant in Slovenia [24], marking the first application of an additively manufactured component in an operating nuclear power plant.

sonotrode and the ultrasonic vibration of the horn yield highly localized strain and plastic deformation in the interfaces between consecutive tape layers, enabling bonding [26-28].




Figure 2.3. UAM Sonotrode (left) and Mill (right) [18]

Dehoff and Babu [27] and Fuji et al. [28] have shown that the bonding between tape layers arises from recrystallization of large grains in the highly deformed regions along tape interfaces. However, it has been found that this bonding mechanism yields mechanical properties which are weaker than those of the bulk foil material, particularly in directions perpendicular to the plane of the tape interface [26, 27, 29]. Thus, much of the ongoing UAM research activity has been directed toward mitigating such deficiencies via process parameter optimization and post-process heat treatment. Schick et al. [26] investigated the impact of the UAM process parameters, including the sonotrode travel speed, normal force, oscillation amplitude, and material preheat temperature, on bond quality and shear and tensile strength in different orientations. The results indicate that while the optimal parameter settings are application-specific, bond quality and material strength may generally be improved with a slower sonotrode travel speed, greater normal force, greater oscillation amplitude, and higher material preheat temperatures. Dehoff and Babu [27] and Wolcott et al. [30] have found that the surface roughness of the sonotrode horn also plays a role in improving bond quality by enhancing local plastic deformation. Post-process annealing was investigated by Gussev et al. [29] and Wolcott et al. [30]; their results show that build strength was improved by a range of annealing heat

treatment schedules, arising from partial elimination of the tape interfaces as well as formation of strengthening intergranular precipitates. Wolcott et al. [30] suggest one additional process consideration in overlapping tape layers to improve material strength by reducing voids.

Acceptance of the eventual practical application of UAM-fabricated components in a reactor environment by the Nuclear Regulatory Commission (NRC) would require demonstration of their integrity under normal operation as well as anticipated operational occurrences (AOOs) and some postulated accident scenarios insofar as they maintain reactor safety [31]. The research discussed here on process parameter optimization and post-process heat treatment is therefore likely to be critical to any potential future role of UAM in the nuclear industry, so that the methods investigated can be developed to commercial scale to yield UAM-fabricated components that meet NRC standards. Additional research would be needed to test such materials in demanding reactor environments, but the results of work thus far suggest that the challenges associated with UAM can be solved.

2.2.1. Application to HFIR Control Elements

Fabrication of the HFIR CEs with UAM has been successfully demonstrated; indeed, the process has even been able to directly print curved CE plates, promising potential for significant savings in the cost, time, and labor required for CE fabrication [32]. In the UAM fabrication mode, an Al-6061 matrix is built up in layers. Holes are then milled in the Al build, and small discrete cylindrical compacts of neutron absorber powders blended with Al, illustrated in Figure 2.4, are directly embedded. The compacts nominally achieve 92 percent theoretical bulk density, with an absorber concentration of 60 percent by volume. The Al matrix buildup continues around the compacts, and additional compacts are embedded when the proper Al thickness is reached. The Al deposition then continues around the new compacts until the CEs are completed [18].




Figure 2.4. Discrete Absorber Compact

This entire process takes place in the solid phase at room temperature. No hot-rolling is required, and therefore the absorber compacts retain their discrete arrangement as shown in Figure 1.1 in the previous chapter. Though this fabrication demonstration was successful, further testing on the additively manufactured CEs is necessary to verify their mechanical and material integrity in the HFIR core environment, while the present research seeks to affirm the neutronic adequacy of the new CE design. Additionally, the UAM fabrication mode introduces new design variables which were not applicable to the original CE design. These new variables include the absorber compact size, compact array spacing, and the absorber concentration in the compacts, whereas in the original design only the absorber concentration could be varied. Each of these variables can impact the CE performance, providing ample opportunity for optimizing the design.

CHAPTER 3. CORE PHYSICS ANALYSIS

This chapter discusses the development of computational models used for neutronic evaluation of the additively manufactured CEs. These models are used to investigate the impact of the new CE design on the HFIR core physics, providing physical insight into the CE neutronic behavior as well as a preliminary feasibility assessment.

3.1. Methodology

The neutronics analysis is carried out using the Monte Carlo N-Particle code, MCNP [33]. MCNP is a well-known stochastic transport code widely used in the nuclear field and beyond. MCNP operates by explicitly simulating particle transport through a user-defined model of a physical system, probabilistically evaluating particle interactions with the medium using pseudo-random number generation and continuous-energy cross section data. The recorded particle track histories from the model can then be used to infer the average behavior of particles in the physical system, providing a statistical estimate of physical quantities such as particle flux, energy deposition, and neutron multiplication [33]. Monte Carlo transport carries the advantage of handling highly heterogeneous systems with high-fidelity 3D geometry, and readily lends itself to parallel computation. However, the primary challenge of Monte Carlo methods arises from the stochastic nature of the calculations; in most cases, detailed system models require very large numbers of particle history simulations to smooth out statistical noise and obtain reliable results, yielding long computation times even in parallel simulations. Simulation efficiency can be improved by employing appropriate variance reduction techniques to bias the source sampling or particle track directions while conserving the “weight” of the particles – that is, the magnitude of the contribution of a particle to a calculated quantity. Considering such variance reduction techniques as well as the broad availability of high-performance computing capabilities at

research institutions worldwide, Monte Carlo methods remain an effective and favored means of transport analysis for a range of fields spanning power and research reactors, remote radiation detection, criticality safety, medical physics, and even astrophysics.

3.1.1. Model Requirements

The basis for the neutronics analysis is the existing MCNP model of the HFIR core at beginning of cycle (BOC) with the experimental loading of cycle 400 [19]. The documented calculations using this model were carried out using MCNP5; thus, while a more recent MCNP version (MCNP6) is available, MCNP5 is used for consistency with the documentation [19].

The CE representation in the base model must first be converted to an equivalent discrete version to represent the additively manufactured CE design. As the different designs are expected to yield unequal isotopic depletion rates, fresh (unirradiated) CE compositions are assumed at BOC 400 for consistent comparison. The equivalent discrete CE model is constructed to preserve the absorber mass content of the base model in each CE absorber zone, thereby isolating the effects of absorber discretization in the following analysis.

The primary modeling challenge of the discrete CEs lies in capturing the geometry of the array of absorber compacts with sufficient fidelity. When the discrete CEs are fully formed, the absorber compacts are arranged in a cylindrically symmetric lattice, with regular spacing in the radial, azimuthal, and axial directions. Two radial layers of compacts are present, each offset from the other azimuthally and axially to maximize absorber coverage of the CEs. In each radial layer, the compacts take on a triangular lattice structure, such that each compact is equidistant from its six nearest neighbors (achieving the maximum possible packing). A rough sketch of this arrangement is shown in Figure 3.1, which illustrates a section of the compact lattice viewed from the side as if the CEs were unrolled into flat sheets; the red compacts form the inner radial

layer, and the blue compacts form the outer radial layer. This pattern is repeated throughout the CE absorbing regions.




Figure 3.1. Planar View of Absorber Compact Array

Tens of thousands of these small compacts, each 5.59 mm in diameter and 1.71 mm in thickness, are present with this cylindrically symmetric arrangement in the CEs. The MCNP geometry specification allows for shorthand input to simplify the definition of rectangular and hexagonal lattices, saving computer memory access and storage requirements. However, no such simplification is available for cylindrically symmetric arrays. Repeated structures may also be defined using transformations of existing cells, but this method is impractical for this application due to the sheer number of compacts present exceeding the limits of MCNP5's numerical labeling scheme of geometric units [33]. Methods for representing other unusual geometries with regularly repeated small units, such as the random distribution of spherical fuel units in a pebble bed reactor core [34, 35], have been developed and successfully implemented, but they are not applicable to this cylindrical structure. Therefore, to retain fidelity, the detailed geometry of the discrete CE models must be defined explicitly.

The detailed heterogeneity defined in this way yields a model that is highly computationally taxing, requiring large quantities of memory and frequent memory access as

simulated particles enter and leave numerous unique geometric units; thus the transport simulation is significantly slowed in the CE regions with no simplifications. In light of this, it is worthwhile to consider an approximate model – one which sacrifices geometric fidelity for improvements in computational efficiency [36]. Such a model could be useful for the design evaluation if it succeeds in capturing the same discrete CE physics phenomena as the explicit model without sacrificing detail which may impact performance calculations.

3.1.2. *Explicit Model*

The high-fidelity explicit model [37] captures the exact geometry of the CEs by first defining the position and orientation of each compact in the top rows of each absorbing zone, and then enclosing these in a lattice cell which is repeated axially to fill the rest of the absorbing regions. Figure 3.2 gives a cross section of this CE representation showing the radial and azimuthal compact arrangement, and Figure 3.3 gives a side view showing their axial arrangement. In this model, the absorber mass content depends on the azimuthal spacing between the compacts; the geometric settings needed to reach the desired absorber content are shown in Table 3.1. The precision in absorber mass achieved in this model is considered acceptable and remains within the mass tolerance limits for actual CE production [17].




Figure 3.2. Cross Section of Explicit Model CE Representation




Figure 3.3. Side View of Explicit Model CE Representation

Table 3.1. Explicit Model Absorber Mass Matching and Geometry

Absorber Zone	Homogeneous Model Absorber Mass (g)	Azimuthal Compact Pitch (degrees between compact centers)	Minimum Spacing between Compact Edges (mm)	Explicit Model Absorber Mass (g)	Total Number of Compacts	Compact Packing Fraction (Volumetric)
ICE Gray	4937	1.634	0.7227	4946	9391	0.5002
ICE Black	8853	1.585	0.5332	8867	43853	0.4872
OCE Gray	4970	1.600	0.8597	4969	9376	0.4807
OCE Black	8905	1.554	0.6765	8900	43275	0.4672

Though the extent of the exact compact array geometry definition is minimized by repeating the most basic unit (the first two rows of compacts) in each absorber zone, the additional required detail is still substantial; about 45,000 new cells and 9,000 new surfaces must be defined. This significantly increases the memory used by the model as well as the frequency of memory access during transport simulations in the CE regions, considerably slowing these calculations. Calculations with the explicit model take about six times the amount of wall time required for the homogeneous model under identical calculation parameters on identical machines.

3.1.3. Approximate Model

The computationally efficient approximate model is developed for the early phase of analysis to identify the neutronics effects of absorber discretization at a high level. The geometry of the absorber compact array is simplified by fitting a Cartesian lattice of compacts into the cylindrical CE regions. In each CE quadrant, the compacts are rotated so that they face approximately normally outward. A portion of this representation is shown in Figure 3.4. Though not visible in Figure 3.4, the compact lattice takes on a body-centered cubic (BCC) structure to

maximize packing. The absorber mass content depends on the lattice pitch, which is parametrically adjusted in each absorbing zone to achieve the absorber mass of the original homogeneous CE model; due to the irregular compact loading resulting from truncating the Cartesian lattice with the cylindrical CE surfaces, the absorber mass is calculated using ray tracing with MCNP [33]. The lattice pitch and corresponding absorber mass in each absorbing zone is given in Table 3.2 for comparison with the homogeneous CE model. The precision obtained is acceptable, and is again within the absorber mass tolerance limits set for the actual production of the HFIR CEs [17].




Figure 3.4. Approximate Model CE Representation

Table 3.2. Approximate Model Absorber Mass Matching and Geometry

Absorber Zone	Homogeneous Model Absorber Mass (g)	Approximate Model Lattice Pitch (cm)	Approximate Model Absorber Mass (g)	Compact Packing Fraction (Volumetric)
ICE Gray	4937	0.8824	4934	0.4990
ICE Black	8853	0.8596	8853	0.4864
OCE Gray	4970	0.8990	4969	0.4807
OCE Black	8905	0.8753	8904	0.4674

The geometry specification for this model is relatively simple, requiring about 50 new cell definitions and 100 new surface definitions. As the Cartesian lattice construct minimizes the computational memory requirements [33], calculations with the approximate model take essentially the same amount of wall time to run as those with the homogeneous CEs under identical calculation parameters on identical machines.

3.2. Physics Results

The HFIR BOC 400 models detailed above are applied to compare the HFIR core neutronics behavior between CE designs [36, 37]. Table 3.3 gives the neutron multiplication eigenvalue yielded by each model. Neutron multiplication is slightly higher in both discrete cases, suggesting that the discrete design loses some absorption effectiveness relative to the homogeneous CEs. This is in agreement with physical intuition, as absorber discretization yields greater self-shielding relative to the homogeneous case, thereby effectively reducing the amount of absorber available to interact with incident thermal neutrons.

Table 3.3. Neutron Multiplication of Homogeneous and Discrete CE Models

Model	k_{eff}	Deviation from Homogeneous (pcm)
Homogeneous	1.00829 ± 0.00007	-
Approximate	1.01075 ± 0.00007	244.0 ± 9.9
Explicit	1.00960 ± 0.00007	129.9 ± 9.9

For additional insight into the discrete CE neutronics, it is instructive to observe the impact on the thermal flux distribution in the core. The thermal neutron energy range is defined with an upper limit of 0.625 eV (for consistency with established model documentation [19]), and thermal flux is tallied in each model in the r-z plane integrated about the azimuth. It is difficult to glean information from plotting the thermal flux in each case alone, as differences between cases are not immediately visually discernable; these plots are given in Appendix A for the sake of completeness. Instead, the pointwise relative difference is taken between the discrete and homogeneous results:

$$\delta(\mathbf{r}) = \frac{F_{\text{het}}(\mathbf{r}) - F_{\text{hom}}(\mathbf{r})}{F_{\text{hom}}(\mathbf{r})} \quad (3.1)$$

where δ is the relative difference, \mathbf{r} is the position vector, F is the physical quantity of interest (thermal flux in this case), and the subscripts “het” and “hom” refer to the discrete and homogeneous cases, respectively. Error is propagated appropriately [38]. Plotting the pointwise relative difference gives immediate clarity as to the differences between designs. Figure 3.5 gives the plots of the relative difference in thermal flux between the discrete and homogeneous CE cases, as well as the associated uncertainty.

Differences in thermal flux between CE designs are generally very limited except in the black CE regions, where the thermal flux in the discrete cases is approximately 20 percent greater than in the homogeneous case. This provides additional evidence for the loss in absorption effectiveness of the discrete CEs, which are now seen to house excess free thermal neutrons relative to the homogeneous CEs. This effect is attributed to a combination of the self-shielding arising from absorber discretization, reducing the effective absorption cross section, and the inherent geometry of the discrete CEs, which contain paths of absorber-free aluminum through which neutrons can travel with significantly reduced absorption probability. In contrast, the homogeneous CEs contain uniformly distributed absorbers in lesser concentrations than in the discrete absorber compacts, mitigating the impact of self-shielding and maintaining a large absorption probability throughout the absorbing region. Similar increases in thermal flux are not seen in the gray regions, which have significantly lower absorption cross sections than the black regions and exhibit only a minor impact on incident thermal flux. Additionally, the effect is more pronounced in the approximate model than in the explicit model due to inaccurate compact orientation amplifying the geometric effect.

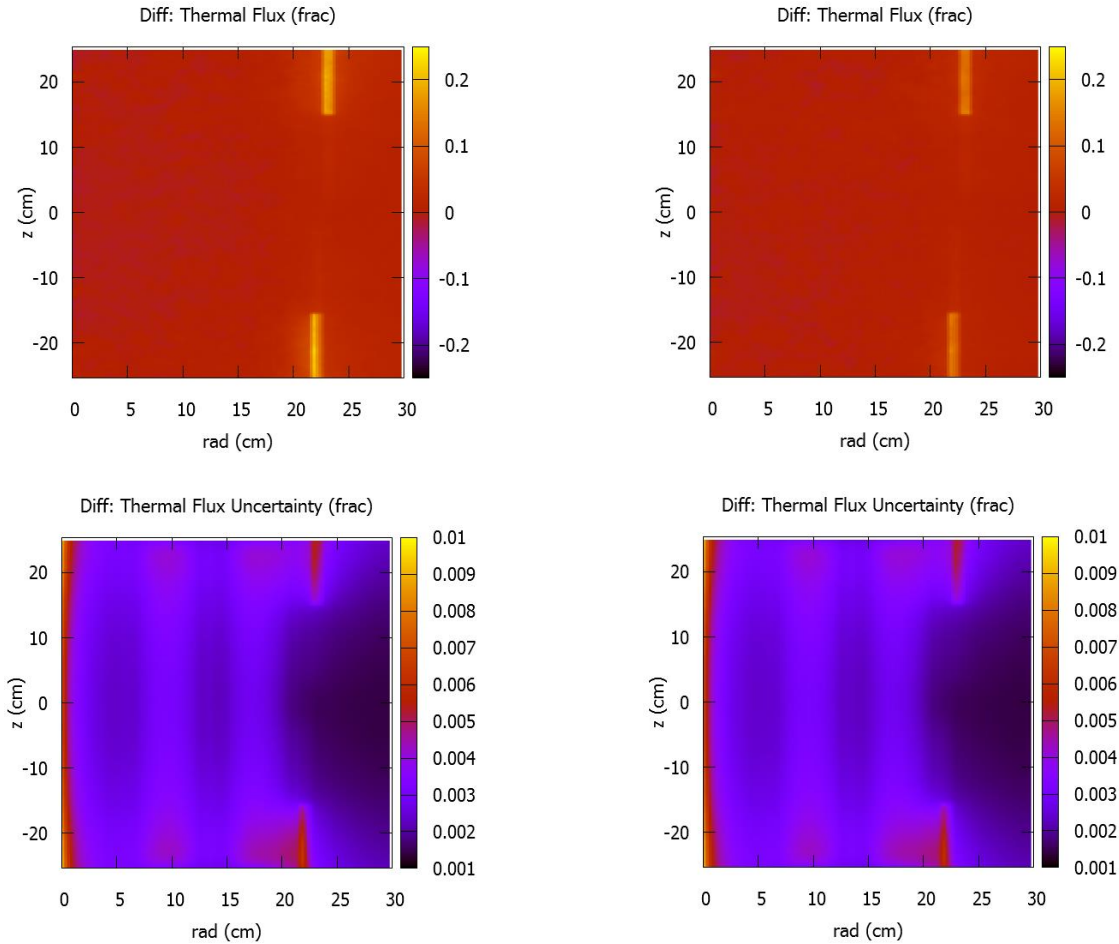


Figure 3.5. Relative Difference in Thermal Flux: Approximate vs. Homogeneous Model (top left) and Uncertainty (bottom left); Explicit vs. Homogeneous Model (top right) and Uncertainty (bottom right)

Another key reactor physics quantity is the fission rate density distribution in the fuel, which indicates power peaking and drives thermal-hydraulic design parameters. Again, differences in plots of the fission rate density alone are difficult to directly observe visually, and therefore the pointwise relative difference between cases is taken and plotted in Figure 3.6 (the fission rate density plots are included in Appendix A). Figure 3.6 reveals the upshot of the previously observed increase in thermal flux in the black regions of the discrete CEs, which is a corresponding increase in the fission rate at the axial extremes of the outer edge of the OFE nearest the black CE regions. The excess thermal neutrons not absorbed by the discrete CEs are

available to induce additional fission reactions in the fuel nearest the black CE regions, yielding this slight shift in the fission rate distribution radially and axially outward and also accounting for the increase in neutron multiplication observed in the discrete CE cases. The effect is again more prominent in the approximate model than in the explicit model due to the non-exact compact orientation allowing the streaming of greater thermal flux. It should be noted that a nonzero fission rate occurs in the beryllium reflector due to trace uranium impurities modeled in this region.

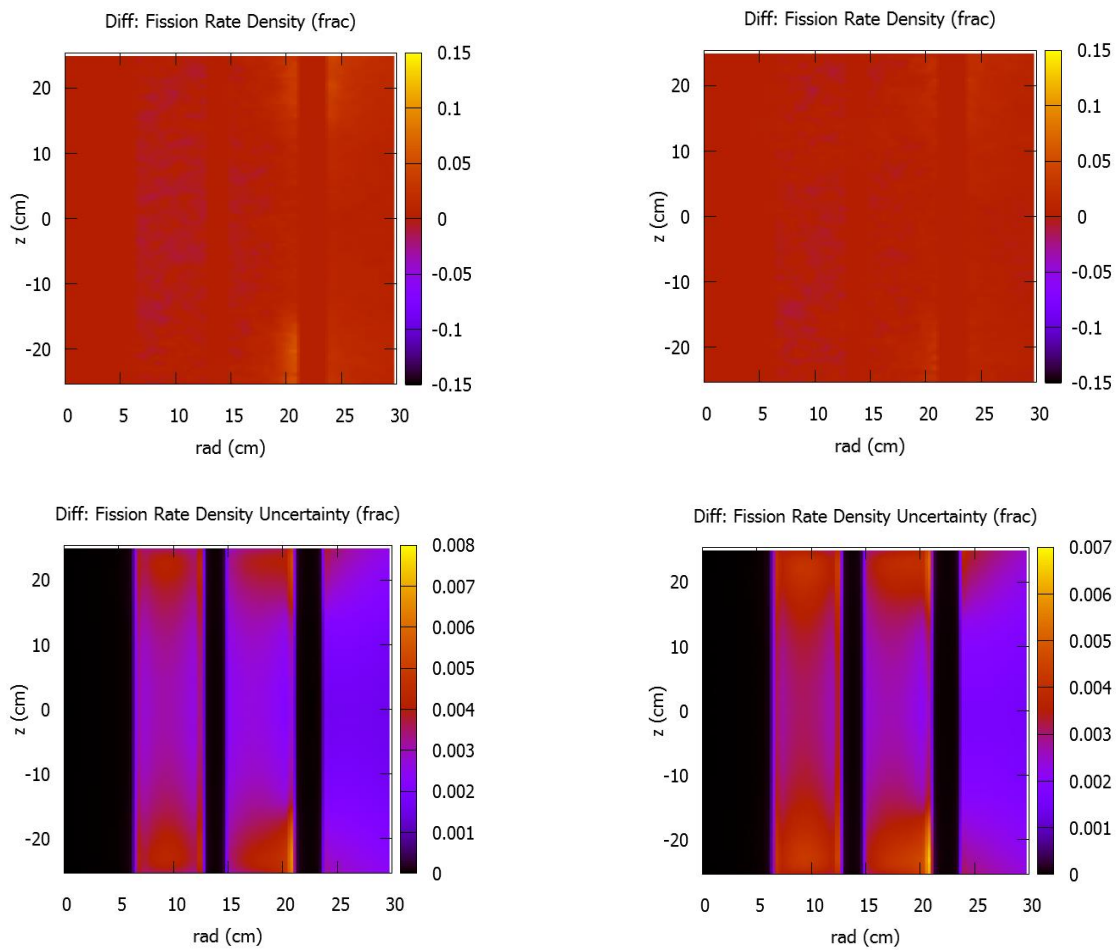


Figure 3.6. Relative Difference in Fission Rate Density: Approximate vs. Homogeneous Model (top left) and Uncertainty (bottom left); Explicit vs. Homogeneous Model (top right) and Uncertainty (bottom right)

As the preceding results reveal some macroscopic effects of the reduced absorption effectiveness of the discrete CEs, it is now desired to observe the absorption behavior of the CEs in greater detail. The absorption rate density is tallied along the axial extent of the CEs; again, as it is difficult to identify differences between cases from plots of the absorption rate density alone, the pointwise relative difference in absorption rate density between cases is taken and plotted in Figure 3.7 (absorption rate density plots are given in Appendix A). Though the differences are noisy due to absorber discretization yielding sections of alternating absorber thickness (with statistics playing a lesser role that is more significant in regions far from the core midplane), the approximate model exhibits little change in absorption on average. However, the discrete model exhibits some curious behavior; though little change on average is observed in the ICE, a significant increase in absorption by the OCE is observed on the order of 5 percent in the gray region and 10-15 percent in the black region. That this increase in absorption rate occurs in tandem with the previously observed increase in thermal flux initially seems counterintuitive; it would normally be expected that an increase in thermal flux would arise from a decrease in absorption rate. Further analysis is needed to explain this phenomenon.

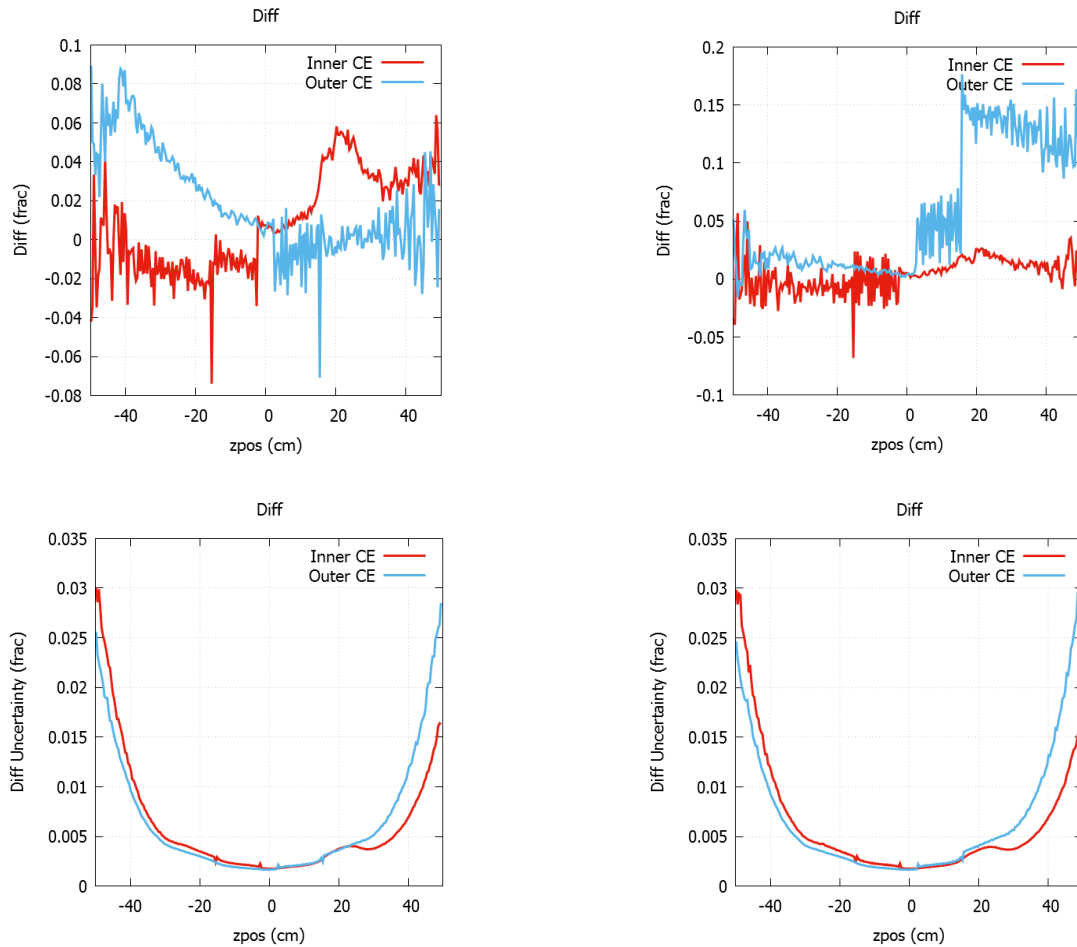


Figure 3.7. Relative Difference in Absorption Rate Density: Approximate vs. Homogeneous Model (top left) and Uncertainty (bottom left); Explicit vs. Homogeneous Model (top right) and Uncertainty (bottom right)

Despite the odd absorption behavior in the explicit model, self-shielding is expected to play a key role in the discrete CE neutronics in all absorbing regions. To assess the extent of its effect, thermal flux is tallied separately in the absorber compacts and the surrounding aluminum matrix in the discrete CE models so as to measure the resultant flux depression. Each of the gray and black regions in the computational models is divided into five axial sections of equal length to capture flux gradients [19]. The thermal flux results in the compacts and the aluminum in each of these axial segments are given in Tables 3.4 and 3.5 for the approximate and explicit models, respectively. The axial boundaries of each region are specified, corresponding to the startup CE

positions. Additionally, a self-shielding factor is calculated as the ratio of the flux in the compacts and the flux in the aluminum in each region [39, 40]:

$$f = \frac{\varphi_c}{\varphi_0} \quad (3.2)$$

where f is the self-shielding ratio, φ is the thermal flux, and the “c” and “0” subscripts, respectively, refer to the compacts and the aluminum.

Table 3.4. Thermal Flux Depression in Absorber Compacts: Approximate Model

CE	Region	Upper Axial Bound (cm)	Lower Axial Bound (cm)	Thermal Flux in Compacts (10^{13} n/cm ² s)	Thermal Flux in Aluminum Matrix (10^{13} n/cm ² s)	Self-Shielding Ratio f
OCE	Black	71.12	59.94	0.00143 ± 0.000048	0.00816 ± 0.000314	0.175 ± 0.0089
OCE	Black	59.94	48.77	0.0104 ± 0.00130	0.0607 ± 0.00086	0.171 ± 0.0032
OCE	Black	48.77	37.59	0.0883 ± 0.00039	0.513 ± 0.0025	0.172 ± 0.0011
OCE	Black	37.59	26.42	0.501 ± 0.0010	2.86 ± 0.006	0.175 ± 0.0005
OCE	Black	26.42	15.24	0.997 ± 0.0015	5.54 ± 0.009	0.180 ± 0.0004
OCE	Gray	15.24	12.70	30.0 ± 0.05	31.3 ± 0.06	0.961 ± 0.0025
OCE	Gray	12.70	10.16	38.6 ± 0.06	40.6 ± 0.07	0.951 ± 0.0022
OCE	Gray	10.16	7.62	43.8 ± 0.07	45.7 ± 0.07	0.958 ± 0.0021
OCE	Gray	7.62	5.08	47.9 ± 0.07	50.3 ± 0.08	0.951 ± 0.0020
OCE	Gray	5.08	2.54	54.9 ± 0.08	57.4 ± 0.08	0.957 ± 0.0019
ICE	Gray	-2.54	-5.08	49.4 ± 0.07	52.2 ± 0.08	0.947 ± 0.0020
ICE	Gray	-5.08	-7.62	43.6 ± 0.07	45.8 ± 0.07	0.952 ± 0.0020
ICE	Gray	-7.62	-10.16	39.8 ± 0.06	41.9 ± 0.07	0.950 ± 0.0021
ICE	Gray	-10.16	-12.70	35.3 ± 0.06	37.1 ± 0.06	0.952 ± 0.0023
ICE	Gray	-12.70	-15.24	27.5 ± 0.05	28.9 ± 0.05	0.952 ± 0.0026
ICE	Black	-15.24	-26.42	0.931 ± 0.0014	4.95 ± 0.009	0.188 ± 0.0004
ICE	Black	-26.42	-37.59	0.480 ± 0.0010	2.66 ± 0.006	0.181 ± 0.0006
ICE	Black	-37.59	-48.77	0.0740 ± 0.00035	0.420 ± 0.0023	0.176 ± 0.0013
ICE	Black	-48.77	-59.94	0.00847 ± 0.000121	0.0488 ± 0.00078	0.174 ± 0.0037
ICE	Black	-59.94	-71.12	0.00137 ± 0.000054	0.00761 ± 0.000313	0.179 ± 0.0099

Table 3.5. Thermal Flux Depression in Absorber Compacts: Explicit Model

CE	Region	Upper Axial Bound (cm)	Lower Axial Bound (cm)	Thermal Flux in Compacts (10^{13} n/cm ² s)	Thermal Flux in Aluminum Matrix (10^{13} n/cm ² s)	Self-Shielding Ratio f
OCE	Black	71.12	59.94	0.00138 ± 0.000044	0.00859 ± 0.000275	0.161 ± 0.0073
OCE	Black	59.94	48.77	0.00960 ± 0.000122	0.0573 ± 0.00068	0.168 ± 0.0029
OCE	Black	48.77	37.59	0.0832 ± 0.00037	0.483 ± 0.0020	0.172 ± 0.0010
OCE	Black	37.59	26.42	0.474 ± 0.0009	2.73 ± 0.005	0.173 ± 0.0005
OCE	Black	26.42	15.24	0.949 ± 0.0014	5.34 ± 0.007	0.178 ± 0.0004
OCE	Gray	15.24	12.70	30.1 ± 0.05	30.7 ± 0.06	0.982 ± 0.0025
OCE	Gray	12.70	10.16	38.9 ± 0.06	39.9 ± 0.06	0.977 ± 0.0022
OCE	Gray	10.16	7.62	43.9 ± 0.07	44.8 ± 0.07	0.980 ± 0.0021
OCE	Gray	7.62	5.08	48.4 ± 0.07	49.4 ± 0.07	0.980 ± 0.0021
OCE	Gray	5.08	2.54	55.3 ± 0.08	56.5 ± 0.08	0.979 ± 0.0019
ICE	Gray	-2.54	-5.08	49.8 ± 0.07	51.9 ± 0.08	0.960 ± 0.0020
ICE	Gray	-5.08	-7.62	43.5 ± 0.07	45.5 ± 0.07	0.956 ± 0.0020
ICE	Gray	-7.62	-10.16	39.9 ± 0.06	41.3 ± 0.07	0.964 ± 0.0022
ICE	Gray	-10.16	-12.70	35.4 ± 0.06	36.7 ± 0.06	0.966 ± 0.0023
ICE	Gray	-12.70	-15.24	27.4 ± 0.05	28.4 ± 0.05	0.964 ± 0.0026
ICE	Black	-15.24	-26.42	0.866 ± 0.0013	4.98 ± 0.007	0.174 ± 0.0004
ICE	Black	-26.42	-37.59	0.446 ± 0.0009	2.63 ± 0.005	0.169 ± 0.0005
ICE	Black	-37.59	-48.77	0.0685 ± 0.00033	0.416 ± 0.0019	0.165 ± 0.0011
ICE	Black	-48.77	-59.94	0.00788 ± 0.000113	0.0462 ± 0.00062	0.170 ± 0.0034
ICE	Black	-59.94	-71.12	0.00124 ± 0.000044	0.00754 ± 0.000248	0.164 ± 0.0080

As expected, self-shielding is significantly greater (lower self-shielding ratio) in the black regions than in the gray regions; the black compacts exhibit a significantly greater absorption cross section than do the gray compacts, yielding a correspondingly shorter mean free path for thermal neutrons and therefore a greater flux depression. Self-shielding in the black regions is greater in the explicit model than in the approximate model, as the improper compact orientation in the approximate model allows neutrons to pass through the compacts along short chord lengths with less likelihood of absorption, thereby bolstering the tallied flux in the compacts in this case. Correcting the compact orientation in the explicit model increases the distance that neutrons must travel through the absorber material, yielding a greater chance of absorption and therefore a greater flux depression. In the gray regions, the flux in the aluminum is generally greater in the approximate model than in the explicit model, while the flux in the compacts is comparable between models; thus the calculated flux depressions are greater in the approximate

model than in the explicit model. This is a geometric effect, as a compact arrangement in the approximate model poses wider paths of aluminum through which neutrons travel, yielding a larger flux tally.

To verify these self-shielding results, the methodology of [41] is adopted to calculate the expected self-shielding based on material properties:

$$G = \frac{1 - e^{-x_{op}}}{x_{op}} \quad (3.3)$$

where G is the self-shielding factor and x_{op} is the material optical thickness, equal to the geometric thickness expressed in units of the mean free path. Using rough thermal cross sections from [42], the self-shielding factors are estimated as 0.963 in the gray compacts and 0.149 in the black compacts. These values agree very well with the results of Tables 3.4 and 3.5. Though Equation 3.3 was actually developed for application to monodirectional flux fields, which is not the case in HFIR, good agreement is obtained regardless, granting confidence in the ability of the computational models to accurately capture the physics phenomena of the discrete CEs.

As the effects of self-shielding are found to be consistent in each discrete CE model, further insight into the CE absorption behavior is sought by investigating the effective absorption cross sections of the absorbing CE regions. Thermal cross sections are calculated by tallying absorption rate density and flux in the CEs in the thermal group and taking their ratio:

$$\Sigma_{th} = \frac{a_{th}}{\varphi_{th}} \quad (3.4)$$

where Σ is the macroscopic absorption cross section, a is absorption rate density, φ is flux, and the subscript “th” refers to the thermal energy group. Error is again propagated appropriately [38]. Differences between cases are now more readily visible with plots of the cross sections

alone, but plots of the relative difference in the cross sections are given here anyway for maximum clarity in the results. These are shown in Figure 3.8, while the cross section plots themselves are given in Appendix A.

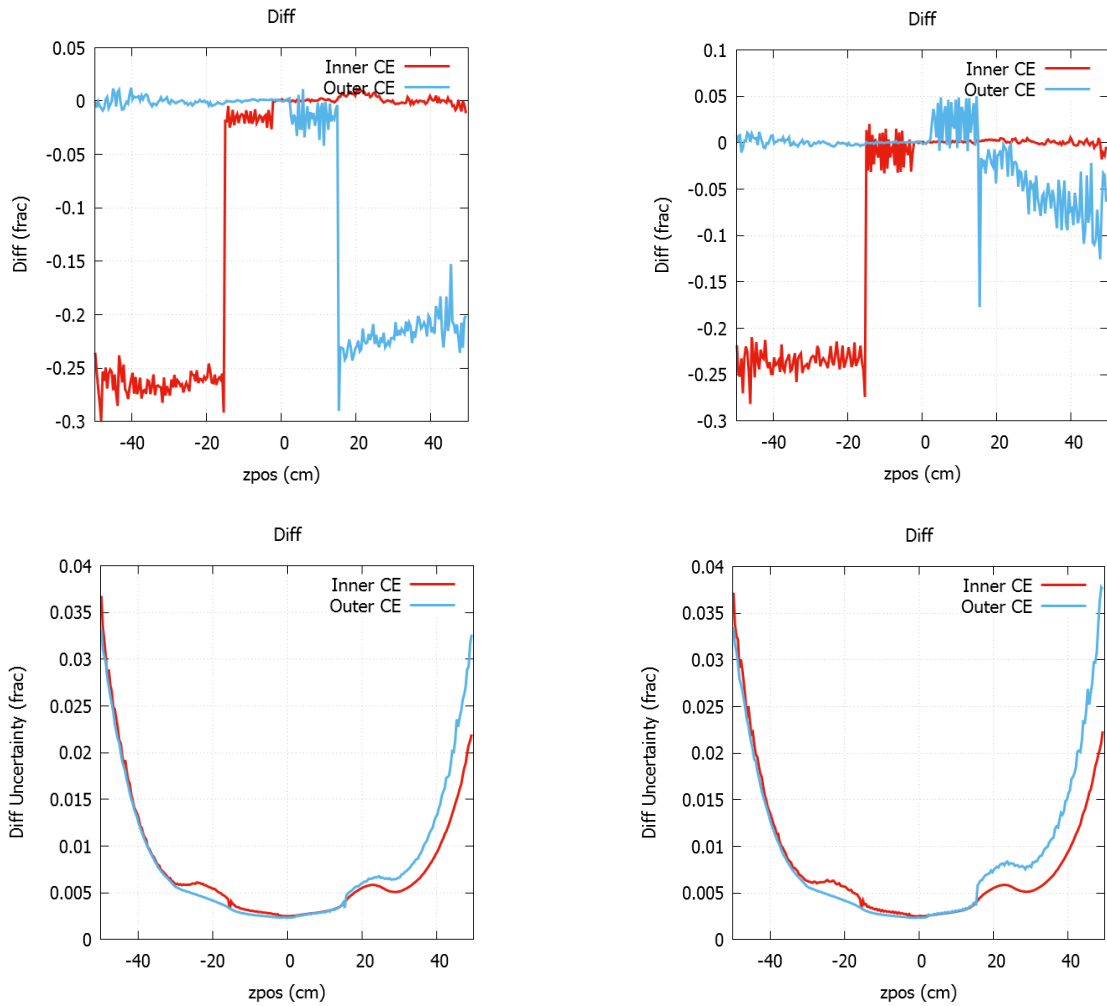


Figure 3.8. Relative Difference in Effective Thermal Absorption Cross Section: Approximate vs. Homogeneous Model (top left) and Uncertainty (bottom left); Explicit vs. Homogeneous Model (top right) and Uncertainty (bottom right)

Figure 3.8 clearly exhibits both self-shielding and geometric effects in each discrete CE model. The self-shielding arising from absorber discretization acts to reduce the effective absorption cross sections in all CE regions in the approximate model. The geometric impact is clear in the explicit model, where correcting the compact orientation bolsters the effective

absorption cross sections in all CE regions relative to the approximate model. However, the geometric correction has a greater impact in the OCE than in the ICE, serving to actually increase the effective absorption cross section in the OCE gray region while significantly mitigating its reduction in the black region. This is a spectral effect – the OCE experiences a softer neutron energy spectrum than does the ICE, thereby amplifying changes to the effective cross sections. The spectrum in each CE absorbing region is given in Figure 3.9, which indicates a greater thermal neutron population in the OCE absorbing regions than in the corresponding ICE regions (while the opposite is true at fast energies). It can therefore be concluded that the increase in absorption rate by the OCE in the explicit model arises from the interplay of these self-shielding, geometric, and spectral effects. Specifically, the modest absorption cross section in the gray regions yields little self-shielding impact, and ultimately the consolidation of absorbers in discretization coupled with the soft spectrum in the OCE yields an overall increase in absorption in the gray region; self-shielding is far more prominent in the black regions, but its impact is mitigated by the soft OCE spectrum and the geometry-derived increase in thermal neutron population. That this effect is captured by the explicit model and not the approximate model indicates that the accurate representation of the absorber compact array geometry is indeed significant to the discrete CE neutronics evaluation.

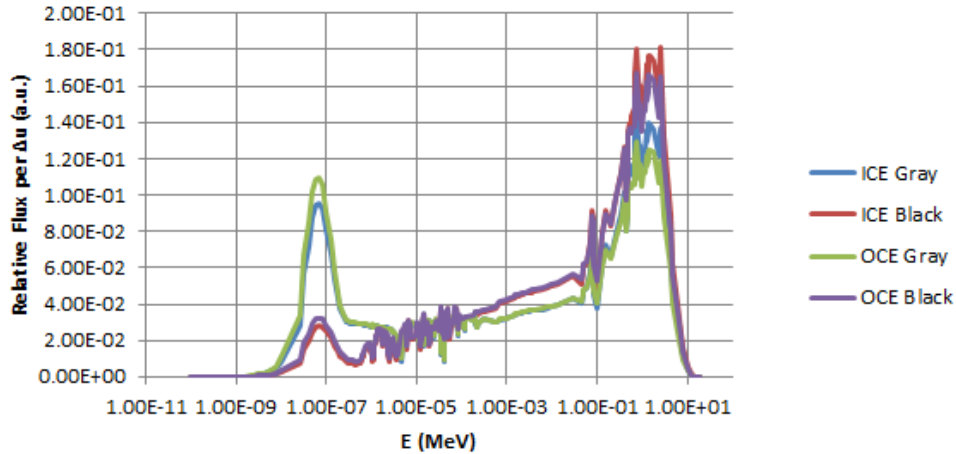


Figure 3.9. Neutron Energy Spectra in CE Absorbing Regions

3.3. Summary and Discussion

While the preceding analysis has identified some interesting neutronics phenomena which give rise to nontrivial differences in the absorption behavior of the CEs, the overall HFIR core characteristics are not significantly altered. The geometry and self-shielding effects arising from absorber discretization were found to yield an increase in free thermal flux in the black CE regions and their immediate vicinity. These excess thermal neutrons induce additional fission reactions at the axial extremes of the outer edge of the OFE, slightly increasing the neutron multiplication eigenvalue and shifting the core power distribution radially and axially outward. Though the relative gain in free thermal flux is nontrivial, it occurs in low-importance regions of the core, yielding minimal impact overall; indeed, the increase in neutron multiplication is on the order of 100 pcm, and the outward shift in the core power distribution does not raise any concern. Therefore, these initial results suggest that the additively manufactured CEs should perform reasonably similarly to the original CE design, with no significant loss in reactivity control capability or thermal margins.

On the microscopic scale, absorber discretization gives rise to the multifaceted intermingling of self-shielding, geometric, and spectral effects. The discrete CEs lose absorption

effectiveness due to self-shielding, but the absorption rate is bolstered by the increased thermal flux housed in the absorber-free aluminum between compacts. The explicit model reveals the spectral component, wherein the geometric boon to the absorption rate is significantly amplified in the OCE relative to the ICE due to a softer spectrum. This phenomenon may have significant implications for the discrete CE performance analyses to follow; thus, while the approximate model was able to identify the high-level macroscopic effects of absorber discretization with minimal computational effort, it must be jettisoned in favor of the explicit model in order to accurately capture all of the neutronics phenomena in the discrete CEs.

CHAPTER 4. PERFORMANCE EVALUATION

The previous chapter revealed no neutronics impediments to the use of additively manufactured CEs in HFIR based on computational models under BOC 400 conditions. There is no reason to believe that the neutronics conditions will drastically change over additional cycles, as the absorptive characteristics of the Ta and Eu transmutation products ensure consistent CE absorption strength, reactivity hold-down, and power shaping. However, these conditions alone do not necessarily guarantee long-term maintenance of the required standards of reactor safety. Thus, in order to verify the long-term safety capability of the additively manufactured CEs, it is necessary to explicitly quantify their lifetime performance. Material damage considerations limit the CE useful lifetime to 100 GWd of reactor operation [22], which corresponds to roughly 50 full-power HFIR cycles. This chapter therefore turns to depletion analysis of the CEs over this time frame for input to performance calculations.

4.1. Methodology

While many options for computational depletion analysis are available, it remains difficult to capture exactly the continuously evolving conditions in time-dependent problems. In a physical system under neutron irradiation, the flux spectrum and reaction cross sections are tightly coupled to isotopic compositions. Initial compositions determine the initial flux spectrum and the corresponding reaction cross sections. These cross sections and the flux magnitude drive the isotopic transmutation rates; the accumulation of new transmutation products in turn impacts the flux spectrum and cross sections, and so on in continuous fashion.

The general transmutation rate equation for a single isotope “i” in a system under neutron irradiation takes the following form:

$$\frac{dn_i}{dt} = \sum_j [\sigma_j(t)\varphi(t) + \lambda_j]n_j(t) - [\sigma_i(t)\varphi(t) + \lambda_i]n_i(t) \quad (4.1)$$

where n is the isotopic content on a total atomic or volumetric basis, t is time, σ is an energy-averaged reaction cross section, ϕ is total flux, λ is a decay constant, and the summation is taken over all isotopes “ j ” which transmute into isotope “ i ” via decay or any neutron interaction with cross section σ_j . The time dependence of isotopic content, spectral-averaged reaction cross sections, and total flux is explicitly shown. Some simplifying assumptions may allow cross sections and flux to be held constant, at least over a particular subinterval within the depletion time, yielding closed-form solutions in the Bateman equations [43, 44]. Indeed, computational depletion solvers generally assume that the analyst has supplied sufficiently fine temporal discretization of the depletion time such that fixed average flux and cross sections adequately capture depletion behavior in each time step, allowing direct application of closed-form solutions to Equation 4.1. Estimation of the average conditions in each time step may be improved by employing various predictor-corrector methods [44-47]; such methods use conditions at the beginning of a time step to estimate conditions at the middle or end of the time step (the predictor step), and then repeat the entire time step using updated conditions (the corrector step). The manner in which conditions are updated varies among methods. The general predictor-corrector computational flow is illustrated in Figure 4.1, and will be discussed in more detail later.

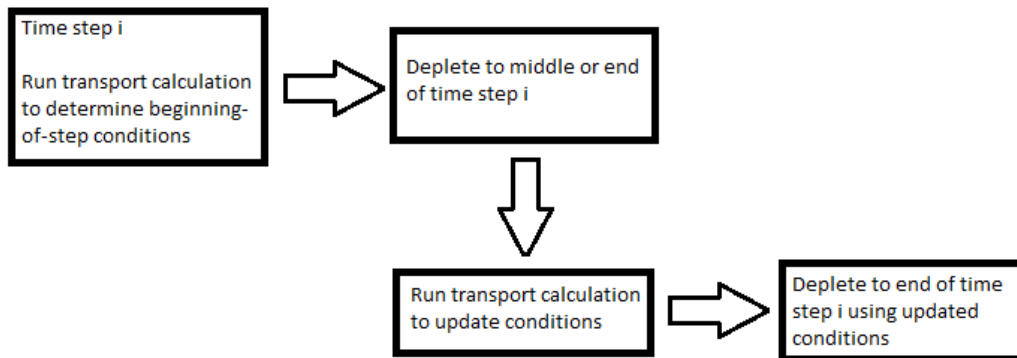


Figure 4.1. Predictor-Corrector Depletion Calculation Flow

4.1.1. Depletion Model

The depletion tool used for the present analysis is the VESTA code [48], a generic Monte Carlo depletion code. Though it is still undergoing development, VESTA is intended to interface with any Monte Carlo transport code and provide several options for depletion solvers for broad, flexible application. In the present application, VESTA runs transport simulations with MCNP to determine spectral conditions in the modeled system in an ultra-fine energy group structure and passes them into the ORIGEN-2.2 depletion and decay module from the SCALE code [44]. The isotopics results from the ORIGEN calculation are then used to update the transport model for the subsequent time step. This process repeats to the end of the depletion time. The use of conditions at the beginning of each time step for the depletion calculation comprises a predictor-only scheme, although the option to use a predictor-corrector scheme is available. VESTA is also capable of modeling system changes such as control rod motion during a cycle and fuel shuffling between cycles using various transformation specifications.

Little additional input is required for VESTA calculations beyond the original MCNP model and specification of the materials to be depleted [50]. In addition to the finely spatially divided fuel materials, separate absorber materials are defined in each axial division of the CE absorbing regions identified in the previous chapter to capture varying depletion rates arising from axial variation in the flux magnitude and spectrum. Time steps are defined for each day in the 25-day HFIR cycle, yielding 25 individual transport simulations with MCNP punctuated by ORIGEN depletion calculations. The reactor power in the simulation of the first day in the cycle is reduced to 56.5 MW to account for the gradual approach to full power (85 MW) during startup [19]. CE withdrawal is simulated using the day-by-day positions of the cycle 400 withdrawal schedule, shown graphically in Figure 4.2 [19].

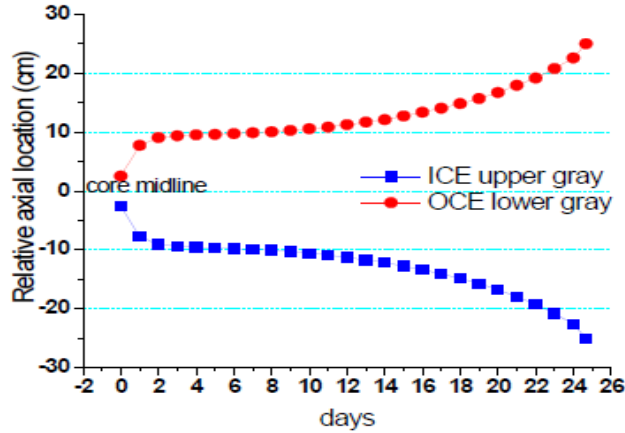


Figure 4.2. Cycle 400 Withdrawal Schedule [19]

It was determined in the previous chapter that the explicit model of the discrete CE design must be used to completely capture the relevant neutronics phenomena and yield reliable performance evaluations. As even a single transport calculation with the explicit model is highly computationally demanding in its own right, the depletion simulation of an entire 25-day cycle is multiply so. Moreover, CE depletion must be evaluated over the entire 50-cycle lifetime of the CEs. As such, explicit simulation of 50 depletion cycles is computationally infeasible; this is the case even with the homogeneous CE models. It is therefore necessary to construct depletion approximations to enable lifetime depletion evaluation with reasonable computational effort.

4.1.2. Long-Term Depletion Approximations

In evaluating lifetime CE depletion, it is desired to run as few explicit depletion simulations as possible so as to minimize computational burden. The approximations to be developed are therefore based on the assumption that a limited set of representative conditions (i.e. flux magnitude, spectrum, and cross sections) determined from explicit simulations can be applied to approximately represent many cycles [49]. Once the representative conditions are fixed, the costly repetition of transport calculations is eliminated, and only ORIGEN is needed to

assess long-term depletion. The key to the approximation is therefore the appropriate selection of representative conditions.

The conditions applied in the approximation scheme must adequately capture the rapid evolution of conditions in a single cycle, which arise in large part from CE withdrawal, as well as more gradual changes over the course of many cycles. Thus there are two levels to the depletion approximation: the selection of conditions which best represent a single cycle, and the selection of a cycle whose conditions best represent multiple cycles through the CE life. Additionally, either level may be temporally subdivided to apply different sets of conditions through different subintervals and, if desired, incorporate a predictor-corrector scheme [47]. At the single-cycle level, previous work has identified the conditions of day 15 of the cycle as those which best represent cycle-average conditions in terms of CE position and flux spectrum [49]. These are determined to high fidelity from an explicit full-cycle depletion simulation, which also yields the day-by-day cross sections and total flux over the cycle for each depletion material. Consolidating this information yields a reasonably accurate single-cycle depletion approximation; for each depletion material, the day 15 spectrum and cycle-averaged cross sections are passed to the COUPLE cross section processing code [44], and the day-by-day total flux is specified in ORIGEN. Use of the day 15 flux spectrum and cycle-averaged cross sections is the best available fixed representation of evolving conditions in the cycle, and thus varying the total flux in each time step in ORIGEN approximately captures changes in conditions over the cycle [49].

As construction of the single-cycle depletion approximation requires an explicit simulation of the same cycle, computational burden must be reduced by avoiding additional explicit simulations and applying these single-cycle conditions to many cycles. Carrying out the

explicit simulation with the existing base models developed in the previous chapter yields the representative cycle 400 conditions. Thus, the absolute minimum computational burden is realized by assuming that the cycle 400 conditions can be applied to adequately represent all 50 cycles in the CE life (cycles 400-449). This comprises an extremely coarse predictor-only scheme, and may be improved by employing one or more corrector steps. Thus, additional multi-cycle approximations are developed based on the predictor-corrector schemes detailed in [47]. Though these schemes were originally applied only to irradiation periods with no intermediate decay time, the selection of representative conditions to correct depletion steps in a single cycle is analogous to the selection of representative cycles within multiple cycles; these schemes are therefore equally applicable on the multi-cycle level as on the single-cycle level.

Applying the cycle 400 conditions to all 50 cycles in the CE life corresponds to an explicit Euler method in the nomenclature of [47]; this approximation is diagrammed in Figure 4.3. The information yielded from this approximation may then be used to apply a corrector step. An implicit Euler method [47] can be constructed by running an explicit simulation to determine the cycle 449 (day 15) conditions based on the BOC 449 isotopics predicted by the explicit Euler approximation, and then applying these cycle 449 conditions to all 50 cycles. This approximation comprises a single iteration of the implicit Euler method, and is diagrammed in Figure 4.4. Additional iterations of the implicit Euler method may be carried out by recalculating the cycle 449 conditions using the results of the previous iteration, applying these updated conditions to all 50 cycles, and then averaging the isotopics results of each iteration [47]. An implicit Euler scheme in two iterations is diagrammed in Figure 4.5; any number of additional iterations may be carried out, but for this study no more than two iterations are assessed to conserve computational effort.



Figure 4.3. Explicit Euler Method

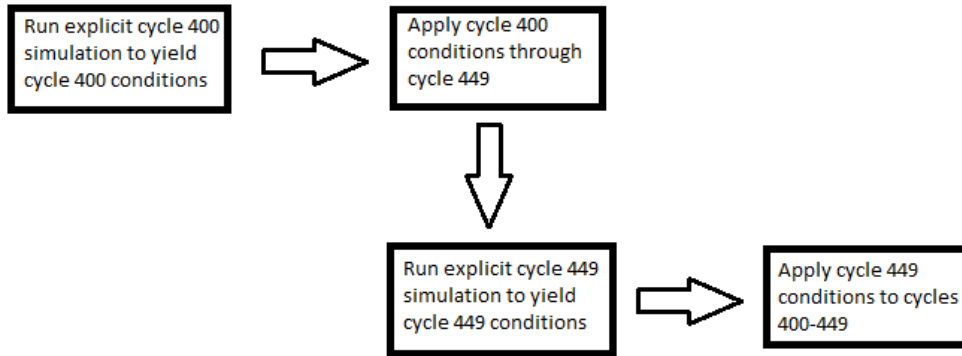


Figure 4.4. Implicit Euler Method in One Iteration

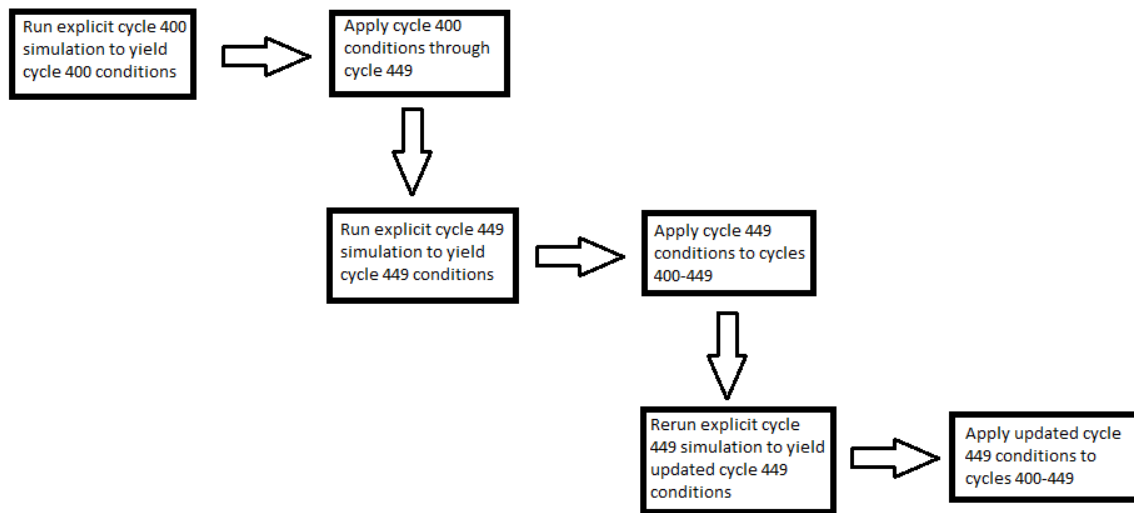


Figure 4.5. Implicit Euler Method in Two Iterations

Each of these preceding schemes relies on conditions at the extremes of the simulation; the use of intermediate conditions can be incorporated to construct an implicit midpoint scheme. In order to obtain reasonably accurate isotopics at intermediate points in the CE life, the 50 cycles are quartered. Specifically, the cycle 400 conditions are applied to cycles 400-412, at

which point an explicit simulation determines the cycle 412 conditions, which are then applied to cycles 400-424. A new explicit simulation determines the cycle 424 conditions, which are applied to cycles 424-437. A final explicit simulation is run to determine the cycle 437 conditions, which are then applied to cycles 424-449. This process is diagrammed in Figure 4.6; it is the most computationally intensive of the approximations considered here, requiring four explicit depletion simulations, and is therefore considered the best-estimate approximation.

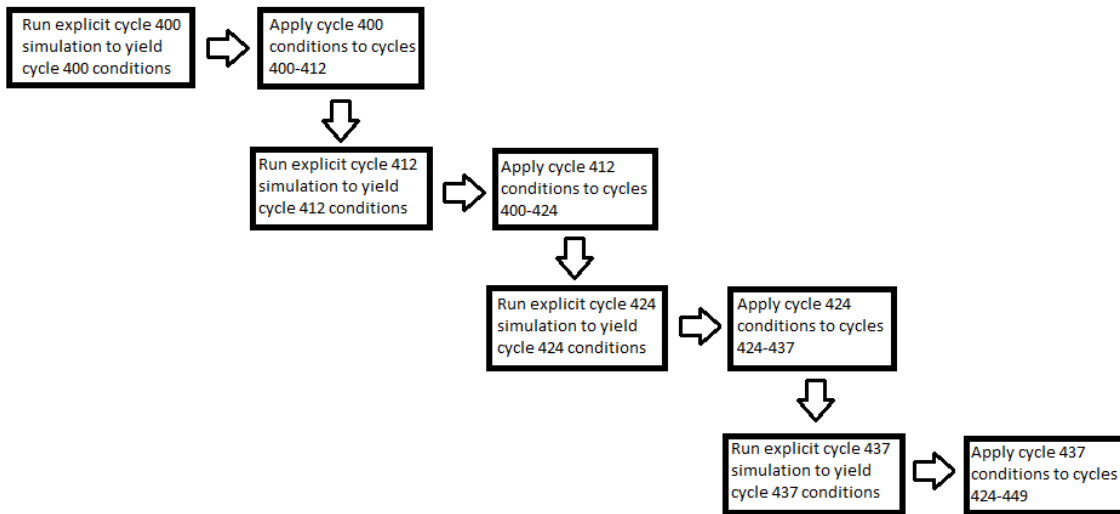


Figure 4.6. Implicit Midpoint Method

4.1.3. Performance Characterization

The depletion approximations developed above may be applied to yield lifetime CE isotopics estimates of varying reliability. The results may then be used to construct standalone transport models for use in performance evaluation. The additively manufactured CEs may be considered acceptable in terms of reactor safety if they meet two criteria: 1) they sustain sufficient reactivity worth over their lifetime, thereby maintaining shutdown margin and reactor power adjustment capability, and 2) they do not introduce excessive power peaking in the fuel, thereby maintaining thermal margins. Thus the key parameters to be assessed over the CE life are differential and integral reactivity worth and power peaking. Some deviations are acceptable

as long as they do not violate limits dictated by the HFIR safety documentation and/or can be compensated for by operational changes, e.g. adjusting the CE position. Differential and integral reactivity worth are calculated as:

$$DW(z) = \frac{100}{\beta} \frac{1}{k(z)} \frac{\partial k(z)}{\partial z} (\text{¢/cm}) \quad (4.2)$$

$$IW(z) = \frac{100}{\beta} \frac{k(z) - k_{in}}{k(z)k_{in}} (\text{¢}) \quad (4.3)$$

where DW and IW are respectively differential and integral worth, z is the nominal CE position at a given point in the reactor cycle, $k(z)$ is the neutron multiplication eigenvalue at this CE position, k_{in} is the neutron multiplication eigenvalue when the CEs are fully inserted, and the factor of $100/\beta$, where β is the delayed neutron fraction, converts the worth to units of cents (¢). Though not necessarily physically representative, the CEs are moved symmetrically in these calculations.

One additional parameter that deserves attention is photon heating in the CEs. Significant heat generation occurs in the CEs due to photon interactions with the high-Z absorber materials, contributing to material damage by driving corrosion of the CE clad surface. It is anticipated that absorber discretization will yield increased photon heating relative to the homogeneous CEs by locally concentrating the high-Z absorber material into regions of greater density and thus increasing the probability of photon interaction. Though no strict limits on photon heating are defined, it is an important design characteristic, and it should be shown that no excessive increases in photon heating are introduced by absorber discretization. Thermal-hydraulics calculations are beyond the scope of this study, but some engineering judgment shall assess whether additional heat generation can be reliably offset by the cooling systems currently in place.

4.2. Depletion and Performance Results

The four depletion approximations developed above (explicit scheme, implicit scheme in one and two iterations, and midpoint scheme) are now separately applied to estimate the lifetime CE isotopics for input to standalone transport models for use in evaluating the key CE performance parameters, reactivity worth and core power peaking. These evaluations are carried out at the beginning, middle, and end of cycles 400, 424, and 449 (BOC/MOC/EOC 400/424/449) to capture a representative cross section of the CE life. A 25-day decay period is modeled between cycles to account for refueling outages. Additionally, in practice, the OCE plates are removed after every seven cycles of use and allowed to cool for one year; during this time they are replaced with a cooled set of plates. This decay time is accounted for as well. The CE withdrawal schedule for cycles beyond cycle 400 cannot be definitively determined due to practical operational uncertainties, such as changes in experimental loading, unplanned shutdowns, etc.; thus it is assumed that the CE positions in later cycles are approximated by the cycle 400 withdrawal schedule.

4.2.1. Approximating Spectra

The spectra calculated in developing the multi-cycle predictor-corrector depletion approximations discussed in the previous section are plotted in Figure 4.7 below; these were taken in the homogeneous OCE black region closest to the core midplane for illustrative purposes. The spectra are labeled by the cycle in which the spectrum was calculated; in the two-iteration implicit scheme, the second simulation of cycle 449 is indicated with a prime mark (449'). Though not used in any of the approximations, the cycle 401 spectrum is included for comparison and was calculated from an explicit simulation.

The spectra exhibit a clear softening trend as CE use progresses due to the burnup of highly absorptive isotopes. The spectral shift over the course of the CE life is nontrivial, underscoring the difficulty of using the conditions of individual cycles for application to multiple cycles with sufficiently coarse temporal discretization as to maintain computational tractability. Along this line, it is also worth noting that the cycle 449 spectrum, calculated from a single iteration of the implicit scheme, is comparable to and even slightly harder than the cycle 437 spectrum calculated in the midpoint scheme. This suggests that the explicit scheme is inadequate in capturing long-term isotopic and spectral evolution with much fidelity; the true cycle 449 spectrum is expected to be much softer, likely closer to that calculated in the second iteration of the implicit scheme. This variation in spectral characteristics over the CE life is expected to yield markedly different isotopic transmutation rates in each approximation, which may subsequently impact the performance evaluations.




Figure 4.7. Approximating Spectra

4.2.2. Isotopics

The results of this section retain the focus on the homogeneous OCE black region closest to the core midplane. This zone is chosen because it exhibits the greatest absorption cross section with a relatively soft incident flux spectrum, and as such it exhibits the clearest changes in isotopics (the same trends are exhibited by the discrete CEs). To begin, it is of interest to examine the adequacy of the single-cycle approximation by comparing against the isotopics yielded by the explicit cycle 400 simulation. As it is again difficult to visually discern differences in isotopics from their plots alone, the relative difference between isotopics from the explicit simulation and the single-cycle approximation at each day of the cycle is calculated and plotted in Figure 4.8. The day 15 spectrum is relatively soft and thus yields accelerated transmutation rates early in the cycle, where the initial inventory of highly absorptive isotopes would in reality yield a hard spectrum. However, by the end of the cycle, as the true spectrum softens, the differences fall to within 50 percent. Differences between the most absorptive isotopes, Eu-151 and Eu-153, are quite limited throughout the cycle.

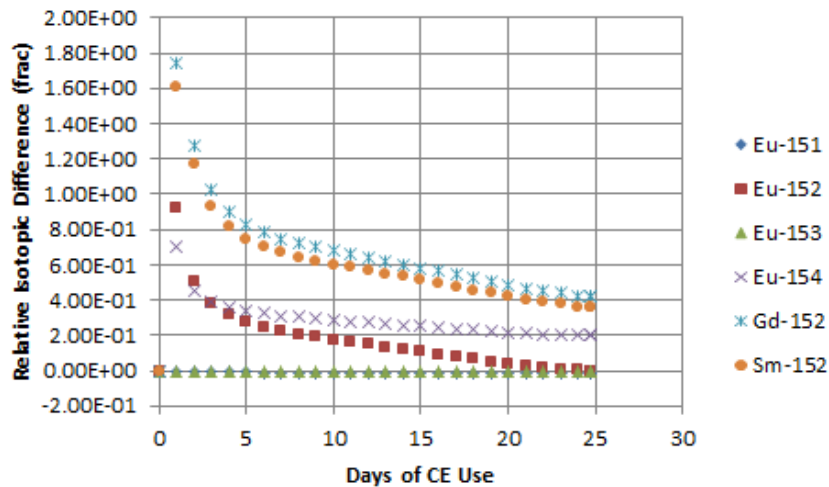


Figure 4.8. Relative Difference in Isotopics: Single-Cycle Approximation vs. Explicit Simulation

Deeming the single-cycle approximation adequate, attention is now turned to multi-cycle depletion. Figure 4.9 below shows the lifetime depletion and accumulation in mass of the most

important isotopes (in terms of absorption contribution) as determined from the best-estimate midpoint depletion scheme. The one-year OCE cooling time is visible every seven cycles. The sudden change of approximating spectra halfway through the depletion time is also visible, particularly for Gd-155 and Gd-157 (both of which have very large thermal absorption cross sections, on the order of 10^5 barns, and thus are highly sensitive to spectral changes). Differences in isotopic content between the approximation schemes are generally not immediately visually discernable, and thus it is of interest again to explicitly plot the relative difference in isotopic content between approximations.




Figure 4.9. Midpoint Scheme Lifetime Isotopics in Homogeneous Black OCE Region Closest to Core Midplane

Figure 4.10 plots the relative difference in isotopics between the explicit and midpoint schemes. Discrepancies are limited until the approximating spectrum of the midpoint scheme changes at cycle 424; at this point, the overestimation of the Eu-151 and Eu-152 content by the explicit scheme skyrockets beyond the range of the plot, which is limited in order to show the

trends in the remaining isotopes. The latter half of the midpoint scheme uses the cycle 437 spectrum, which is significantly softer than the cycle 400 spectrum. The hard cycle 400 spectrum inhibits rates of neutron capture, thereby yielding significantly slower transmutation rates and retaining greater quantities of the highly absorptive isotopes in the explicit scheme.




Figure 4.10. Relative Difference in Isotopics: Explicit Scheme vs. Midpoint Scheme

Figures 4.11 and 4.12 compare the midpoint scheme isotopics to the implicit schemes in one and two iterations, respectively. The soft cycle 449 spectra in both implicit schemes yield significant discrepancies in the first half of the depletion time, where the midpoint scheme uses the relatively hard cycle 412 spectrum. The implicit schemes therefore artificially accelerate depletion of highly absorptive isotopes and yield more rapid accumulation of transmutation products. These differences are ameliorated for most isotopes when the midpoint scheme switches to the softer cycle 437 spectrum, generally falling within a factor of two by the time cycle 449 is reached.




Figure 4.11. Relative Difference in Isotopics: Implicit Scheme in One Iteration vs. Midpoint Scheme




Figure 4.12. Relative Difference in Isotopics: Implicit Scheme in Two Iterations vs. Midpoint Scheme

4.2.3. Reactivity Worth

While the isotopics results presented above reveal some insight into the physical mechanisms of each depletion approximation, it is of greater practical relevance to examine their impact on the reactivity worth of the CEs. Tables 4.1-4.4 compare the differential worth estimations of each CE design under each approximation scheme, and Tables 4.5-4.8 give the same for integral worth.

Comparing across approximation schemes, it is apparent that those with softer approximating spectra yield lesser predictions of differential worth than those with harder spectra. Soft spectra as are present in the implicit schemes bolster neutron capture and transmutation rates, thereby yielding more rapid depletion of the most highly absorptive isotopes. CEs under soft approximating spectra therefore exhibit weaker reactivity hold-down (higher nominal neutron multiplication) and correspondingly lesser differential worth than those under harder approximating spectra. Conversely, CEs under hard approximating spectra, as in the explicit scheme, exhibit hampered absorption rates and slower transmutation, retaining more of their initial inventory of strongly absorbing isotopes and bolstering reactivity hold-down and differential worth. However, the opposite is true for integral worth; full insertion of CEs under softer approximating spectra yields a greater relative change in reactivity than CEs under harder approximating spectra, which exhibit greater initial reactivity at the nominal CE position. Thus, the softer approximating spectra of the implicit schemes yield greater predictions of integral worth than does the harder approximating spectrum of the explicit scheme. The greatest spread in worth predictions between approximation schemes is about 12 percent for differential worth and about 7 percent for integral worth; these discrepancies are much more limited than the significantly divergent isotopics results might suggest.

Table 4.1. Differential Worth Comparison under Explicit Depletion Approximation Scheme

Time in Cycle	Homogeneous CE Diff. Worth (€/cm)	Discrete CE Diff. Worth (€/cm)	Δ (%)
BOC 400	102.8 ± 1.6	103.9 ± 1.6	1.07 ± 2.21
MOC 400	79.2 ± 0.7	78.1 ± 0.7	-1.39 ± 1.77
EOC 400	22.6 ± 0.7	21.3 ± 0.7	-5.75 ± 4.26
BOC 424	106.7 ± 0.7	105.5 ± 0.7	-1.12 ± 0.92
MOC 424	76.6 ± 0.7	73.9 ± 0.7	-3.52 ± 1.27
EOC 424	19.7 ± 0.7	19.8 ± 0.7	0.51 ± 5.04
BOC 449	102.7 ± 0.7	99.7 ± 0.6	-2.92 ± 0.88
MOC 449	70.1 ± 0.6	68.3 ± 0.7	-2.57 ± 1.30
EOC 449	16.8 ± 0.6	16.9 ± 0.7	0.60 ± 5.50

Table 4.2. Differential Worth Comparison under Implicit Depletion Approximation Scheme in One Iteration

Time in Cycle	Homogeneous CE Diff. Worth (€/cm)	Discrete CE Diff. Worth (€/cm)	Δ (%)
BOC 400	102.8 ± 1.6	103.9 ± 1.6	1.07 ± 2.21
MOC 400	79.1 ± 0.7	77.7 ± 0.7	-1.77 ± 1.24
EOC 400	21.6 ± 0.7	20.7 ± 0.7	-4.02 ± 4.49
BOC 424	101.7 ± 0.7	101.9 ± 0.7	0.20 ± 0.97
MOC 424	74.7 ± 0.7	72.3 ± 0.7	-3.21 ± 1.30
EOC 424	18.7 ± 0.7	19.0 ± 0.7	1.60 ± 5.34
BOC 449	97.0 ± 0.7	95.9 ± 0.7	1.13 ± 1.01
MOC 449	66.2 ± 0.7	63.8 ± 0.6	-3.63 ± 1.36
EOC 449	14.9 ± 0.7	15.9 ± 0.7	6.71 ± 6.87

Table 4.3. Differential Worth Comparison under Implicit Depletion Approximation Scheme in Two Iterations

Time in Cycle	Homogeneous CE Diff. Worth (€/cm)	Discrete CE Diff. Worth (€/cm)	Δ (%)
BOC 400	102.8 ± 1.6	103.9 ± 1.6	1.07 ± 2.21
MOC 400	78.9 ± 0.7	79.9 ± 0.9	1.27 ± 1.45
EOC 400	20.9 ± 0.7	19.5 ± 0.9	-6.70 ± 5.32
BOC 424	102.2 ± 0.7	101.3 ± 1.1	-0.88 ± 1.27
MOC 424	72.3 ± 0.6	70.8 ± 1.0	-2.07 ± 1.60
EOC 424	18.6 ± 0.7	18.4 ± 0.9	-1.08 ± 6.11
BOC 449	97.5 ± 0.6	89.7 ± 0.9	-8.0 ± 1.08
MOC 449	62.5 ± 0.6	62.3 ± 1.0	-0.32 ± 1.86
EOC 449	16.0 ± 0.6	14.3 ± 0.9	-10.63 ± 6.55

Table 4.4. Differential Worth Comparison under Midpoint Depletion Approximation Scheme

Time in Cycle	Homogeneous CE Diff. Worth (€/cm)	Discrete CE Diff. Worth (€/cm)	Δ (%)
BOC 400	102.8 ± 1.6	103.9 ± 1.6	1.07 ± 2.21
MOC 400	78.9 ± 0.7	78.1 ± 0.8	-1.01 ± 1.34
EOC 400	20.9 ± 0.7	21.1 ± 0.8	0.96 ± 5.11
BOC 424	105.7 ± 0.7	103.5 ± 0.5	-2.08 ± 0.80
MOC 424	77.3 ± 0.8	75.8 ± 0.5	-1.94 ± 1.20
EOC 424	19.7 ± 0.7	19.8 ± 0.5	0.51 ± 4.38
BOC 449	99.4 ± 0.6	96.8 ± 0.8	-2.62 ± 1.00
MOC 449	68.2 ± 0.4	67.4 ± 0.6	-1.17 ± 1.05
EOC 449	16.7 ± 0.4	16.8 ± 0.5	0.60 ± 3.84

Table 4.5. Integral Worth Comparison under Explicit Depletion Approximation Scheme

Time in Cycle	Homogeneous CE Int. Worth (€)	Discrete CE Int. Worth (€)	Δ (%)
BOC 400	1407 \pm 3	1414 \pm 3	0.50 \pm 0.30
MOC 400	1988 \pm 2	1983 \pm 2	-0.25 \pm 0.14
EOC 400	2806 \pm 2	2796 \pm 2	-0.36 \pm 0.10
BOC 424	1242 \pm 2	1250 \pm 1	0.64 \pm 0.18
MOC 424	2028 \pm 2	2025 \pm 2	-0.15 \pm 0.14
EOC 424	2822 \pm 2	2805 \pm 2	-0.60 \pm 0.10
BOC 449	1357 \pm 1	1363 \pm 1	0.44 \pm 0.10
MOC 449	2093 \pm 2	2087 \pm 2	-0.29 \pm 0.13
EOC 449	2826 \pm 2	2810 \pm 2	-0.57 \pm 0.10

Table 4.6. Integral Worth Comparison under Implicit Depletion Approximation Scheme in One Iteration

Time in Cycle	Homogeneous CE Int. Worth (€)	Discrete CE Int. Worth (€)	Δ (%)
BOC 400	1407 \pm 3	1414 \pm 3	0.50 \pm 0.30
MOC 400	1982 \pm 2	1979 \pm 2	-0.15 \pm 0.14
EOC 400	2809 \pm 2	2799 \pm 2	-0.36 \pm 0.10
BOC 424	1287 \pm 2	1294 \pm 1	0.54 \pm 0.17
MOC 424	2047 \pm 2	2044 \pm 2	-0.15 \pm 0.14
EOC 424	2819 \pm 2	2806 \pm 2	-0.46 \pm 0.10
BOC 449	1431 \pm 1	1439 \pm 2	0.56 \pm 0.16
MOC 449	2123 \pm 2	2115 \pm 2	-0.38 \pm 0.13
EOC 449	2824 \pm 2	2809 \pm 2	-0.53 \pm 0.10

Table 4.7. Integral Worth Comparison under Implicit Depletion Approximation Scheme in Two Iterations

Time in Cycle	Homogeneous CE Int. Worth (¢)	Discrete CE Int. Worth (¢)	Δ (%)
BOC 400	1407 \pm 3	1414 \pm 3	0.50 \pm 0.30
MOC 400	1982 \pm 2	1983 \pm 2	0.05 \pm 0.14
EOC 400	2807 \pm 2	2793 \pm 2	-0.50 \pm 0.10
BOC 424	1309 \pm 2	1315 \pm 2	0.46 \pm 0.22
MOC 424	2055 \pm 2	2048 \pm 2	-0.34 \pm 0.14
EOC 424	2818 \pm 2	2796 \pm 2	-0.78 \pm 0.10
BOC 449	1452 \pm 2	1491 \pm 2	2.69 \pm 0.20
MOC 449	2139 \pm 2	2133 \pm 2	-0.28 \pm 0.13
EOC 449	2818 \pm 2	2798 \pm 2	-0.71 \pm 0.10

Table 4.8. Integral Worth Comparison under Midpoint Depletion Approximation Scheme

Time in Cycle	Homogeneous CE Int. Worth (¢)	Discrete CE Int. Worth (¢)	Δ (%)
BOC 400	1407 \pm 3	1414 \pm 3	0.50 \pm 0.30
MOC 400	1983 \pm 1	1980 \pm 2	-0.15 \pm 0.11
EOC 400	2808 \pm 5	2793 \pm 2	-0.53 \pm 0.19
BOC 424	1235 \pm 1	1245 \pm 1	0.81 \pm 0.12
MOC 424	2025 \pm 2	2024 \pm 1	-0.05 \pm 0.07
EOC 424	2817 \pm 1	2802 \pm 1	-0.53 \pm 0.05
BOC 449	1380 \pm 1	1388 \pm 1	0.58 \pm 0.10
MOC 449	2099 \pm 1	2091 \pm 1	-0.38 \pm 0.07
EOC 449	2821 \pm 1	2805 \pm 1	-0.57 \pm 0.05

Regardless of the approximations employed, the discrete CE reactivity worth is consistently comparable to that of the homogeneous CEs. Discrepancies in differential worth are generally within 4 percent and are never greater than 11 percent, with these greatest discrepancies occurring at EOC when the CEs are completely withdrawn. Indeed, the best agreement in differential worth between designs is observed in the best-estimate midpoint scheme, where discrepancies are generally within 2 percent. Agreement between designs is even better for integral worth, where discrepancies are generally within 1 percent and never greater

than 3 percent. Furthermore, these calculations were carried out with the fixed cycle 400 withdrawal schedule, which only approximates the true CE positions in later cycles. It can therefore be expected these minor discrepancies can be recovered by adjusting the CE positions. Even so, the fact that the discrete CE design performs comparably to the homogeneous design under the fixed withdrawal schedule grants significant confidence in their operational feasibility.

4.2.4. Power Peaking

Having characterized the impact of each depletion scheme on the CE isotopics and reactivity control characteristics, the depletion scheme applied for the remaining analyses is now fixed to the best-estimate midpoint scheme.

Some insight into the effect of absorber discretization on core power peaking was granted in the previous chapter in observing the fission rate density distribution. This is roughly analogous to the absolute power distribution; however, a more precise metric for power peaking is the relative fission density (RFD) distribution – that is, the ratio of the pointwise fission rate density to the core average fission rate density. This yields the peaking factor distribution in the core.

Once again, it is difficult to observe differences between cases from plots of RFD alone; in this case the pointwise absolute difference in RFD is calculated since the RFD is already a normalized quantity. Plots of RFD itself are given in Appendix A. Figures 4.13-4.15 respectively give the absolute difference in RFD between the discrete and homogeneous CE cases and the corresponding uncertainty at BOC 400, 424, and 449 plotted in the r-z plane. Only BOC times are considered because this is the limiting case for power peaking, as the power distribution flattens as the cycle progresses. Differences in RFD remain limited throughout the CE life, not exceeding 2 percent. Further, as observed in the previous chapter and confirmed here, the

discrete CE design introduces a generally radially outward shift in the power distribution arising from the increased free thermal flux in the vicinity of the black CE regions. This is expected to improve thermal margins relative to the homogeneous CE case, where peak power occurs at the inner edge of the OFE [51]. The fixed use of the cycle 400 CE withdrawal schedule also impacts the RFD calculations, and adjusting the CE positions in practice will yield RFD distributions different from those calculated here. However, these results ultimately indicate that the discrete CEs carry power shaping capability that is comparable to the homogeneous CEs.




Figure 4.13. Absolute Difference in RFD between Homogeneous and Discrete CE Cases (left) and Uncertainty (right) at

BOC 400




Figure 4.14. Absolute Difference in RFD between Homogeneous and Discrete CE Cases (left) and Uncertainty (right) at BOC 424




Figure 4.15. Absolute Difference in RFD between Homogeneous and Discrete CE Cases (left) and Uncertainty (right) at BOC 449

As a final comment on power peaking, there is reasonable concern regarding the possibility of introducing small-scale local peaking in the fuel arising from uneven coverage of the CE absorbing region by the discrete absorber compacts. Significant “micro-peaking” may yield large local thermal gradients that could challenge the fuel integrity. To investigate this, the BOC 400 RFD is tallied in a very fine azimuthal mesh in the outermost radial edge of the OFE. This local RFD and its uncertainty are plotted in the ϕ - z plane, viewed as if the OFE were

unrolled into a flat sheet, in Figure 4.16; no local peaking is observed, and thus it is concluded that absorber compact coverage of the CE absorbing regions is sufficient, and no challenging thermal gradients are introduced. The peaks that are present arise from the absorber-free welds between CE quadrants.

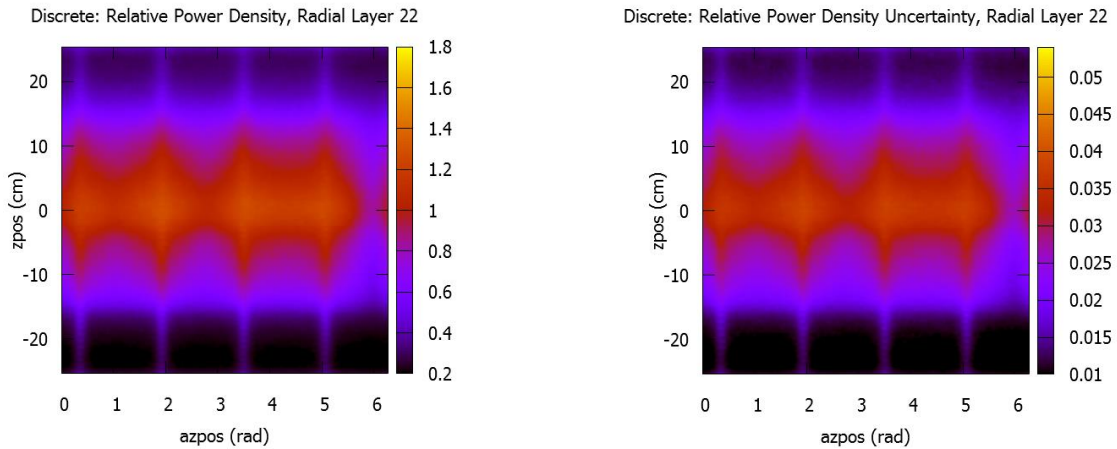


Figure 4.16. Local Peaking in Outermost OFE Edge (left) and Uncertainty (right) at BOC 400

4.2.5. Photon Heating

Photon heating is tallied along the axial extent of the CEs in each design case. Again, for maximum clarity, only the pointwise relative difference in photon heating is given here; photon heating plots are included in Appendix A. Figures 4.17-4.19 give the relative difference in photon heating between the discrete and homogeneous CEs and the corresponding uncertainty at BOC 400, 424, and 449, respectively. BOC is again the limiting case for photon heating, as the CE absorbing regions are positioned closest to the core midplane and thus receive the greatest photon flux. Absorber discretization yields oscillations in the discrete CE photon heating, but as expected, the discrete CEs experience consistently greater photon heating than the homogeneous CEs due to denser local concentration of the high-Z absorbing materials. The greatest difference occurs in the OCE gray region and is limited within 10 percent in each cycle considered. This does not raise any immediate prohibitive concern, but it is an important consideration in

transitioning to the implementation of the new CE design. Additional thermal-hydraulic analysis is needed to assess the adequacy of the existing cooling system settings to prevent accelerated corrosion and truncation of the useful CE life.




Figure 4.17. Relative Difference in Photon Heating between Homogeneous and Discrete CE Cases (left) and Uncertainty (right) at BOC 400




Figure 4.18. Relative Difference in Photon Heating between Homogeneous and Discrete CE Cases (left) and Uncertainty (right) at BOC 424

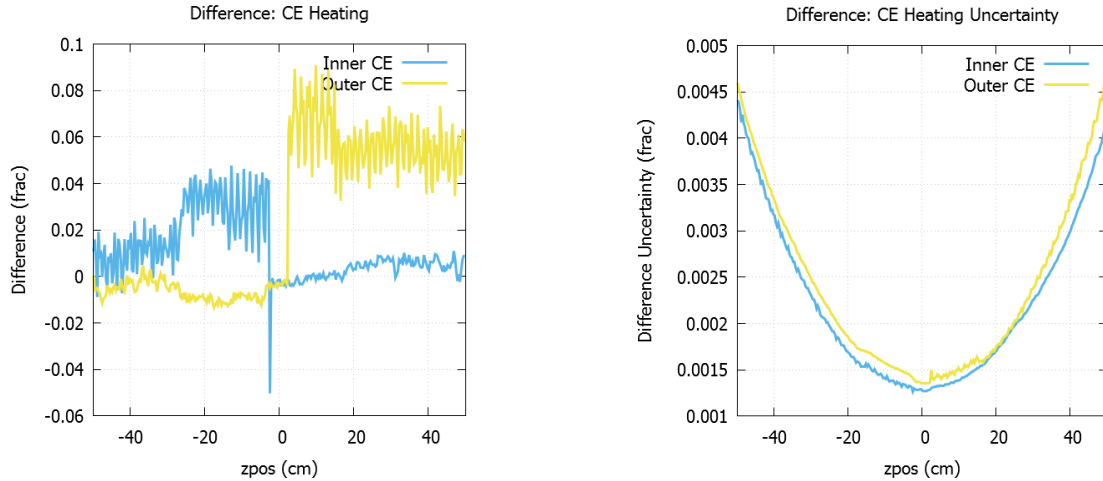


Figure 4.19. Relative Difference in Photon Heating between Homogeneous and Discrete CE Cases (left) and Uncertainty (right) at BOC 449

4.3. Summary and Discussion

Single- and multi-cycle depletion approximations have been developed to facilitate depletion analysis over the 50-cycle lifetime of the HFIR CEs and enable lifetime performance evaluation of the additively manufactured CEs to determine their operational feasibility. These approximations are based on fixed representative conditions determined from individual explicit depletion simulations, and it is found that the lifetime isotopics predictions arising from these approximations exhibit significant sensitivity to these representative conditions (particularly the flux spectra). However, this spread in isotopics predictions has only limited impact on the corresponding reactivity worth predictions. Furthermore, regardless of the depletion scheme applied, the discrete CE design consistently exhibits reactivity worth comparable to that of the homogeneous CEs even without correcting for changes to the critical CE positions. Power shaping by the discrete CEs is also seen to be comparable to that of the homogeneous CEs, again before correcting the CE positions.

It is therefore concluded that there are no immediate operational impediments to the use of additively manufactured CEs in HFIR, and that they are capable of maintaining reactor safety

over their full lifetime. One issue that needs to be addressed outside of this neutronics analysis is the increase in photon heating in the discrete CE design, which, if left unchecked, may accelerate corrosion of the CE clad surface. This accelerated damage could potentially shorten the useful CE life and partially offset the economic advantage introduced by employing UAM fabrication. However, this increase in photon heating is limited and not expected to be prohibitive. Thus it is maintained that the additively manufactured CEs provide sufficient reactor safety in terms of reactivity control and power peaking over their useful life, and therefore appear operationally feasible for use in HFIR.

CHAPTER 5. DESIGN OPTIMIZATION

The preceding analyses have demonstrated acceptable performance of the additively manufactured CEs over long-term operation of HFIR with no challenges to core physics behavior or reactor safety, concluding that they are feasible for practical application. This was determined based on computational models of the discrete CEs that were constructed to be “equivalent” to the original CE design by matching the total absorber content in each absorbing zone. Though this equivalent discrete CE design was determined to be neutronic and operationally feasible for use in HFIR, it still exhibits discrepancies in performance and core physics relative to the homogeneous CEs. These discrepancies may be mitigated by adjusting the discrete CE design, taking advantage of new design parameters introduced by the UAM process that are not applicable to the homogeneous CE design. Only the uniform absorber concentration in the gray and black regions could be adjusted in the homogeneous CEs, whereas absorber discretization arising from UAM fabrication allows the discrete absorber compact size and compact array spacing, in addition to the absorber concentration in the compacts, to be varied, providing multidimensional means to adjust the discrete CE performance and corresponding core physics. As it is desired to avoid reevaluation of the established HFIR safety documentation and provide as seamless of a transition as possible to the new CE design, this additional design freedom can be engaged to minimize changes in the HFIR core physics introduced by the discrete CEs and ensure conformity with existing safety margins. This endeavor is formulated as an optimization problem that is the focus of this chapter.

5.1. Optimization Problem Formulation

5.1.1. Objective

The discrete CE design and its impact on the HFIR core physics can be characterized by any of a broad collection of performance indicators, including (but not limited to) those explored in the previous chapter, namely, the core power distribution and CE reactivity worth. While reactivity worth is critical to reactor safety, it does not make for a valuable optimization objective; increasing the discrete CE reactivity worth or closely matching it to that of the homogeneous CEs is not necessary as long as sufficient shutdown margin and power level control capability are maintained. The only other reactor safety consideration that is directly affected by the CE design is the core power distribution and the associated maintenance of sufficient core cooling. Thermal safety margins are narrowest at the core coolant outlet, where the coolant temperature is greatest and pressure is least. These outlet conditions depend on the core power shape in both the axial and radial dimensions, which drives the spatially-dependent heating of the coolant as it flows through the core (azimuthal variations exist as well but are far lesser in magnitude than those in the radial and axial dimensions). Compliance with thermal margins can therefore be guaranteed if the core power distribution under the new CE design is matched to that under the homogeneous CEs, ideally yielding identical core outlet conditions.

Thus, the optimization objective is to minimize discrepancies in the core power distribution between CE designs. The objective function therefore takes the following form:

$$f(\mathbf{x}) = \sum_i (\Delta RFD_i)^2 \quad (5.1)$$

where f is the objective, \mathbf{x} is the vector of design variable settings, ΔRFD_i is the absolute difference in relative fission density between the discrete and homogeneous CE designs at mesh

point i , and the summation is taken over the entire tally mesh of the MCNP fission rate density calculation. Squaring the difference in RFD penalizes large individual pointwise discrepancies more heavily than multiple small pointwise discrepancies, favoring uniform conformity. RFD is assessed only at BOC 400, as practical operational uncertainties preclude high-confidence predictions of power distribution in cycles beyond cycle 400. The design vector \mathbf{x} is composed of the following parameters:

$$\mathbf{x} = \langle s_{ig}, s_{ib}, s_{og}, s_{ob}, r, t, a_g, a_b \rangle \quad (5.2)$$

where s is the compact array angular spacing, the subscripts “i” and “o” respectively refer to the ICE and the OCE, the subscripts “g” and “b” respectively refer to the gray and black absorber compacts, r is the compact radius, t is the compact thickness, and a is the absorber concentration.

5.1.2. Constraints

The design optimization is constrained by safety and operational considerations. It was shown in the previous chapter that the equivalent discrete CE design exhibits reactivity worth that is comparable to the homogeneous CEs, and it is expected that deficiencies could be recovered by adjusting the CE withdrawal schedule. Further, small design adjustments are not expected to prohibitively impact reactivity worth. Therefore, though it is a key performance indicator, reactivity worth is considered to be a soft constraint in this optimization problem. It will not be explicitly included in the problem formulation, but it should be verified that the optimized design maintains reactor safety and does not significantly stray from the reactivity worth of the homogeneous CEs.

Instead, a hard constraint arises from the practical necessity of maintaining the HFIR cycle length. Experimentalists depend on the sustained neutron source provided by HFIR for

research activities; any truncation of the neutron source arising from a loss in the HFIR cycle length would preclude a proposed CE design. To avoid running computationally costly depletion calculations, the HFIR cycle length will be represented by the BOC 400 neutron multiplication eigenvalue with the CE position fixed. The BOC neutron multiplication reflects the total BOC core reactivity content, including the fuel, CEs, and experiments, and therefore serves as a suitable surrogate for the HFIR cycle length. The cycle length constraint therefore takes the following form:

$$k \geq k_{hom} \quad (5.3)$$

where k is the neutron multiplication eigenvalue and the subscript “hom” refers to the homogeneous CE design. The value of k_{hom} was determined in Chapter 3 to be 1.00829 ± 0.00007 .

Geometric constraints should be included as well, as structural integrity requires that the compacts cannot be in contact with one another:

$$s_{ig}, s_{ib}, s_{og}, s_{ob} > 0 \quad (5.4)$$

No additional hard constraints are necessary. The cycle length constraint ensures sufficient reactivity control and reactivity worth. As observed in the previous chapter, photon heating of the CEs remains a concern, but is not expected to be prohibitive.

5.1.3. Problem Statement

The full optimization problem is stated as follows:

$$\begin{aligned}
& \text{Minimize } f(\mathbf{x}) = \sum_i (\Delta RFD_i)^2 \\
& \text{subject to} \\
& k \geq k_{hom} \\
& s_{ig}, s_{ib}, s_{og}, s_{ob} > 0
\end{aligned} \tag{5.5}$$

The functional forms of the objective and the neutron multiplication eigenvalue in terms of the design variables are unknown; though transport theory provides a mathematical framework for determining these reactor characteristics, the resultant equations can only practically be solved computationally. As the objective and the neutron multiplication eigenvalue depend on intricate transport phenomena, it is anticipated that they are highly nonlinear functions of the design variables. Traditional linear programming optimization methods are therefore not applicable, and it is necessary to construct a practical optimization approach that can accommodate computationally demanding transport calculations.

5.2. Methodology

The equivalent discrete CE design characterized in the preceding analyses serves as an initial feasible solution (from Chapter 3, it yields a neutron multiplication eigenvalue of 1.00960 ± 0.00007 , thus satisfying the cycle length constraint). The design variable settings for this nominal equivalent design are specified in Table 5.1.

Table 5.1. Nominal Equivalent Design Variable Settings

Design Variable	Value
s_{ig}	0.028516 rad
s_{ib}	0.027660 rad
s_{og}	0.027921 rad
s_{ob}	0.027128 rad
r	0.2795 cm
t	0.171 cm
a_g	90 w/o Ta
a_b	81.5 w/o Eu_2O_3

Intuition provides a general sense of the qualitative impact of each design variable; for example, adjustments of decreasing the compact spacing, increasing the compact size, and increasing the absorber concentration all introduce greater absorber content in the CEs, reducing neutron multiplication and shifting the power distribution radially inward. However, explicit calculations at discrete points in the design space are required to determine the precise magnitudes of these effects and their interactions. Optimization methods generally rely on well-defined mathematical representations of the objective and constraints as functions of the design variables to search throughout the multidimensional problem geometry, but no such precise formulations are available in the present application. The Nelder-Mead algorithm [52] provides an option for problems lacking formulaic representations, requiring only calculation of function values at test points in design space. This method constructs a simplex in the design space that is iteratively transformed according to prescribed rules regarding the objective values at its vertices, moving it through the design space toward an optimum solution. However, the Nelder-Mead algorithm exhibits slow convergence, and is only intended for application to unconstrained problems [52]. Instead, the mathematical behavior of the objective and the eigenvalue constraint in this problem shall be statistically inferred by employing response surface methodology (RSM) to guide the optimal design search [53, 54].

RSM is a tool commonly employed in the analysis of systems which cannot be readily described with closed-form mathematical expressions based on theory, requiring simulation or experimentation for response measurement [53-57]. Response surfaces are fit to data using least squares regression [53] and are accordingly statistical in nature; unless there is theoretical justification, they do not typically carry physical significance, only reflecting statistical trends in the data to which they are fit [53]. In simulation optimization [54, 55], RSM is employed to

guide solution searches based on limited sets of experimental runs. First-order surfaces are typically fit to initial sets of data far from the optimum, and higher-order terms are introduced as the search nears the optimum and the simulation data suggest increasing curvature [53, 54, 56]. Second-order response surfaces are typically considered sufficient to capture curvature near the optimum for practical purposes [53, 54, 56].

Construction of statistically reliable response surfaces is achieved by selecting design test points to which the surface is fit with adherence to principles of experimental design [53-57]. Such principles are invaluable for experiments with limited observational resources, guiding the selection of a limited number of test points to yield sufficient information about the system under inspection with the greatest possible efficiency [53, 57]. In this application, with limited computational resources, it is practically desirable to limit the number of design perturbations for which transport calculations are executed while still collecting sufficient data for response surface fitting. With 8 design variables, a full two-level factorial design [53, 54, 57] requires $2^8 = 256$ simulations at each design iteration, which is far too many to be considered computationally feasible. The computational burden is reduced by running a 1/16 fraction of the full factorial design, yielding a 2^{8-4} experimental design [53, 54, 57] that requires only 16 simulations at each design iteration. This is a “resolution V” design [53], meaning that all first- and second-order effects can be isolated with no confounding (the “V” designator indicates the number of design variables used to define the test matrix) [53]. The fractional factorial case matrix employed is shown in Table 5.2; the design variable settings are labeled as multiples of the perturbation magnitude and are referred to as “coded” variables in [53], such that the physical variable setting is equal to the nominal setting plus the specified multiple of the perturbation magnitude.

Table 5.2. Fractional Factorial Experiment Cases

Perturbation	Sig	Sib	Sog	Sob	r	t	a_g	a_b
Design Case	0.001 rad	0.001 rad	0.001 rad	0.001 rad	0.005 cm	0.01 cm	2 w/o	2 w/o
1	-1	-1	-1	-1	1	-1	-1	-1
2	1	-1	-1	-1	-1	1	-1	-1
3	-1	1	-1	-1	-1	-1	1	-1
4	1	1	-1	-1	1	1	1	-1
5	-1	-1	1	-1	-1	-1	-1	1
6	1	-1	1	-1	1	1	-1	1
7	-1	-1	-1	1	-1	-1	-1	-1
8	1	-1	-1	1	1	1	-1	-1
9	-1	1	1	-1	1	-1	1	1
10	1	1	1	-1	-1	1	1	1
11	-1	1	-1	1	1	-1	1	-1
12	1	1	-1	1	-1	1	1	-1
13	-1	-1	1	1	1	-1	-1	1
14	1	-1	1	1	-1	1	-1	1
15	-1	1	1	1	-1	-1	1	1
16	1	1	1	1	1	1	1	1

The experimental design specified in Table 5.2 most optimally allows for fitting first-order response surfaces. When necessary (i.e., as the search nears the optimum and the data indicate significant curvature), this design can be augmented to form a central composite design (CCD) that is optimal for fitting second-order surfaces. The additional cases needed for the CCD include a “center” point, corresponding to the nominal design variable settings, and 16 “axial” points chosen to maintain rotatability of the experimental design [53, 54]; these cases are specified in Table 5.3.

Table 5.3. CCD Cases

Design Case	Sig	Sib	Sog	Sob	r	t	ag	ab
17	0	0	0	0	0	0	0	0
18	2	0	0	0	0	0	0	0
19	-2	0	0	0	0	0	0	0
20	0	2	0	0	0	0	0	0
21	0	-2	0	0	0	0	0	0
22	0	0	2	0	0	0	0	0
23	0	0	-2	0	0	0	0	0
24	0	0	0	2	0	0	0	0
25	0	0	0	-2	0	0	0	0
26	0	0	0	0	2	0	0	0
27	0	0	0	0	-2	0	0	0
28	0	0	0	0	0	2	0	0
29	0	0	0	0	0	-2	0	0
30	0	0	0	0	0	0	2	0
31	0	0	0	0	0	0	-2	0
32	0	0	0	0	0	0	0	2
33	0	0	0	0	0	0	0	-2

A second-order response surface equation with 8 design variables will have 45 regression coefficients, thus requiring at least this many data points to carry out a least squares fit. However, the small fractional factorial employed here limits the total number of observations, with the CCD only providing 33 data points. To provide the additional necessary data, a sensitivity study is carried out by perturbing each design variable one at a time (this sensitivity data provides valuable physical insight on its own as well); these cases are specified in Table 5.4. Appending these sensitivity observations to the CCD yields 49 data points, enabling fitting of a second-order response surface.

Table 5.4. Sensitivity Cases

Design Case	Sig	Sib	Sog	Sob	r	t	ag	ab
34	1	0	0	0	0	0	0	0
35	-1	0	0	0	0	0	0	0
36	0	1	0	0	0	0	0	0
37	0	-1	0	0	0	0	0	0
38	0	0	1	0	0	0	0	0
39	0	0	-1	0	0	0	0	0
40	0	0	0	1	0	0	0	0
41	0	0	0	-1	0	0	0	0
42	0	0	0	0	1	0	0	0
43	0	0	0	0	-1	0	0	0
44	0	0	0	0	0	1	0	0
45	0	0	0	0	0	-1	0	0
46	0	0	0	0	0	0	1	0
47	0	0	0	0	0	0	-1	0
48	0	0	0	0	0	0	0	1
49	0	0	0	0	0	0	0	-1

The design cases listed in Tables 5.2-5.4 are incorporated into the optimization computational flow as follows. The initial nominal equivalent design, actually labeled as case 17 above, is assessed first, as well as the sensitivity cases (even though they were listed last above) to give a clear indication of curvature in the objective and constraining eigenvalue in the design space. Some statistical tests of the adequacy of the first-order fit are prescribed in [53]. If curvature is limited, only the fractional factorial cases are carried out to allow fitting of a first-order response surface. The next design test point is selected along the path of steepest descent of this surface within the range of design variables to which the surface was fit, taking care not to cross the feasibility boundaries set by the constraints (the eigenvalue boundary is set by its own response surface). The search carries on in this manner until significant curvature is identified, at which point the CCD cases are carried out to fit a second-order response surface. An optimum is identified if a local minimum of the response surface occurs in the range of perturbations evaluated, and is then assessed against the optimality condition:

$$f(x^*) \leq f(x) \forall x \text{ s.t. } \|x - x^*\| \leq \delta \tag{5.6}$$

where \mathbf{x}^* is the optimal design vector and δ is the distance in design space defining locality. Equation 5.6 defines a local optimum by limiting the design space to that within the range of design perturbations evaluated. A global optimum search is considered neither feasible nor necessary, as it would require additional searches far from the nominal design. This process is diagrammed in Figure 5.1.

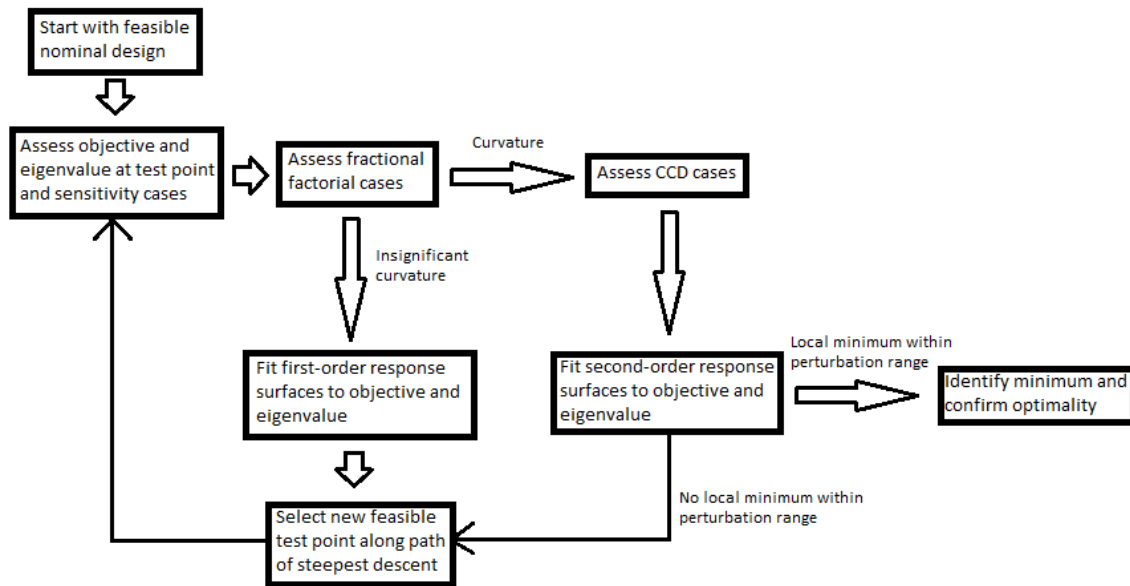


Figure 5.1. Optimization Computational Flow

5.3. Results

5.3.1. Nominal Equivalent Design Sensitivities

The sensitivity of the nominal equivalent CE design to the design variables is of particular characterizing interest. The objective and eigenvalue results of the nominal design and sensitivity cases are listed in Table 5.5 below; each case is labeled consistent with those in Tables 5.2-5.4. In this section, though the transport calculations yield uncertainties in the quantities of interest, only the best estimate numbers are reported, as uncertainties are not factored into response surface fitting.

Table 5.5. Nominal Equivalent Design Sensitivity Results

Design Case	f	k
17	0.019058	1.00961
34	0.056558	1.01074
35	0.046272	1.00849
36	0.029758	1.00994
37	0.025002	1.00936
38	0.079689	1.01038
39	0.049964	1.00874
40	0.082432	1.00971
41	0.046290	1.00936
42	0.016070	1.00833
43	0.055705	1.01097
44	0.055844	1.00759
45	0.099701	1.01174
46	0.019261	1.00845
47	0.047878	1.01074
48	0.019456	1.00928
49	0.024094	1.00982

These results reveal invaluable information about the behavior of the objective and eigenvalue responses to design perturbations in the neighborhood of the nominal equivalent design; specifically, all perturbations except for one (the positive compact radius perturbation, case 42) yielded an increase in the objective value relative to the nominal equivalent design. This indicates significant curvature and suggests that the optimum is very near in the design space to the nominal equivalent design. Such a fortunate circumstance is likely a consequence of closely matching the discrete CE absorber content to that of the homogeneous CEs, necessitating only minimal design adjustment to compensate for the effects of absorber discretization and match the RFD distribution yielded by the homogeneous CEs. Thus the CCD cases can be immediately carried out for fitting a second-order response surface and identifying an optimal design in a single iteration.

5.3.2. Response Surface Fitting

The objective and eigenvalue results of all cases from the CCD and sensitivity study are compiled in Table 5.6. These results are fed into MATLAB to automate least squares fitting of

second-order response surfaces for the objective and eigenvalue [58]. These surfaces take the following forms:

$$R_f = \alpha_0 + \sum_{i=1}^8 \alpha_i x_i + \sum_{i=1}^8 \alpha_{ii} x_i^2 + \sum_{i=1}^7 \sum_{j>i} \alpha_{-ij} x_i x_j \quad (5.7)$$

$$R_k = \beta_0 + \sum_{i=1}^8 \beta_i x_i + \sum_{i=1}^8 \beta_{ii} x_i^2 + \sum_{i=1}^7 \sum_{j>i} \beta_{-ij} x_i x_j \quad (5.8)$$

where R is the response, the “f” and “k” subscripts respectively refer to the objective and eigenvalue, α and β are the regression coefficients, and x_i refers to component “i” of the design vector defined in Equation 5.2. The design variable settings are kept in their coded units (multiples of the design perturbation) for response surface fitting so as to achieve reasonable scaling across all design variables, which take on nominal values spanning three orders of magnitude. The resulting fit coefficients are listed in Table 5.7 for the objective surface and Table 5.8 for the eigenvalue surface.

The values of these regression coefficients indicate the extent of the effect of each design variable and their interactions on the objective and eigenvalue. The results indicate particular sensitivity of the objective and eigenvalue to the compact size. It is noteworthy that zero coefficients are obtained for most of the interaction terms involving the compact thickness and absorber concentration variables, indicating that the effects of these variables are largely independent of the others. The exception is a surprisingly significant interaction between the absorber compact spacing in the ICE black region and the gray absorber concentration.

Table 5.6. Compiled Design Perturbation Results

Design Case	f	k
1	0.021018	1.00920
2	0.082287	1.01002
3	0.074430	1.00990
4	0.184653	1.00550
5	0.209609	1.01316
6	0.039689	1.00888

Design Case	f	k
7	0.109894	1.01241
8	0.044107	1.00788
9	0.049549	1.00902
10	0.046291	1.00966
11	0.072683	1.00798
12	0.104516	1.00882
13	0.184725	1.01132
14	0.135259	1.01214
15	0.169715	1.01206
16	0.047764	1.00756
17	0.019058	1.00961
18	0.158885	1.01188
19	0.134841	1.00739
20	0.076767	1.01034
21	0.055985	1.00901
22	0.145056	1.01152
23	0.096151	1.00759
24	0.073253	1.01019
25	0.040676	1.00900
26	0.039536	1.00707
27	0.130558	1.01234
28	0.077040	1.00599
29	0.205474	1.01399
30	0.049521	1.00724
31	0.117593	1.01198
32	0.034972	1.00915
33	0.056534	1.01023
34	0.056558	1.01074
35	0.046272	1.00849
36	0.029758	1.00994
37	0.025002	1.00936
38	0.079689	1.01038
39	0.049964	1.00874
40	0.082432	1.00971
41	0.046290	1.00936
42	0.016070	1.00833
43	0.055705	1.01097
44	0.055844	1.00759
45	0.099701	1.01174
46	0.019261	1.00845
47	0.047878	1.01074
48	0.019456	1.00928
49	0.024094	1.00982

Table 5.7. Objective Response Surface Regression Coefficients

Coefficient	Value
α_0	3.03922×10^{-2}
α_1	1.01422×10^{-2}
α_2	7.30865×10^{-3}
α_3	1.42172×10^{-2}
α_4	1.00974×10^{-2}
α_5	-1.95997×10^{-2}
α_6	-2.57677×10^{-2}
α_7	-1.37994×10^{-2}
α_8	-3.31243×10^{-3}
α_{11}	2.86415×10^{-2}

Coefficient	Value
α_{22}	8.28961×10^{-3}
α_{33}	2.32517×10^{-2}
α_{44}	8.25039×10^{-3}
α_{55}	1.31832×10^{-2}
α_{66}	2.88729×10^{-2}
α_{77}	1.26963×10^{-2}
α_{88}	3.10742×10^{-3}
α_{12}	1.50407×10^{-2}
α_{13}	-3.01271×10^{-2}
α_{14}	-1.27241×10^{-2}
α_{15}	1.14645×10^{-2}
α_{16}	0
α_{17}	0
α_{18}	0
α_{23}	-2.71901×10^{-2}
α_{24}	-5.10799×10^{-3}
α_{25}	1.29567×10^{-2}
α_{26}	0
α_{27}	-5.81673×10^{-2}
α_{28}	0
α_{34}	1.39757×10^{-2}
α_{35}	-1.19113×10^{-2}
α_{36}	0
α_{37}	0
α_{38}	0
α_{45}	-3.28108×10^{-3}
α_{46}	0
α_{47}	0
α_{48}	0
α_{56}	0
α_{57}	0
α_{58}	0
α_{67}	0
α_{68}	0
α_{78}	0

Table 5.8. Eigenvalue Response Surface Regression Coefficients

Coefficient	Value
β_0	1.00958
β_1	1.11543×10^{-3}
β_2	3.03476×10^{-4}
β_3	9.78810×10^{-4}
β_4	2.90769×10^{-4}
β_5	-1.30808×10^{-3}
β_6	-2.02257×10^{-3}
β_7	-1.19752×10^{-3}
β_8	-2.41190×10^{-4}
β_{11}	1.51096×10^{-5}
β_{22}	2.65802×10^{-5}
β_{33}	-6.94921×10^{-6}
β_{44}	9.91962×10^{-7}
β_{55}	3.36390×10^{-5}
β_{66}	1.01580×10^{-4}
β_{77}	8.05079×10^{-6}
β_{88}	2.42273×10^{-5}
β_{12}	-1.56250×10^{-5}

Coefficient	Value
β_{13}	-3.12500×10^{-6}
β_{14}	-9.37500×10^{-6}
β_{15}	5.06250×10^{-5}
β_{16}	0
β_{17}	0
β_{18}	0
β_{23}	6.87500×10^{-6}
β_{24}	-9.37500×10^{-6}
β_{25}	4.37500×10^{-6}
β_{26}	0
β_{27}	-6.34821×10^{-5}
β_{28}	0
β_{34}	-6.87500×10^{-6}
β_{35}	2.18750×10^{-5}
β_{36}	0
β_{37}	0
β_{38}	0
β_{45}	-3.43750×10^{-5}
β_{46}	0
β_{47}	0
β_{48}	0
β_{56}	0
β_{57}	0
β_{58}	0
β_{67}	0
β_{68}	0
β_{78}	0

5.3.3. Optimal Design

The optimal design variable settings are initially sought at a local minimum in the objective response surface, where the gradient is zero in all components. The design variable settings satisfying this requirement are given in Table 5.9, and are now converted to physical units rather than coded units. This point in design space is the only critical point of the objective response surface, and is therefore a global minimum of this response surface (though the surface may only reasonably be applied within the range of design perturbations used for its fitting). The objective value predicted at this point by the objective response surface is 0.0035885, indicating excellent agreement of the predicted power distribution with respect to that of the homogeneous CE design; however, the predicted neutron multiplication eigenvalue at this location is 1.00526, which violates the cycle length constraint. This global optimum therefore lies outside the region of feasibility.

Table 5.9. Objective Response Surface Global Minimum

Variable	Value
S _{ig}	0.027875 rad
S _{ib}	0.027574 rad
S _{og}	0.027558 rad
S _{ob}	0.026465 rad
r	0.28359 cm
t	0.17546 cm
a _g	90.6906 w/o Ta
a _b	82.5660 w/o Eu ₂ O ₃
Objective f	0.00335885
Eigenvalue k	1.00526

In order to identify a feasible optimum design, the neutron multiplication eigenvalue is fixed to its minimum feasible value of 1.00829, and the minimum objective value is sought in the region where the objective response is bound by the constraining eigenvalue surface $k(\mathbf{x}) = 1.00829$. The method of Lagrange multipliers [59] is applied to identify the location of the minimum objective value in this intersection. The resulting optimum feasible design variables are listed in Table 5.10; via explicit transport calculation (rather than the response surface estimation), the corresponding objective value is 0.018528, which is a small yet nontrivial improvement over the nominal equivalent design. This result highlights the difficulty of achieving a power distribution with the discrete CEs that is identical to that under the homogeneous CEs without truncating the HFIR cycle length, which is a fundamental limitation of the discrete CE design. The objective value predicted by the response surface is 0.014663, indicating some deficiency in the quality of the fit. Improvements could be realized by including higher order terms in the response surface, but this would of course introduce additional computational burden. The eigenvalue response surface is an imperfect fit as well, predicting the minimum eigenvalue of 1.00829 as opposed to the explicitly calculated 1.00853.

Table 5.10. Optimal Design Variable Settings

Variable	Value
S _{ig}	0.028325 rad
S _{ib}	0.027669 rad
S _{og}	0.027937 rad
S _{ob}	0.026574 rad
r	0.28118 cm
t	0.17241 cm
a _g	90.2965 w/o Ta
a _b	81.9078 w/o Eu ₂ O ₃
Objective f	0.018528
Eigenvalue k	1.00853

To confirm optimality of the design identified in Table 5.10, a final sensitivity study is carried out, perturbing each variable from its optimal setting one at a time. Smaller perturbations are employed in this case, as the preceding results indicate significant sensitivity of the objective to most of the design variables. These results are given in Table 5.11. Every perturbation tested yields an increase in the objective value, and many also violate the eigenvalue constraint. Though an additional response surface fit could be carried out in the neighborhood of this design, this is deemed unnecessary, and the design identified in this section is considered optimal within the range of first-order perturbations tested.

Table 5.11. Optimal Design Sensitivity Results

Design Variable	Perturbation	f	k
S _{ig}	0.0005 rad	0.029842	1.00880
S _{ig}	-0.0005 rad	0.041560	1.00768
S _{ib}	0.0005 rad	0.050050	1.00812
S _{ib}	-0.0005 rad	0.035409	1.00828
S _{og}	0.0005 rad	0.061768	1.00881
S _{og}	-0.0005 rad	0.048745	1.00777
S _{ob}	0.0005 rad	0.042890	1.00847
S _{ob}	-0.0005 rad	0.047678	1.00833
r	0.001 cm	0.044376	1.00796
r	-0.001 cm	0.047533	1.00871
t	0.001 cm	0.061449	1.00731
t	-0.001 cm	0.050533	1.00936
a _g	1 w/o	0.049793	1.00749
a _g	-1 w/o	0.047467	1.00896
a _b	1 w/o	0.046003	1.00835
a _b	-1 w/o	0.046957	1.00834

5.4. Optimal Design Assessment

This section now assesses the core physics effects of the optimal discrete CE design identified above, as well as its operational performance. Figure 5.2 gives the absolute difference in RFD between the optimal design and the original homogeneous CEs and the corresponding uncertainty at BOC 400; the plot of the RFD itself is given in Appendix A. As suggested by the objective value, only a slight improvement in the power distribution is achieved. Even so, the differences are sufficiently small (well within 2 percent) that thermal margins are not expected to be violated. The optimized CEs are also withdrawn to the nominal BOC 400 position in this calculation, which would not be the case in practice; adjusting the CE position to an appropriate withdrawal length will yield a different power distribution. However, given the close match in RFD distribution, it is sufficient to conclude that the optimized CE design retains its power shaping capability.




Figure 5.2. Absolute Difference in RFD between Homogeneous and Optimal CE Designs (left) and Uncertainty (right) at BOC 400

Finally, the differential and integral worth of the optimized CEs are evaluated at BOC using the methodology of the previous chapter. The results are presented in Table 5.12; the worth results of the homogeneous and nominal equivalent CE designs are repeated here for

comparison. The optimal CE design retains reactivity worth comparable to the homogeneous and nominal equivalent CE designs, and is thus not expected to sacrifice power level control or shutdown margin, yielding the conclusion that it is feasible for use in the operation of HFIR. Depletion analysis of the optimal design is not considered necessary at this point, as there is no reason to expect divergent behavior in later cycles (further, uncertainty in the withdrawal schedule of the optimized CEs precludes reliable assessment of their long-term behavior).

Table 5.12. Optimal Design BOC 400 Reactivity Worth Comparison

Design	Diff. Worth (¢/cm)	Int. Worth (¢)
Optimal	104.5 ± 0.7	1412 ± 2
Nominal Equivalent	103.9 ± 1.6	1414 ± 3
Homogeneous	102.8 ± 1.6	1407 ± 3

5.5. Summary and Discussion

RSM has been employed to facilitate development of an optimal CE design which matches the HFIR core power distribution under the original homogeneous CEs as closely as possible without truncating the cycle length. The second-order response surfaces are seen to sacrifice accuracy in estimating the objective and constraint responses, but the selected design nevertheless satisfies the optimality condition within the range of design perturbations evaluated. Only small design adjustments were implemented, and these yielded only limited improvement in the power distribution; even so, the remaining discrepancy in the power distribution is not expected to be prohibitive. The optimal CE design also exhibits power shaping capability and reactivity worth that is comparable to the homogeneous CEs and the nominal equivalent discrete CE design, and thus appears feasible for the operation of HFIR. This optimal CE design is therefore expected to provide as seamless a transition as possible in implementing the additively manufactured CEs in HFIR.

CHAPTER 6. CONCLUSION

Early demonstrations of UAM fabrication of the HFIR CEs have verified the significant potential offered by this technology in reducing the time, cost, and labor associated with the traditional HFIR CE fabrication process [17, 18, 32], providing impetus to move forward with more detailed feasibility studies. This particular work has taken on a holistic neutronic assessment of the unique CE design with discretized absorbers that is yielded by UAM, addressing its operational feasibility in terms of the impact on the HFIR core physics, operational performance, and reactor safety. The results of these analyses indicate no neutronic or performance impediments to the practical application of additively manufactured CEs in HFIR, concluding that they are indeed neutronic and operationally feasible. An additional effort was undertaken to optimize the additively manufactured CE design, taking advantage of new design variables introduced by UAM fabrication that are not applicable to the homogeneous CEs. An optimal design was developed that minimizes changes in the HFIR core physics, thereby eliminating the need to reconsider the well-established HFIR safety documentation and carry out new safety analyses.

The impact of absorber discretization on the HFIR core physics was assessed using computational models of the discrete CEs constructed to be “equivalent” to the original homogeneous CEs by preserving the absorber content in each CE absorbing region. Self-shielding arising from absorber discretization yields a loss in absorption effectiveness and a corresponding increase in free thermal neutron flux, particularly in the vicinity of the black CE regions. These extra free thermal neutrons yield a limited increase in neutron multiplication and a slight shift in the core power distribution radially outward. The absorption physics of the discrete CEs were found to be complicated by several competing effects arising from the discrete

CE geometry, discrete absorber self-shielding, and spectral differences between the ICE and OCE. The net result is a loss in absorption rate by the ICE but a surprising increase in absorption rate by the OCE. However, changes in the macroscopic core characteristics of HFIR were concluded to be limited and not prohibitive, confirming the neutronic feasibility of employing additively manufactured CEs in HFIR.

The analysis then turned to assessment of the discrete CE operational performance over their useful lifetime. Methodology was developed for computationally efficient approximations of CE absorber depletion over long periods of operation (50 HFIR cycles), minimizing the requisite number of computationally costly explicit depletion simulations. Though these methods yielded varying precision in CE isotopics, the discrete CEs were found to exhibit performance comparable to the homogeneous CEs regardless of the approximation employed. Discrepancies in differential and integral reactivity worth were limited to less than 10 percent, and shifts in the core power distribution were limited to less than 2 percent difference in RFD. Further, these evaluations were carried out with the CE designs in identical withdrawal positions; this would not practically be the case due to differences in the reactivity control capability of each CE design. Thus it is expected that the identified performance discrepancies can be recovered by employing a CE withdrawal schedule developed for specific application to the discrete CEs. Absorber discretization was found to yield an increase in photon heating of the CE absorbing regions, but this is not expected to be prohibitive, and the associated thermal-hydraulic analysis is reserved for a separate endeavor. It is therefore concluded that the additively manufactured CEs are feasible for the operation of HFIR.

Finally, to provide a seamless transition between CE designs in HFIR, the discrete CE design was optimized with the intent of minimizing its impact on the core power distribution, so

that maintenance of sufficient core cooling is guaranteed and the HFIR safety documentation does not need to be reevaluated. An optimization methodology is developed based on application of RSM to statistically infer the mathematical relationship between the CE performance and the design variables. It is found that the optimal design was very near in the design space to the nominal equivalent design, and a brief performance assessment confirms the operational feasibility of the optimal design and its minimized impact on the core power distribution.

A great deal of additional analysis and testing is required before additively manufactured CEs can practically be employed in HFIR, including (but not limited to) transient analyses, thermal-hydraulic studies, and mechanical testing of physical samples. Nevertheless, the conclusions of this work instill confidence and stoke enthusiasm for the potential role of UAM in the nuclear industry. The analyses undertaken here have verified the neutronic and operational feasibility of additively manufactured CEs in HFIR, and their fabrication has been successfully demonstrated [18, 32]. Further, the body of UAM research literature suggests that solutions are available to address mechanical and material deficiencies of UAM-fabricated parts [26, 27, 29, 30]. The application of UAM considered in this work is narrow in scope, and thus future endeavors should seek to broaden the application of UAM by taking on other reactor components and incorporating metal matrix materials beyond the Al-6061 structure in the HFIR CEs. Beyond such fabrication demonstrations, it must also be shown that UAM facilities can be developed for industrial-scale manufacturing while maintaining product reliability.

Though much additional work remains before UAM can be adopted prolifically in the nuclear industry, the findings of the analyses carried out here are promising. Nuclear technology is a powerful tool that can provide abundant energy to support high standards of living nationally and globally with minimal environmental impact, but it must overcome unfavorable market

forces to improve its practical attractiveness. UAM represents a promising potential solution in addressing this challenge by simplifying the fabrication of components and subsequent construction to cut the upfront capital costs of new nuclear power projects, which remain perhaps the greatest hindrance to investment. It is hoped that the results of this and related work spur wider interest in additional applications of UAM in the nuclear industry and ultimately help to accelerate the global advancement of nuclear power technology. This is a crucial humanitarian endeavor, as the world is operating on borrowed time in combating climate change.

APPENDICES

APPENDIX A. MISCELLANEOUS NEUTRONICS PLOTS

This appendix includes plots of the neutronics quantities discussed in Chapters 3 and 4.




Figure A.1. Thermal Flux (left column) and Uncertainty (right column) in Homogeneous (top row), Approximate (middle row), and Explicit (bottom row) Models




Figure A.2. Fission Rate Density (left column) and Uncertainty (right column) in Homogeneous (top row), Approximate (middle row), and Explicit (bottom row) Models

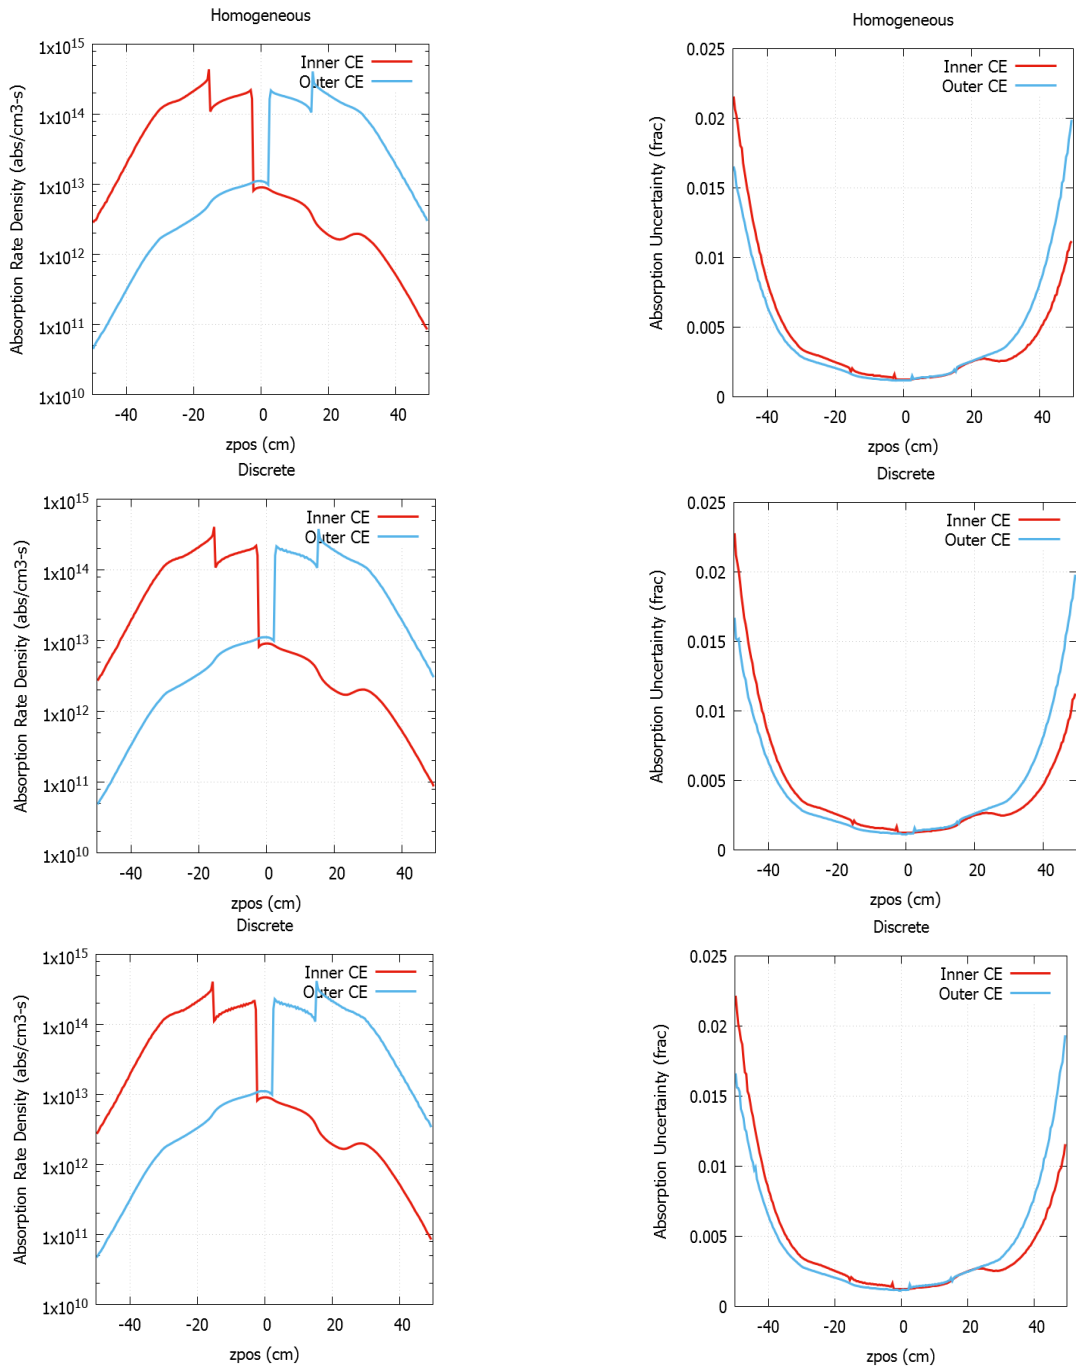


Figure A.3. Absorption Rate Density (left column) and Uncertainty (right column) in Homogeneous (top row), Approximate (middle row), and Explicit (bottom row) Models

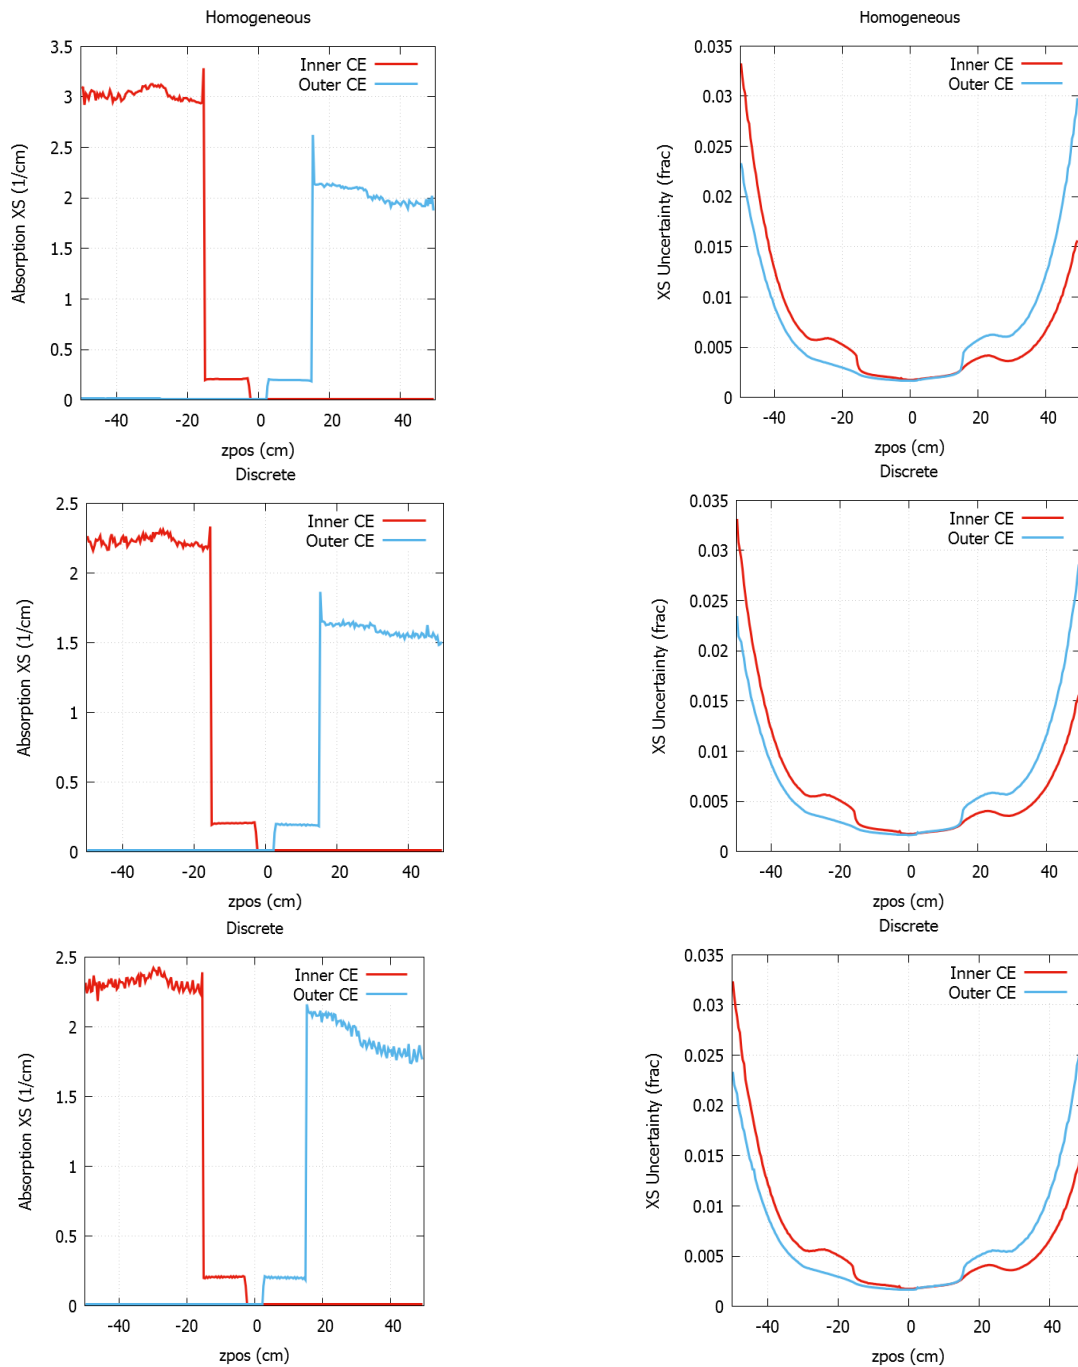


Figure A.4. Effective Thermal Absorption Cross Section (left column) and Uncertainty (right column) in Homogeneous (top row), Approximate (middle row), and Explicit (bottom row) Models




Figure A.5. BOC 400 Relative Fission Density (left column) and Uncertainty (right column) in Homogeneous (top row) and Explicit (bottom row) Models




Figure A.6. BOC 424 Relative Fission Density (left column) and Uncertainty (right column) in Homogeneous (top row) and Explicit (bottom row) Models




Figure A.7. BOC 449 Relative Fission Density (left column) and Uncertainty (right column) in Homogeneous (top row) and Explicit (bottom row) Models

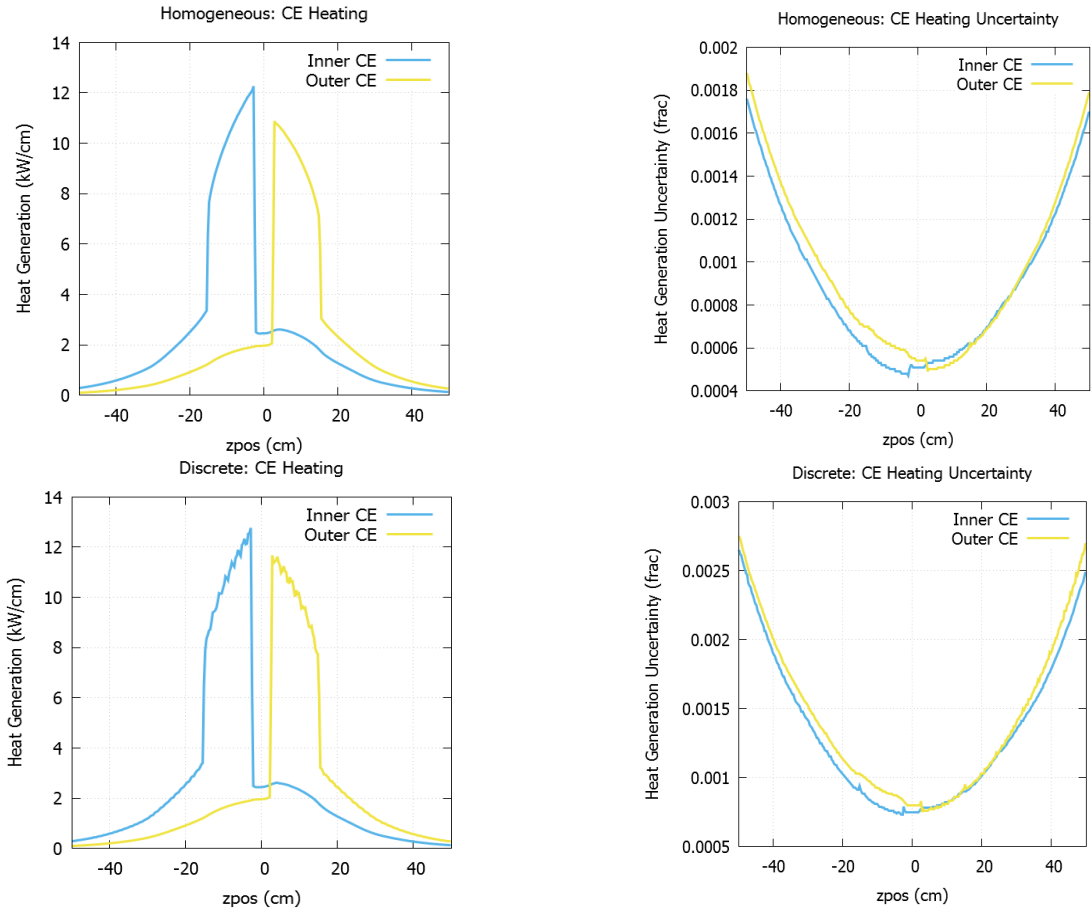


Figure A.8. BOC 400 Photon Heating (left column) and Uncertainty (right column) in Homogenous (top row) and Explicit (bottom row) Models

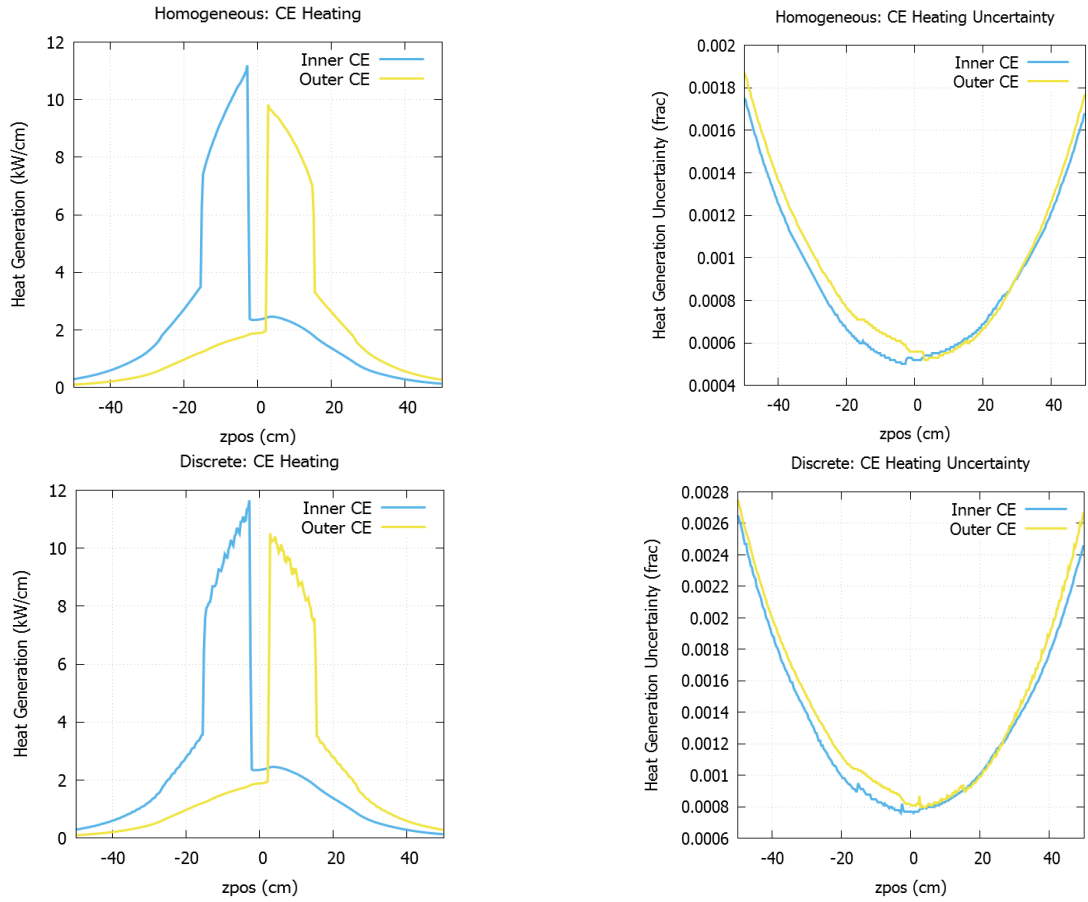


Figure A.9. BOC 424 Photon Heating (left column) and Uncertainty (right column) in Homogeneous (top row) and Explicit (bottom row) Models

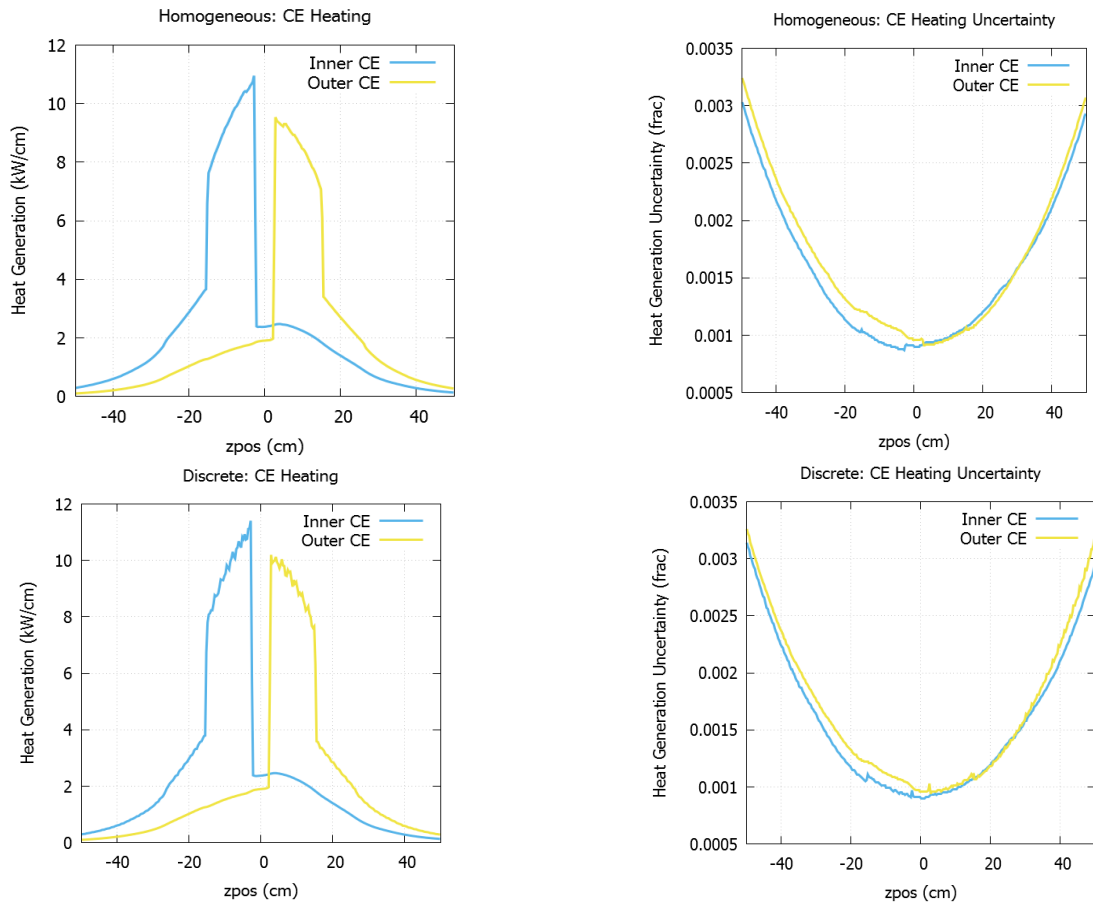


Figure A.10. BOC 449 Photon Heating (left column) and Uncertainty (right column) in Homogeneous (top row) and Explicit (bottom row) Models

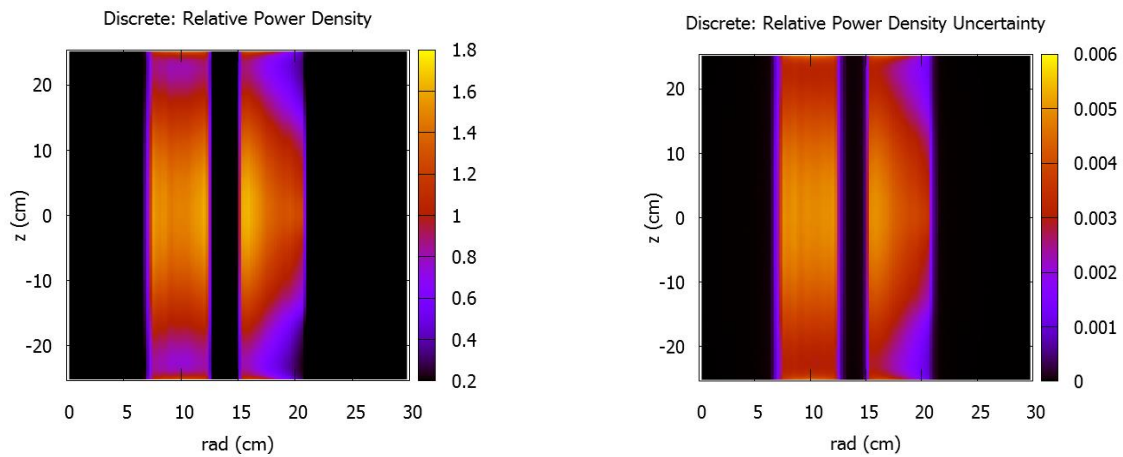


Figure A.11. BOC 400 Relative Fission Density (left) and Uncertainty (right) of Optimal CE Design

REFERENCES

1. Bierbaum, R. and M. Stults. "Adaptation to Climate Change: Context Matters." *Michigan Journal of Sustainability* 1 (Fall 2013).
2. Paris Agreement, FCCC/CP/2015/L.9/Rev.1. Paris Climate Change Conference (COP21), Twenty-first Session, November 30 to December 11, 2015, Paris, France.
3. Taylor, J. "The Real Climate Consensus: Nuclear Power." [Editorial]. *Forbes* (3 August 2017). Retrieved from <https://www.forbes.com/sites/jamestaylor/2017/08/03/the-real-climate-consensus-nuclear-power/#efe7fe62ef51>.
4. Asafu-Adjaye, J. et al. "An Ecomodernist Manifesto." (April 2015). Retrieved from <http://www.ecomodernism.org/manifesto-english/>.
5. "Examining the International Climate Change Negotiations: Hearing before the U.S. Senate Committee on Environment and Public Works," 114th Cong. (18 November 2015) (Testimony of David Waskow). Retrieved from http://www.wri.org/sites/default/files/Waskow_WRI_Senate_EPW_testimony_FINAL.PDF.
6. U.S. Department of Energy. "Staff Report to the Secretary on Electricity Markets and Reliability." (August 2017). Retrieved from <https://energy.gov/downloads/download-staff-report-secretary-electricity-markets-and-reliability>.
7. Government Accountability Office. "Commercial Nuclear Waste: Effects of a Termination of the Yucca Mountain Repository Program and Lessons Learned." GAO-11-229 (8 April 2011). Retrieved from <http://www.gao.gov/products/GAO-11-229>.
8. International Atomic Energy Agency. "Spent Fuel Reprocessing Options." IAEA-TECDOC-1587 (August 2008). Retrieved from http://www-pub.iaea.org/MTCD/publications/PDF/te_1587_web.pdf.

9. Ackerman, J.P. et al. "Treatment of Wastes in the IFR Fuel Cycle." *Progress in Nuclear Energy* 31 (1997) 141-154.
10. Weaver, K.D. et al. "A Once-Through Fuel Cycle for Fast Reactors." *Proceedings of the 17th International Conference on Nuclear Engineering*, Brussels, Belgium, 12-16 July 2009.
Retrieved from <http://terrapower.com/publications>.
11. Westinghouse Electric Company LLC. "AP1000 Overview." Presented at the Interregional Workshop on Advanced Nuclear Reactor Technology, Vienna, Austria, 4-8 July 2011. Retrieved from https://www.iaea.org/NuclearPower/Downloads/Technology/meetings/2011-Jul-4-8-ANRT-WS/2_USA_UK_AP1000_Westinghouse_Pfister.pdf.
12. Modro, S.M. et al. "Multi-Application Small Light Water Reactor Final Report." INEEL/EXT-04-01626, Idaho National Laboratory (December 2003). Retrieved from <https://indigitallibrary.inl.gov/sites/sti/sti/2546531.pdf>.
13. International Atomic Energy Agency. "Integral Design Concepts of Advanced Water Cooled Reactors." IAEA-TECDOC-977 (November 1997). Retrieved from http://www-pub.iaea.org/MTCDD/publications/PDF/te_977_prn.pdf.
14. Brown, N.R., Todosow, M., and Cuadra, A. "Screening of Advanced Cladding Materials and UN-U₃Si₅ Fuel." *Journal of Nuclear Materials* 462 (2015) 26-42.
15. Petrovic, B. et al. "The Pioneering Role of IRIS in Resurgence of Small Modular Reactors (SMRs)." *Nuclear Technology* 178 (2012) 126-152.
16. "Southern Nuclear Takes Over from Bankrupt Westinghouse to Finish Vogtle Nuclear Plant." *Chattanooga Times Free Press* (29 July 2017). Retrieved from <http://www.timesfreepress.com/news/business/aroundregion/story/2017/jul/29/southern-nuclear-takes-over-bankrupt-westingh/440893/>.

17. Sease, J.D. "Fabrication of Control Rods for the High Flux Isotope Reactor." ORNL/TM-9365/R1, Oak Ridge National Laboratory (March 1998).
18. Terrani, K. et al. "Demonstration of Advanced Manufacturing Techniques for Production of Nuclear Core Structures: Ultrasonic Additive Manufacturing of Hybrid Structures Resembling HFIR Control Plates." *Transactions of the American Nuclear Society* (June 2015).
19. Ilas, G. et al. *Modeling and Simulations for the High Flux Isotope Reactor Cycle 400*. ORNL/TM-2015/36, Oak Ridge National Laboratory (March 2015).
20. Oak Ridge National Laboratory. "High Flux Isotope Reactor (HFIR) User Guide." (November 2015). Retrieved from <https://neutrons.ornl.gov/>.
21. Korea Atomic Energy Research Institute. "ENDFPLOT-2.0." Retrieved from <http://atom.kaeri.re.kr:8080/cgi-bin/endfplot.pl>.
22. Betzler, B.R. "High Flux Isotope Reactor Control Element Operating Conditions." ORNL/TM-2010/168, Oak Ridge National Laboratory (August 2010).
23. Coin News Media Group LLC. "US Inflation Calculator." (2008). Retrieved from <http://www.usinflationcalculator.com/>.
24. Siemens AG. "Siemens Sets Milestone with First 3D-Printed Part Operating in Nuclear Power Plant." [Press release]. (9 March 2017). Retrieved from <https://www.siemens.com/press/en/pressrelease/?press=/en/pressrelease/2017/powergenerationservices/pr2017030221psen.htm>.
25. Friel, R.J. and R.A. Harris. "Ultrasonic Additive Manufacturing – A Hybrid Production Process for Novel Functional Products." *Procedia CIRP* 6 (2013) 35-40.
26. Schick, D.E. et al. "Microstructural Characterization of Bonding Interfaces in Aluminum 3003 Blocks Fabricated by Ultrasonic Additive Manufacturing." *Welding Journal* 89 (2010) 105-115.

27. Dehoff, R.R. and S.S. Babu. "Characterization of Interfacial Microstructures in 3003 Aluminum Alloy Blocks Fabricated by Ultrasonic Additive Manufacturing." *Acta Materialia* 58 (2010) 4305-4315.
28. Fuji, H.T. et al. "High-Strain-Rate Deformation in Ultrasonic Additive Manufacturing." *Scripta Materialia* 135 (2017) 125-129.
29. Gussev, M.N. et al. "Effect of Post Weld Heat Treatment on the 6061 Aluminum Alloy Produced by Ultrasonic Additive Manufacturing." *Materials Science and Engineering A* 684 (2017) 606-616.
30. Wolcott, P.J. et al. "Process Improvements and Characterization of Ultrasonic Additive Manufactured Structures." *Journal of Materials Processing Technology* 233 (2016) 44-52.
31. United States Nuclear Regulatory Commission. "General Design Criteria for Nuclear Power Plants." 10 CFR Part 50 Appendix A.
32. Hehr, A. et al. "Five-Axis Ultrasonic Additive Manufacturing for Nuclear Component Manufacture." *JOM* (2 December 2016), 10.1007/s11837-016-2205-6.
33. X-5 Monte Carlo Team, "MCNP – A General Monte Carlo N-Particle Transport Code, Version 5." LA-UR-03-1987, Los Alamos National Laboratory (2003).
34. Hosseini, S.A. "A New Random Distribution method to Simulate PBM Type Reactors with MCNP." *Annals of Nuclear Energy* 50 (2012) 215-219.
35. Kim, H. et al. "A New Strategy to Simulate a Random Geometry in a Pebble-Bed Core with the Monte Carlo Code MCNP." *Annals of Nuclear Energy* 38 (2011) 1877-1883.
36. Burns, J.R. et al. "Neutronic Modeling of Additively Manufactured HFIR Control Elements." *Transactions of the American Nuclear Society* (November 2015).

37. Burns, J.R. et al. "A Comparison of Modeling Strategies for Additively Manufactured HFIR Control Elements." *Transactions of the American Nuclear Society* (June 2016).
38. Tsoufanidis, N. and S. Landsberger. *Measurement and Detection of Radiation* (3rd ed.), p. 43. CRC Press Taylor and Francis Group (2011).
39. Tzika, F. and I.E. Stamatelatos. "Thermal Neutron Self-Shielding Correction Factors for Large Sample Instrumental Neutron Activation Analysis Using the MCNP Code." *Nuclear Instruments and Methods in Physics Research B* 213 (2004) 177-181.
40. Trkov, A. et al. "On the Self-Shielding Factors in Neutron Activation Analysis." *Nuclear Instruments and Methods in Physics Research A* 610 (2009) 553-565.
41. Liu, Y.H. et al. "Monte Carlo Determined Self-Shielded Groupwise Cross-Sections for the Activation Foil Stack Applied in the Epithermal Neutron Spectrum Adjustment." *Nuclear Instruments and Methods in Physics Research A* 602 (2009) 557-563.
42. Duderstadt, J.J. and L.J. Hamilton. *Nuclear Reactor Analysis*, pp. 606-609. John Wiley & Sons (1976).
43. Sjoden, G.E. *Foundations in Applied Nuclear Engineering Analysis*, p. 83. World Scientific Publishing Co. Pte. Ltd. (2009).
44. "Scale: A Comprehensive Modeling and Simulation Suite for Nuclear Safety Analysis and Design." ORNL/TM-2005/39, Version 6.1, June 2011. Available from Radiation Safety Information Computational Center at Oak Ridge National Laboratory as CCC-785.
45. Kotlyar, D. and E. Shwageraus. "On the Use of Predictor-Corrector Method for Coupled Monte Carlo Burnup Codes." *Annals of Nuclear Energy* 58 (2013) 228-237.
46. Kotlyar, D. et al. "A Perturbation-Based Substep Method for Couple Depletion Monte Carlo Codes." *Annals of Nuclear Energy* 102 (2017) 236-244.

47. Kotlyar, D. and E. Shwageraus. "Stochastic Semi-Implicit Substep Method for Coupled Depletion Monte-Carlo Codes." *Annals of Nuclear Energy* 92 (2016) 52-60.
48. Haeck, W. et al. "VESTA User's Manual – Version 2.1.0." DSU/SEC/T/2011-81 – Index A, Institut de Radioprotection et de Surete Nucleaire (2012).
49. Chandler, D. "Control Cylinder and Safety Plate Heat Generation, Thickness Variation, and Reduced Tantalum Loading Studies." C-HFIR-2014-002, Oak Ridge National Laboratory (July 2014).
50. Burns, J.R. et al. "Depletion Analysis of Additively Manufactured Control Elements in HFIR." *Transactions of the American Nuclear Society* (November 2016).
51. Primm, R.T. "Reactor Physics Studies of Reduced-Tantalum-Content Control and Safety Elements for the High Flux Isotope Reactor." ORNL/TM-2003-65, Oak Ridge National Laboratory (December 2003).
52. Singer, S. and J. Nelder. "Nelder-Mead Algorithm." *Scholarpedia* 4 (7) 2928.
53. Box, G.E.P. et al. *Statistics for Experimenters*, pp. 510-537. New York: John Wiley & Sons (1978).
54. "Response Surface Methodology Handbook for Nuclear Reactor Safety," EUR-9600 EN, Commission of the European Communities (1984).
55. Kim, U. and P. Seong. "Optimization of the Worth Shape of Axially Variable Strength Control Rods with Simulation Optimization Methodology for the Power Maneuvering of Pressurized Water Reactors." *Nuclear Engineering and Design* 225 (2003) 27-35.
56. Neddermeijer, H.G. et al. "A Framework for Response Surface Methodology for Simulation Optimization." *Proceedings – Winter Simulation Conference* (February 2000).

57. Anderson, V.L. and R.A. McLean. *Design of Experiments: A Realistic Approach*, pp. 225-278. Marcel Dekker, Inc. (1974).
58. Mathworks. "Multiple Linear Regression." MATLAB R2017b Documentation (2017). Retrieved from <https://www.mathworks.com/help/stats/regress.html>.
59. Stewart, J. *Multivariable Calculus: Concepts and Contexts*, pp.813-814. Belmont, CA: Thomson Brooks/Cole (2005).