

Methodology for Coupled Reactor Physics, Thermal Hydraulics, and Depletion
of a Fluoride Salt-Cooled High-Temperature Reactor (FHR)

A Dissertation
Presented to
The Academic Faculty

By

Kyle Michael Ramey

In Partial Fulfillment
Of the Requirements for the Degree
Doctor of Philosophy in Nuclear and Radiological Engineering

Georgia Institute of Technology

August 2021

Copyright © Kyle M. Ramey 2021

Methodology for Coupled Reactor Physics, Thermal Hydraulics, and Depletion
of a Fluoride Salt-Cooled High-Temperature Reactor (FHR)

Approved by:

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

Dr. Dan Kotlyar
Nuclear and Radiological Engineering Program
G. W. Woodruff School of Mechanical Engineering
Georgia Institute of Technology

Dr. Davor Grgić
Department of Energy and Power Systems
University of Zagreb, Croatia

Dr. Dingkang Zhang
Nuclear and Radiological Engineering Program
G. W. Woodruff School of Mechanical Engineering
Georgia Institute of Technology

Dr. G. Ivan Maldonado
Department of Nuclear Engineering
Tickle College of Engineering
University of Tennessee

Date Approved: April 26, 2021

TABLE OF CONTENTS

LIST OF TABLES	viii
LIST OF FIGURES	xi
LIST OF ABBREVIATIONS	xix
SUMMARY	xxi
CHAPTER 1 THE FHR-AHTR SYSTEM.....	1
1.1 Historical Background	3
1.2 3D Reference Core Geometry.....	4
1.2.1 Fuel Assembly Geometry and Dimensions.....	5
1.2.1.a TRISO Fuel Particles.....	6
1.2.1.b Fuel Plank.....	7
1.2.1.c Fuel Assembly	8
1.2.2 Reactivity Control.....	10
1.2.2.a Burnable Poison Spheres	10
1.2.2.b Control Blades.....	11
1.2.3 Materials and Compositions	11
1.3 Previous Studies	13
1.3.1 Previous Work Conducted on Neutronics.....	14
1.3.2 Previous Work Conducted on Thermal Hydraulics	16
1.3.3 Previous Work Conducted on Thermal Expansion.....	17
1.3.3.a Geometric Thermal Expansion of TRISO and TRISO Lattice.....	17
1.3.3.b Density Changes in Graphite Components	19
CHAPTER 2 MOTIVATION AND OBJECTIVES.....	21
2.1 Motivation.....	21
2.2 Objectives	22
2.3 Scope.....	24
2.4 Approach.....	25
2.5 Overview of Dissertation	28
CHAPTER 3 THERMAL EXPANSION METHODOLOGY	30
3.1 Assumptions and Justifications.....	30
3.1.1 Free Expansion	30
3.1.2 Integrity of Shapes	31
3.1.3 Priority in Expanding Components.....	31
3.1.4 Axial Expansion.....	32
3.1.5 Radial Expansion	32
3.1.6 Conservation of Media.....	33
3.1.7 Fuel Assembly Pitch	34
3.1.8 Components Radially Beyond the Active Core	35
3.2 Treatment of Particles	36
3.2.1 TRISO Particles	36
3.2.2 Europium Oxide Spheres (Burnable Poison).....	38
3.3 Axial and Radial Treatment for All Remaining Regions.....	39

3.3.1 TRISO Fuel Stripe Matrix	39
3.3.2 Fuel Planks.....	40
3.3.3 Fuel Plank Spacers.....	41
3.3.4 Control Blade.....	43
3.3.5 Structural Y-Shape.....	43
3.3.6 Assembly Pitch	44
3.4 Extra-Assembly Components.....	45
3.4.1 Radial Layers Just Beyond the Core.....	45
3.4.2 Radial Layers Beyond the Coolant Downcomer.....	46
3.4.3 Lower and Upper Axial Reflectors	47
3.4.4 Bottom and Top Support Plates	47
CHAPTER 4 HEAT TRANSFER AND THERMAL HYDRAULICS METHODOLOGY	49
4.1 Fuel Stripe Homogenization	51
4.2 1D Heat Transfer Across Fuel Plate and Coolant Channel.....	52
4.2.1 Convection to Coolant	54
4.2.2 Conduction through Graphite Plank Sleeve.....	56
4.2.3 Conduction through Fuel Stripes	56
4.2.4 Central Graphite Region (Adiabatic Boundary)	57
4.2.5 Dependence on the Number of TRISO Layers	57
4.2.5.a Reactivity Impact.....	57
4.2.5.b Thermal Impact on Fuel Plate Temperature Distribution.....	58
4.2.6 Dependence on Axial Location.....	61
4.3 Neutronic Impacts of Temperature Variation	62
4.3.1 Use of the Average Fuel Temperature within Fuel Stripes.....	62
4.3.2 TRISO-Wise Fuel Temperature Impact on Power Distribution	64
4.4 Temperature-Dependence of Fuel Thermal Conductivity	67
4.5 Impact of Neutron Fluence on Graphite Thermophysical Properties.....	68
4.5.1 Volumetric Change	69
4.5.2 Thermal Expansion Coefficient	71
4.5.3 Thermal Conductivity	72
CHAPTER 5 TRIANGULAR MESH TALLYING IN SERPENT.....	76
5.1 Current Tallying Efficiencies in Serpent.....	77
5.2 Introduction of Triangular Mesh Tally	80
5.3 Implementation of Triangular Mesh Tally.....	81
5.3.1 Changes Made to the Serpent Source Code.....	81
5.3.2 Division Scheme for Triangular Bins	83
5.3.3 Reporting of the Triangular Mesh Bin Location.....	84
5.4 Results and Testing from Using Triangular Mesh Tally.....	85
5.4.1 Interpreting the Output of the Triangular Mesh Tally	85
5.4.2 Radial Testing of the Triangular Mesh Tally.....	87
5.4.2.a Even Lattice Size.....	88
5.4.2.b Odd Lattice Size.....	90
5.4.3 Axial Testing of the Triangular Mesh Tally	90
5.4.4 Testing Compatibility with the 3D AHTR Serpent Model.....	91
CHAPTER 6 NEUTRONICS METHODOLOGY AND CORE REACTIVITY CONTROL	93
6.1 Reactivity Sensitivity Studies	94
6.1.1 Control Blade Reactivity Impact	94
6.1.2 Control Blade Axial Power Impact.....	97

6.1.3 Depletion Using Various Burnable Poison Densities	99
6.2 Control Blade Movement Schedule for Core reactivity Control.....	103
6.2.1 Reactivity Control by Using Full Density Burnable Poison Europia Spheres.....	103
6.2.2 Reactivity Control by Using Quarter Density Burnable Poison Europia Spheres	109
6.3 Automated CB Selection and Insertion Schedule	109
6.4 Automated Control Blade Withdrawal Description.....	112
6.4.1 Details and Testing of Automated Control Blade Withdrawal Capability	113
6.4.2 Reaching Critical Core Configuration through Control Blade Movement	116
CHAPTER 7 USING THE C++ SCRIPT ATOMICS.....	118
7.1 Modifying How the Script is Executed via the User Input File.....	118
7.1.1 Selecting the Module Type for the Script	119
7.1.2 Controlling Geometry Features.....	121
7.1.2.a Fuel Kernel.....	122
7.1.2.b Fuel Lattice.....	122
7.1.2.c Burnable Poison Lattice	123
7.1.2.d Planks and Assembly	124
7.1.2.e Reflector Assemblies.....	124
7.1.2.f Axial Discretization	125
7.1.2.g Beyond the Active Core Region.....	126
7.1.3 Providing Material Parameters	126
7.1.3.a Setting Uniform or Region-Specific Materials.....	127
7.1.3.b Material Densities	127
7.1.3.c Linear Thermal Expansion Coefficients.....	128
7.1.3.d Thermal Conductivities	128
7.1.4 Providing Thermal Hydraulic Parameters	129
7.1.4.a Core Power and Flow Properties	129
7.1.4.b FLiBe Properties	130
7.1.5 Depletion Simulation Options.....	131
7.2 Implementation of the Methodology in ATOMICS.....	133
7.2.1 Reading of Input Files.....	134
7.2.1.a Main Serpent File	134
7.2.1.b User Input File.....	135
7.2.1.c Spatial Mapping Files.....	136
7.2.1.d Power Tallies.....	136
7.2.1.e Burnup Schedule	136
7.2.1.f CB Schedule	138
7.2.1.g Eigenvalue.....	138
7.2.1.h Iteration History	138
7.2.1.i Neutron Fluence.....	139
7.2.2 Code Execution.....	139
7.2.3 Writing of Output Files.....	140
7.2.3.a Iteration History	140
7.2.3.b CB Insertion Schedule.....	142
7.2.3.c CB Withdrawal Schedule	142
7.2.3.d Axial Power Profile.....	143
7.2.3.e Hot Channel Thermal Hydraulic Profile	143
7.2.3.f 2D Radial Power Profile	144
7.2.3.g Materials.....	144
7.2.3.h Geometry.....	145
7.3 Using ATOMICS to Run Simulations	146

7.3.1 Updating Installation of Serpent	146
7.3.2 Compiling ATOMICS	146
7.3.3 Running ATOMICS in Linux Environment	147
CHAPTER 8 RESULTS	149
8.1 Geometric Sensitivity Studies	150
8.1.1 Fuel Plank Width	150
8.1.2 Distance from Fuel Stripe to Coolant	152
8.1.3 Height of Axial Regions beyond the Active Core	154
8.1.4 Number of Burnable Poison Columns	158
8.1.5 Burnable Poison Column Pitch	160
8.2 Thermal Sensitivity Studies	160
8.2.1 Coolant Inlet Temperature	161
8.2.2 Core Power	163
8.2.3 Graphite Thermal Conductivity	166
8.3 Depletion Simulations	169
8.3.1 Numerical Stability in 3D Depletion Simulations	171
8.3.2 Single Tracked Fuel and Burnable Poison Material	177
8.3.2.a Explicit Xenon Treatment	177
8.3.2.b Equilibrium Xenon Treatment	179
8.3.2.c Xenon Treatment Comparison	180
8.3.3 One-Third Assembly and One-Sixteenth Axial (Fine) Tracking	182
8.3.3.a Explicit Xenon Treatment	182
8.3.3.b Impact of Number of Simulated Particles	186
8.3.3.c Equilibrium Xenon Treatment	188
8.3.3.d Xenon Treatment Comparison	189
8.3.3.e Comparison to Single Material Tracking	191
8.3.4 Fine Spatial tracking with Multiphysics	192
8.3.4.a Number of Thermal Hydraulic Iterations between Burnup Steps	192
8.3.4.b Explicit Xenon Treatment	198
8.3.4.c Equilibrium Xenon Treatment	201
8.3.4.d Xenon Treatment Comparison	203
8.3.5 Fine Spatial Tracking with Criticality Control via Control Blade Movement	207
8.3.5.a Explicit Xenon Treatment	208
8.3.5.b Equilibrium Xenon Treatment	213
8.3.5.c Comparison between Xenon Treatments	217
8.3.6 Fine Tracking with Both Multiphysics and Criticality Control via Control Blade Movement	219
8.3.6.a Impact of Number of Simulated Particles	220
8.3.6.b Depletion using Explicit Xenon Treatment	221
8.3.7 Runtime Comparison of Depletion Simulations	224
8.3.8 Spatial Power Performance of Depletion Simulations	227
CHAPTER 9 CONCLUSION	231
9.1 Review of Multiphysics Scripting Efforts	231
9.1.1 Thermal Expansion	231
9.1.2 Heat Transfer and Thermal Hydraulics	232
9.1.3 Neutronics	233
9.2 Use of ATOMICS	233
9.3 Summary of Results	235
9.4 Future Work	237

9.4.1 Further Scripting Efforts	237
9.4.2 Verification and Validation of Results	238
9.4.3 Additional Modeling Capabilities.....	239
APPENDIX A DEVELOPMENT OF GRAPHITE THERMAL CONDUCTIVITY BEHAVIOR	240
APPENDIX B MODIFICATIONS TO SERPENT SOURCE CODE FOR TRIANGULAR MESH	245
B.1 Recognition of Reserved and Code-Specific Keywords	245
B.2 Initialization of New Tally from Input	246
B.3 Creation of Tallying Structure for Recording Scores	248
B.4 Bin Index Searching to Find Where a Score Should be Recorded	250
B.5 Modifying Writing of Results to Detector-Specific Output File	254
B.6 Minor Modifications due to Argument Changes in Functional Calls	257
APPENDIX C MAPPING FROM TRIANGULAR MESH TALLY BINNING TO 3D AHTR SERPENT MODEL LOCATION NOTATION	258
APPENDIX D SUPPLEMENTARY RADIAL POWER PROFILES	272
D.1 Full Density Burnable Poison Loading Control Blade Insertion	272
D.2 Quarter Density Burnable Poison Loading Control Blade Insertion.....	277
D.3 Quarter Density Burnable Poison Loading Control Blade Withdrawal	294
APPENDIX E EXAMPLE INPUT FILE	297
APPENDIX F SCRIPTS FOR WORK CONDUCTED	301
F.1 Hex Map Input File.....	301
F.2 ATOMICS C++ Source Code.....	306
REFERENCES	307

LIST OF TABLES

Table 1.1. Dimensions and compositions of major components of the AHTR core.....	5
Table 1.2. TRISO particle layers and dimensions.....	7
Table 1.3. Parameters and dimensions used to construct the AHTR fuel assembly.	9
Table 1.4. List of materials and their compositions using in the reference AHTR model.....	12
Table 1.5. Reactivity impact of TRISO lattice geometric thermal expansion with no reactivity control. Differences have a statistical uncertainty of 7 pcm.....	18
Table 1.6. Reactivity impact of TRISO lattice geometric thermal expansion with burnable poison spheres. Differences have a statistical uncertainty of 7 pcm.	19
Table 1.7. Reactivity impact of TRISO lattice thermal expansion with both BP spheres and control blade. Differences have a statistical uncertainty of 7 pcm.	19
Table 1.8. Material-wise graphite density sensitivity [24].....	20
Table 3.1. Parameters of AHTR model (at 293 K).	30
Table 3.2. Dimensions and properties of TRISO particle layers.	36
Table 4.1. Parameters used to thermally homogenize TRISO particles in fuel stripe for fuel assembly under core-average operating conditions with conservative fuel and graphite thermal conductivities.	52
Table 4.2. Coolant volume fraction (as percent) by region for cold reference design.	54
Table 4.3. Parameters of FLiBe coolant channel.	56
Table 4.4. Eigenvalues of cases varying the number of TRISO layers (while conserving total fuel particle loading).....	58
Table 4.5. Average and maximum temperatures for both homogenized fuel stripe and explicit fuel kernels. All cases assume core-average specific power and average coolant density and temperature. The cases using four TRISO fuel layers are highlighted as they correspond to the reference design used in this dissertation.	61
Table 4.6. Layer-wise and average fuel kernel temperatures for average and double power density.	65
Table 5.1. Tallying efficiency of fission power in Serpent for the 3D AHTR core.....	79
Table 5.2. List of Serpent source code files modified to create the triangular mesh tally.	83
Table 5.3. Tally comparison between using a cell-based physical tally and a superimposed triangular mesh tally for hexx type geometry test case using an even lattice size.	89

Table 5.4. Tally comparison between using a cell-based physical tally and a superimposed triangular mesh tally for hexx type geometry test case also using axial partitioning.	91
Table 6.1. Normalized power (via fission energy deposition) generated by the assembly groups indicated in Figure 6.1 from each of the first eight radial rings in the AHTR core.	95
Table 6.2. Isotopic data for europium [43].	101
Table 6.3. Eigenvalue and radial power peaking factors for the four considered CB insertion schemes.	106
Table 7.1. Burnup step types used within Serpent [46].	137
Table 8.1. Estimates of reported statistical uncertainty for various power shaping parameters for an example simulation with an eigenvalue relative statistical uncertainty of 10 pcm.	149
Table 8.2. Five fuel plank width cases with associated eigenvalues.	152
Table 8.3. Case summary and eigenvalue results for varying the plank sleeve thickness.	153
Table 8.4. Axial reflector and support plate thickness eigenvalue differences [†] [pcm].	155
Table 8.5. Axial reflector and support plate thickness axial PPF.	157
Table 8.6. Axial reflector and support plate thickness axial PPF relative change [%].	157
Table 8.7. Eigenvalue impact from using a different number of BP columns.	159
Table 8.8. Eigenvalue impact from adjusting BP column pitch.	160
Table 8.9. Summary of results for varying inlet coolant temperature.	161
Table 8.10. Summary of results for varying core power for all CBs withdrawn.	164
Table 8.11. Summary of results for varying core power for critical CB insertion.	164
Table 8.12. Summary of results for varying graphite thermal conductivity.	167
Table 8.13. New Serpent time integration depletion options introduced with version 2.32 [59].	176
Table 8.14. Depletion history for a single tracked fuel material with explicit xenon treatment.	178
Table 8.15. Depletion history for single tracked material with equilibrium xenon treatment.	179
Table 8.16. Comparison of xenon treatments for a single tracked material.	180
Table 8.17. Depletion sequence summary for fine tracking with explicit xenon treatment.	183
Table 8.18. Depletion sequence summary for fine tracking with equilibrium xenon treatment.	188
Table 8.19. Comparison of xenon treatments for fine spatial tracking.	189

Table 8.20. Comparison between single and fine material tracking cases for explicit xenon usage.	191
Table 8.21. Eigenvalue results of different xenon treatments when using one thermal hydraulic substep.....	204
Table 8.22. Relative (to explicit xenon treatment) average differences summary for various spatial power parameters when using one thermal hydraulic substep.	206
Table 8.23. Depletion sequence summary using CB movement and explicit xenon treatment...	209
Table 8.24. Depletion sequence summary using CB movement and equilibrium xenon treatment.	213
Table 8.25. Example statistical uncertainties for cases using different neutron generation sizes.	220
Table 8.26. Depletion sequence summary for using PC, TH feedback, criticality iteration via CB movement, and explicit xenon treatment.....	222
Table 8.27. Runtime summary for depletion cases considered in this dissertation.	226
Supplementary Table A.1. List of inferred actual values from reference [36] and those obtained from the function fit of Equation A.9.	244
Supplementary Table C.1. One-to-one mapping of results from the 3D AHTR SERPENT model using both a physical tally and the “dht” superimposed triangular mesh tally. Note that the “tri” index corresponds to the sum of two adjacent triangular regions as discussed above.	260
Supplementary Table D.1. Complete results to accompany the CB insertion analysis for Scheme 4. Rows highlighted in gray indicate one-sixth radially symmetric assembly group pairs.	272
Supplementary Table D.2. Radial power profiles for Scheme 4.....	273
Supplementary Table D.3. Complete results to accompany the quarter density europia CB insertion scheme.....	278
Supplementary Table D.4. Radial power profiles for the quarter density europia automated CB insertion scheme.....	280
Supplementary Table D.5. Complete results to accompany the quarter density europia CB withdrawal schedule.....	294
Supplementary Table F.1. Hexagon mapping to recreate a 19 x 19 array from ATOMICS phase space.....	301

LIST OF FIGURES

Figure 1.1. AHTR full core with main features [2] [3].	1
Figure 1.2. Hexagonal AHTR assembly and radial core layout.	2
Figure 1.3. 3D depiction of the AHTR core [2].	5
Figure 1.4. Left: TRISO fuel particle with all layers identified [3]. Right: AGR-2 TRISO particle modeled in SERPENT [11] with surrounding matrix material (teal) with average cubic pitch.	6
Figure 1.5. AHTR fuel plank with two fuel stripes.	8
Figure 1.6. AHTR fuel assembly with key features identified.	9
Figure 1.7. Control blade (dark green) inserted in the central location of an AHTR fuel assembly.	11
Figure 1.8. k_{∞} difference between lattice and five unique random heterogeneous dispersions of TRISO particles [17].	15
Figure 1.9. k_{∞} difference versus nine depletion subdivisions per fuel stripe [17].	16
Figure 2.1. Conceptual flowchart of how the multiphysics effects will be coupled and implemented.	26
Figure 2.2. Radial assembly layout of 252 fuel assemblies. Assemblies are shown to be in 84 groups of three for corresponding one-third symmetric locations. Note that the geometry shown was obtained by rotating the reference geometry (shown in Figure 1.2) clockwise by 30° for ease of viewing.	27
Figure 3.1. Visualization of different fuel assembly radial pitch treatments in axial sections. Deviations away from vertical in b. and c. are exaggerated for visual effect and would be relatively less in an actual implementation.	35
Figure 3.2. TRISO particle used in AHTR.	36
Figure 3.3. Fuel plank (red) embedded in structural graphite (gray).	41
Figure 3.4. Two sizes of plank spacer are used in the model, depending on fabricated channel width.	42
Figure 3.5. Structural Y-shape (gray). Fuel planks not shown as embedded to better highlight this structural feature.	44
Figure 3.6. Portions of removable reflector assemblies, permanent reflector region, boron carbide layer, and core barrel.	46
Figure 3.7. Downcomer, pressure vessel inner liner, and pressure vessel wall.	46

Figure 4.1 AHTR half-plank channel with individual features identified.	50
Figure 4.2. Eigenvalue behavior with variable TRISO layer configuration. Left: uncontrolled. Right: controlled with europa BP spheres at the center of the plank.	58
Figure 4.3. Temperature distribution for homogenized fuel stripe depending on the number of TRISO layers. Assumes average core power and coolant temperature.....	59
Figure 4.4. Power distribution (a) and decomposed power distribution (b). [33].....	60
Figure 4.5. Temperature distribution for heterogeneous model through the center of TRISO particles based on the number of TRISO layers. Assumes average core power and coolant temperature.	60
Figure 4.6. Average plank temperature distribution by axial location for four TRISO layers.	62
Figure 4.7. Fuel temperature coefficient for a reflected assembly.....	63
Figure 4.8. Temperature distribution across coolant channel and fuel plate for average power density (blue) and double average power density (orange) of a fuel stripe with four layers.	64
Figure 4.9. Relative differences from Explicit Layer Temperature Average Power case to Average Layer Temperature Average Power case. Maximum: 0.741%. Minimum: -0.673%. 66	66
Figure 4.10. Relative differences from Average Layer Temperature Double Power case to Average Layer Temperature Average Power case. Maximum: 0.679%. Minimum: -0.694%. 67	67
Figure 4.11. Dependence of fuel thermal conductivity on fuel temperature. [3].....	68
Figure 4.12. Volume change versus neutron fluence for specimens irradiated at different temperatures. Each data point is the volume change of an individual specimen. The lines were fitted to the data using a second-order polynomial [36].....	71
Figure 4.13. Normalized change of mean thermal expansion coefficient (referenced to 25 °C) measured at irradiation temperatures, plotted versus neutron fluence [36].....	72
Figure 4.14. Measured thermal conductivity versus measurement temperature for as-received and specimens irradiated at 459 ± 37 °C. The vertical dashed line indicates the irradiation temperature. The labels indicate the specimen orientation and total neutron fluence ($\times 10^{25}$ n/m ² [E > 0.1MeV]) [36].	73
Figure 4.15. Heatmap of the relative change in thermal conductivity of grade G347A graphite when varying both temperature and neutron fluence (E > 0.1 MeV). Change is relative to value for nonirradiated graphite at 459 °C.	74
Figure 5.1. Dependence of transport runtime on the number of material tally bins in Serpent.	78
Figure 5.2. Orientations of the types of hexagons used in SERPENT.....	81
Figure 5.3. Hexagons divided into six equilateral triangles.....	81

Figure 5.4. Partitioning scheme used to divide “dht” tally bins.....	84
Figure 5.5. Geometry and centroid (point O) of an equilateral triangle [41].....	84
Figure 5.6. “dht” tally results from SERPENT detector output file.....	86
Figure 5.7. Test geometries for even lattice size.....	89
Figure 5.8. Test geometries for odd lattice size.	90
Figure 6.1. Eight assembly groups being assessed for control blade insertion.....	95
Figure 6.2. Control blade bank worth of a three-assembly group within radial rings of the first eight rings.....	96
Figure 6.3. Integral rod worth of a representative three-assembly group from each of the first eight radial rings of assemblies of the active core.....	97
Figure 6.4. AO caused by one-eighth active core height increment insertion of control blade groups from specified radial rings.....	98
Figure 6.5. Axial power peaking profiles for one-eighth active core height increment insertion of control blade group in the third radial ring (assembly group 8).....	99
Figure 6.6. Change in eigenvalue with respect to burnup for various europia densities. Densities considered are 0.00, 1.25, 2.50, and 5.00 g/cm ³ , respectively.....	100
Figure 6.7. Eigenvalue differences due to the number of tracked material zones for both BP and fuel. All cases use full density europia (5.00 g/cm ³)......	102
Figure 6.8. Radial power profile for full density europia BP and no CBs inserted.	104
Figure 6.9. Four control blade insertion schemes considered for full density europia.	105
Figure 6.10. Radial power profiles for the four considered CB insertion schemes.	107
Figure 6.11. Effective eigenvalue impact from successively inserting CB groups.....	108
Figure 6.12. Radial section- and assembly-wise PPFs.....	108
Figure 6.13. Conceptual flowchart of how automated CB insertion search process functions....	110
Figure 6.14. Effective eigenvalue impact from automated insertion of CB groups.....	112
Figure 6.15. Radial section- and assembly-wise PPFs for automated insertion.	112
Figure 6.16. Effective eigenvalue impact from automated withdrawal of CB groups.....	114
Figure 6.17. Radial section- and assembly-wise PPFs for automated withdrawal.	115
Figure 7.1. Input file portion pertaining to selecting modes for running the script.	119

Figure 7.2. Input file portion pertaining to geometric features, options, and dimensions.	121
Figure 7.3. Input file portion pertaining to materials options and properties.....	126
Figure 7.4. Input file portion pertaining to thermal hydraulic options and properties.	129
Figure 7.5. Input file portion pertaining to depletion options and criticality convergence.	131
Figure 7.6. A visual depiction of how ATOMICS works and a summary of important I/O files.	134
Figure 7.7. Example <i>input.txt</i> file.	135
Figure 7.8. Example <i>dep.txt</i> input file.....	137
Figure 7.9. Example <i>iteration.txt</i> file.....	141
Figure 7.10. Order in which ATOMICS subroutines are executed.....	145
Figure 7.11. Example of how to compile ATOMICS in a Linux environment.	147
Figure 7.12. Example bash shell for-loop used for running ATOMICS.....	148
Figure 8.1. Images for assembly sections with five different plank thicknesses.	151
Figure 8.2. Eigenvalue impact due to changing fuel plank thickness.....	152
Figure 8.3. Varying separation of the fuel stripe from the coolant.	153
Figure 8.4. Eigenvalue impact of varying the sleeve thickness relative to the reference (0.1 cm).	154
Figure 8.5. Normalized axial power profiles for sixteen axial cases considered. The first index in the legend refers to the axial reflector thickness [cm] while the second index refers to the axial support plate thickness [cm].	156
Figure 8.6. Left: Axial offset values for groupings of same axial support plate (ASP) thickness. Right: Axial offset values for groupings of same axial reflector (AR) thickness. Calculated statistical uncertainties for all cases are below 0.007%.	158
Figure 8.7. Fuel planks with different numbers of BP columns in the center.....	159
Figure 8.8. Eigenvalue difference from the reference case for adding BP columns.....	159
Figure 8.9. Eigenvalue impact of inlet coolant temperature for zero power cases.	162
Figure 8.10. Number of CB groups required for criticality when varying inlet coolant temperature for zero power cases.	163
Figure 8.11. Eigenvalue impact due to varying the core power (all CBs withdrawn).	164
Figure 8.12. Peak temperatures based on core power.....	165

Figure 8.13. PPFs resolved at multiple levels when varying thermal core power for no CBs inserted (N) and critical CB insertion (C).....	166
Figure 8.14. Reactivity impact of varying graphite thermal conductivity.....	168
Figure 8.15. Hot channel temperature distribution for different graphite thermal conductivities [W/(m·K)].....	169
Figure 8.16. Eigenvalue (left) and Shannon entropy (right) by cycle for AHTR transport.	171
Figure 8.17. Burnup schemes with and without using depletion substeps [57].....	175
Figure 8.18. Eigenvalue over depletion sequence for single material (core average) tracking with explicit xenon treatment.	178
Figure 8.19. Power results for single material (core average tracking) with explicit xenon treatment. Left: radial PPFs for both the assembly and section levels. Right: axial power profiles over cycle (legend given for each burnup step in units of MWd/kgHM).....	179
Figure 8.20. Eigenvalue over depletion sequence for single material (core average) tracking with equilibrium xenon treatment.	180
Figure 8.21. Reactivity differences over a depletion sequence using a single tracked material for both explicit and equilibrium xenon treatments.	181
Figure 8.22. Normalized axial power distribution over a depletion sequence using fine spatial tracking. Legend values are for cumulative burnup points in units of MWd/kgHM.	184
Figure 8.23. Axial PPF over depletion sequence using fine spatial tracking.....	185
Figure 8.24. AO over depletion sequence using fine spatial tracking.	185
Figure 8.25. Radial PPFs over depletion sequence using fine spatial tracking.....	186
Figure 8.26. Reactivity impact over a depletion sequence due to varying the number of active particles used. Reference used 1.25×10^8 particles whereas the two comparison cases used 10-times fewer (Blue) and 100-times fewer (Orange) particles.....	187
Figure 8.27. Eigenvalue over depletion sequence for fine spatial tracking.	190
Figure 8.28. PPFs over depletion sequence for both xenon treatments using fine spatial tracking.	190
Figure 8.29. Reactivity differences for different material tracking granularity (differences taken with respect to fine material tracking resolution).....	192
Figure 8.30. Eigenvalue over depletion sequence when using a difference number of thermal hydraulic iterations and equilibrium xenon treatment. Legend values are the number of thermal hydraulic iterations used between burnup steps.	193

Figure 8.31. Reactivity differences over depletion sequence due to using a different number of thermal hydraulic iterations when using equilibrium xenon treatment. Legend values are the number of thermal hydraulic iterations used between burnup steps, with differences taken relative to the reference case of four iterations.	194
Figure 8.32. Average of 4032 tracked regions' absolute change in local power divided by local power statistical uncertainty.	195
Figure 8.33. Normalized axial power profiles for thermal hydraulic feedback and zero iteration substeps. Legend corresponds to cumulative burnup amounts given in units of MWd/kgHM for the seventeen depletion steps considered.	196
Figure 8.34. Normalized axial power profiles for final four depletion steps. Legend corresponds to cumulative burnup amounts give in units of MWd/kgHM.	197
Figure 8.35. AO behavior over depletion sequence when using a different number of thermal hydraulic substeps. Cases used equilibrium xenon treatment but no predictor-corrector method.	198
Figure 8.36. Normalized axial power profiles over depletion sequence using one thermal hydraulic substep and explicit xenon treatment. Legend values correspond to the burnup step number.	199
Figure 8.37. PPF values for whole core (4032 regions), section (252), assembly (84), and axial levels (112) over the depletion sequence when using one thermal hydraulic substep, explicit xenon treatment, and predictor-corrector.	200
Figure 8.38. AO of the power profile at each depletion step using one thermal hydraulic substep and explicit xenon treatment.	201
Figure 8.39. Normalized axial power profiles over depletion sequence using one thermal hydraulic substep and equilibrium xenon treatment. Legend values correspond to the burnup step number.	202
Figure 8.40. PPF values for whole core, section, assembly, and axial levels over the depletion sequence when using one thermal hydraulic substep, equilibrium xenon treatment, and predictor-corrector.	202
Figure 8.41. AO of the power profile at each depletion step using one thermal hydraulic substep and equilibrium xenon treatment.	203
Figure 8.42. Reactivity differences between xenon treatments when using one thermal hydraulic substep.	205
Figure 8.43. Relative (to explicit xenon treatment) differences between xenon treatments for various spatial power parameters when using one thermal hydraulic substep.	205
Figure 8.44. AO of the power profile at each depletion step using one thermal hydraulic substep.	207

Figure 8.45. Eigenvalue over depletion sequence using criticality iteration via CB movement and explicit xenon treatment. Red lines correspond to the criticality tolerance for the criticality search (± 150 pcm).....	210
Figure 8.46. PPFs over depletion sequence using criticality iteration via CB movement and explicit xenon treatment.	211
Figure 8.47. Normalized axial power profiles over depletion sequence using criticality iteration via CB movement and explicit xenon treatment. Legend values are for each depletion step.	212
Figure 8.48. AO (left axis) over depletion sequence (solid blue) with criticality iteration substeps (dashed blue) via CB movement and explicit xenon treatment. Partial insertion fraction (right axis) of the last CB group (orange) is shown for reference for impact on the AO.	213
Figure 8.49. Eigenvalue over depletion sequence using criticality iteration via CB movement and equilibrium xenon treatment. Red lines correspond to the criticality tolerance for the criticality search (± 150 pcm).....	215
Figure 8.50. PPFs over depletion sequence using criticality iteration via CB movement and equilibrium xenon treatment.	215
Figure 8.51. Normalized axial power profiles over depletion sequence using criticality iteration via CB movement and equilibrium xenon treatment. Legend values are for each depletion step.	216
Figure 8.52. AO over depletion sequence (solid blue) using criticality iteration (dashed blue line) via CB movement and equilibrium xenon treatment. Partial insertion fraction of the last CB group (orange) is shown for reference for impact on the AO.....	216
Figure 8.53. Critical CB insertion for explicit and equilibrium xenon treatments.....	217
Figure 8.54. AO over depletion sequence using criticality iteration via CB movement for both explicit (orange) and equilibrium (blue) xenon treatments. Previous results from the depletion case with all CBs withdrawn (red) is shown for comparison.	219
Figure 8.55. Number of control blade groups inserted to achieve critical insertion over the depletion sequence for different neutron generation sizes.	221
Figure 8.56. PPFs over the depletion sequence when using thermal hydraulic feedback, criticality iteration via control blade movement, and explicit xenon treatment.	223
Figure 8.57. AO over depletion sequence (solid blue line) when using thermal hydraulic feedback and criticality iteration via control blade movement (dashed blue line) with explicit xenon treatment. CB partial insertion fraction (orange line) shown for reference for power shaping. AO for no TH feedback and no criticality iteration also shown for reference (black line).....	224
Figure 8.58. Whole Core PPFs for all depletion cases.....	228

Figure 8.59. One-third assembly section PPFs for all depletion cases.	228
Figure 8.60. Assembly PPFs for all depletion cases.	229
Figure 8.61. Axial PPFs for all depletion cases.	230
Figure 8.62. AO for all depletion cases.	230
Figure 9.1. A flowchart of how ATOMICS works and a summary of important I/O files.....	234
Supplementary Figure A.1. Measured thermal conductivity versus measurement temperature for as-received and specimens irradiated at 459 ± 37 °C. The vertical dashed line indicates the irradiation temperature. The labels indicate the specimen orientation and total neutron fluence ($\times 10^{25}$ n/m ² [E > 0.1MeV]) [36].	240
Supplementary Figure A.2. Fit of fluence-dependent function f_2 at 459 °C.	242
Supplementary Figure A.3. Fit of fluence-dependent function used to obtain f_3 at 1000 °C.	243
Supplementary Figure C.1. Physical tallying mapping scheme for the first and third columns of Supplementary Table C.1. The first number corresponds to the “Hex Group”, which is for one-third core symmetric assembly locations. The second number corresponds to the “Radial Group”, which loosely starts in the positive y-direction and rotates clockwise.....	258
Supplementary Figure C.2. Triangular element mapping scheme for the sixth column of Supplementary Table C.1. Note that triangular elements (black numbers) have been combined into one-third assembly sections (designated by orange numbers and divisions shown by thick black lines). The notation correspondence is: tri section 1 for triangular regions 1 and 6, tri section 2 for triangular regions 4 and 5, and tri section 3 for triangular regions 2 and 3.	259

LIST OF ABBREVIATIONS

1D	One-Dimensional
2D	Two-Dimensional
3D	Three-Dimensional
AHTR	Advanced High Temperature Reactor
BOC	Beginning of Cycle
BP	Burnable Poison
C-C	Carbon-Carbon (usually in reference to a carbon-carbon composite material)
CB	Control Blade (three-armed control rod inserted in the central location of fuel assemblies)
CE/LI	Constant Extrapolation / Linear Interpolation
CHM	Carbon-to-Heavy-Metal (usually in reference to a ratio between graphite and fuel)
EFPD	Effective Full Power Days
EOC	End of Cycle
FHR	Fluoride salt-cooled High temperature Reactor
FLiBe	Commonly used name for the coolant salt with actual chemical composition Li_2BeF_4
I/O	Input/Output (in reference to code reading from and writing to external files)
LWR	Light Water Reactor
MC	Monte Carlo
MHC	Molybdenum Hafnium Carbide (used for control blades)
MSR	Molten Salt Reactor
MSRE	Molten Salt Reactor Experiment (at Oak Ridge National Laboratory in the 1960s)
ORNL	Oak Ridge National Laboratory
PC	Predictor-Corrector
PPF	Power Peaking Factor

SIE	Stochastic Implicit Euler
TH	Thermal Hydraulic(s)
TRISO	TRi-structural ISOTropic (layering scheme used for fuel particles)

SUMMARY

Recent work on new nuclear reactors has focused on Generation IV designs, which includes molten salt technologies. One type of molten salt reactor with a solid fuel form is a fluoride salt-cooled high-temperature reactor. A prismatic fuel design of this type is the Advanced High Temperature Reactor (AHTR). AHTR uses hexagonal fuel assemblies that contain fuel plates embedded with TRISO fuel particles. The geometry of the AHTR fuel assembly is complex to model, so many previous studies have used simplifications to make simulations more feasible. This work created a parameterized multiphysics framework of the 3D AHTR core to allow for both ease of design change and detail in simulations, all while remaining practical using typically available computational resources. Detailed simulations were executed using the Monte Carlo code Serpent with key features including criticality search, depletion, and multiphysics capabilities coupling neutronics with materials property changes (thermal expansion and heat conduction) and thermal hydraulics. These areas were incorporated into a novel AHTR-specific framework called ATOMICS, which was used to conduct several sensitivity studies and depletion simulations. Results demonstrated the impact of model refinements made possible by ATOMICS as well as provided information for potential future design changes made to AHTR. Despite AHTR being a large system susceptible to spatial numerical instabilities, the depletion processes used by ATOMICS were shown to be mostly numerical stable for the cases considered when appropriate methods and options were used. ATOMICS is a practical and flexible tool enabling realistic analysis of AHTR for statepoint and depletion simulations on conventional computing environments.

CHAPTER 1

THE FHR-AHTR SYSTEM

A subset of molten salt reactors, Fluoride salt-cooled High-temperature Reactors (FHRs) are an advanced reactor concept, which utilizes liquid salt as the primary coolant and has a solid fuel form. An advantage of using fluoride salt as a coolant is that it has a very high boiling point (1430 °C for the specific salt Li_2BeF_4 , commonly referred to by FLiBe, considered in this work [1]), meaning that the reactor can be safely operated at atmospheric pressure with a large temperature margin beyond the normal operating regime. Having a non-pressurized system is advantageous from both a system design and accident mitigation standpoint. A specific FHR design originating from Oak Ridge National Laboratory (ORNL) is the Advanced High Temperature Reactor (AHTR) [2], which will be the focus of this dissertation. A three-dimensional (3D) depiction of the AHTR reactor system can be seen in Figure 1.1 [2] [3].

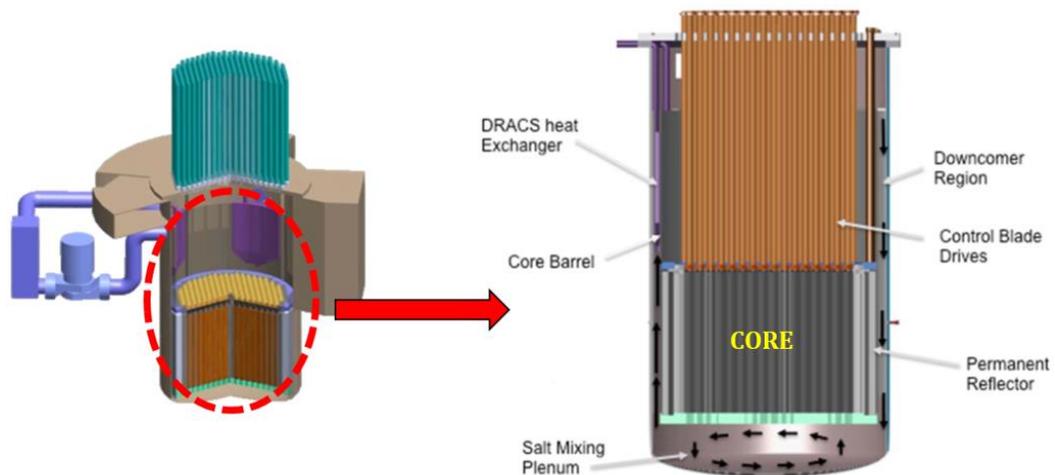


Figure 1.1. AHTR full core with main features [2] [3].

AHTR uses prismatic assemblies with fuel particles embedded in fuel planks placed in one-third symmetric locations and the core has 252 fuel assemblies. The fuel particles of AHTR are TRi-structural ISOtropic (TRISO) particles, which are composed of spherical fuel kernels

surrounded by layers of carbonaceous material which function as an additional fission product barrier. While the fuel form of choice is used in several advanced reactor designs due to TRISO's inherent safety performance, the heterogeneity of randomly dispersed fuel particles in an already complex geometry creates modeling challenges. This so-called "double heterogeneity" of the geometry is possible to explicitly model in existing neutronics tools, such as Monte Carlo (MC) transport codes, but often requires more computational resources than traditional lattice designs. Renderings of the AHTR fuel assembly with the appropriate boundary interfaces with neighboring assemblies and a two-dimensional (2D) radial core layout of the 252 fuel assemblies can be seen in Figure 1.2.

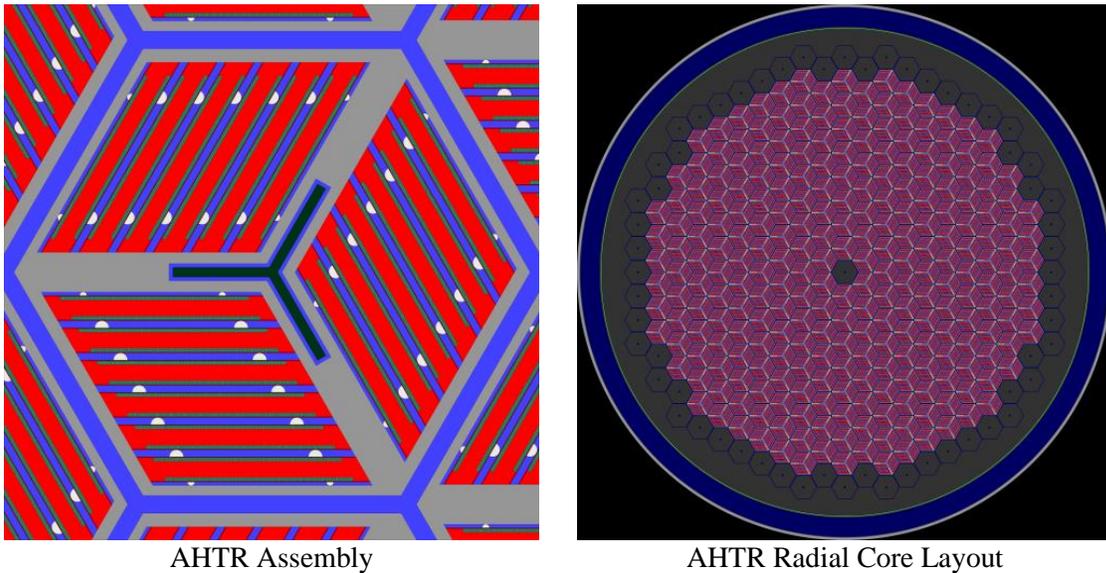


Figure 1.2. Hexagonal AHTR assembly and radial core layout.

In addition to neutronics, another important component for reactor physics studies in this dissertation is the thermal performance of core fuel, structures, and materials. This includes thermal expansion for components in the high temperature (greater than 650 °C) reactor, thermal hydraulics (TH) for heat transfer between components, and corresponding density changes due to both phenomena. The reactivity impact due to the thermal performance is on the order of hundreds of pcm. This warrants consideration in an AHTR model to adequately capture these effects.

The main challenge addressed in this work is to develop a methodology and create a practical tool which can capture the coupled multiphysics (neutronics, thermal expansion of components, and TH for the movement of heat) of the AHTR core as well as run in a timely manner. This is especially relevant for depletion simulations which require several transport calculations for all the burnup steps and feedback iterations considered. The practicality will be reflected by having a flexible tool which will allow for numerous core design modifications to the reference design to be easily modeled and simulated in acceptable time. This will allow for future researchers to make simple input file specifications to create comprehensive changes to a complex 3D system which would otherwise impact hundreds or thousands of lines of code. This dissertation will start by framing the AHTR, discuss previous related studies on core physics, detail the components of relevant multiphysics, and then summarize the results found from using the tool incorporating these multiphysics effects.

1.1 Historical Background

The AHTR design was born out of ORNL in 2003 [4] as revival of work on MSRs. ORNL had previous experience with the Molten Salt Reactor Experiment (MSRE) in the 1960's [5], which was also a FLiBe-cooled system like the proposed AHTR. There was renewed interest in MSRs as it was a candidate advanced reactor technology for providing both electricity and process heat for industrial applications. In addition to the large prismatic FHR design considered in this dissertation, ORNL also created preconceptual designs for a small modular AHTR [6] as well as a pebble bed AHTR [7]. The later evolved into a pebble bed FHR [8], receiving further analysis from multiple collaborating universities and catalyzing a nuclear reactor startup company to bring the design to market [9]. These concepts have been further expanded and developed to better explore their own technical merits but will not be mentioned any further in this document.

The AHTR design eventually evolved from using fuel pins to the fuel plate design considered in this work and shown previously in Figure 1.2. Fuel loadings and core lifetimes had

changed over time [2], and it is entirely possible that it could happen again depending on what performance is desired from the design. For this reason, one of the goals of this work was to have certain design parameters as variables, which can be changed by users of the proposed tool in this dissertation. An example of one such modification was an increase in the Carbon-to-Heavy-Metal (CHM) ratio to go from approximately 200 to about 400; effectively halving the amount of fuel loaded into the core. This increased core Beginning of Cycle (BOC) reactivity while also reducing the cycle length of the design. Making this modification to a static full core model would be significant, but simple for a model that is parameterized. This example justifies the creation of an adaptive tool, since the design requirements and cycle needs of a mature AHTR design can differ significantly from those proposed in early preconceptual documents. Being flexible would lend such a tool to still being relevant especially when future design changes are made.

1.2 3D Reference Core Geometry

The models referenced in this dissertation follow from design documents published by ORNL from 2011 [3] and 2012 [2]. Careful attention was paid toward matching these specifications closely, but some minor simplifying assumptions were made along the way, which will be documented where appropriate. A 3D representation of the AHTR core geometry can be seen in Figure 1.3 [2]. The active core, which is composed of 252 fuel assemblies with replaceable reflector assemblies at the central and peripheral locations, is surrounded by a fixed reflector to reduce neutron leakage and improve neutron economy. Radially, replaceable reflector assemblies (solid graphite hexagonal block the same size as a fuel assembly except for a 2 cm hole in the center) are embedded in a permanent reflector region (dark gray), which is further surrounded by a thin boron carbide layer and then a core barrel layer in contact with the downcomer coolant. Beyond the downcomer region is the reactor pressure vessel wall, which is composed of a thin corrosion-resistant inner liner and thicker steel alloy wall. Axially, the fueled core region is extended by unfueled axial reflectors with the same geometry. On the other sides of these axial reflectors are

lower (green) and upper (blue) core support plates, meant to keep assemblies in their correct locations. General dimensional and material parameters of these sections are summarized in Table 1.1.

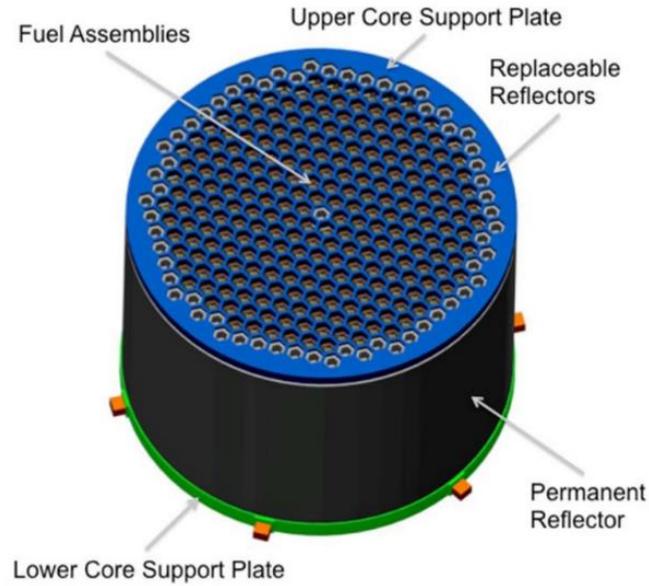


Figure 1.3. 3D depiction of the AHTR core [2].

Table 1.1. Dimensions and compositions of major components of the AHTR core.

Feature	Dimension	Material Composition
Number of Fuel Assemblies	252	-
Fuel Assembly Pitch	46.8 cm	-
Permanent Reflector Outer Radius	478 cm	C-C Composite
Boron Carbide Outer Radius	479 cm	B ₄ C
Core Barrel Outer Radius	481 cm	C-C Composite
Downcomer Outer Radius	519 cm	FLiBe (Li ₂ BeF ₄)
Pressure Vessel Liner Outer Radius	520 cm	Alloy N (INOR-8)
Pressure Vessel Outer Radius	525 cm	Incoloy alloy 800H
Active Core Height	550 cm	-
Lower Axial Reflector Height	25 cm	-
Upper Axial Reflector Height	25 cm	-
Modeled Assembly Height	600 cm	-
Lower Support Plate Height	35 cm	C-C Composite
Upper Support Plate Height	35 cm	SiC
Total Model Height	670 cm	-

1.2.1 Fuel Assembly Geometry and Dimensions

The fuel assemblies of the AHTR are complex to model due to being composed of several fuel planks (or plates) with TRISO fuel particles embedded in fuel stripes. Models involving TRISO

particles are normally referred to as being “double heterogeneous”, because the structures do not repeat as simply as fuel pins in a lattice geometry.

1.2.1.a TRISO Fuel Particles

The TRISO particles of the AHTR are AGR-2 type [10], which feature uranium oxycarbide fuel kernels. The specific molar composition is 71.4% UO_2 , 12.3% $\text{UC}_{1.86}$, and 16.4% UC which can effectively be represented as $\text{UC}_{0.392}\text{O}_{1.427}$ when normalized relative to uranium. The fuel kernels are surrounded by several layers of carbonaceous material meant to serve as fission product barriers. This is one of primary benefits of selecting TRISO particles as a fuel form since the surrounding layers are better able to retain fission products even under adverse conditions than conventionally clad pins. The separate layers of the TRISO particles can be seen in Figure 1.4. The dimensions of the TRISO particle layers can be found in Table 1.2.

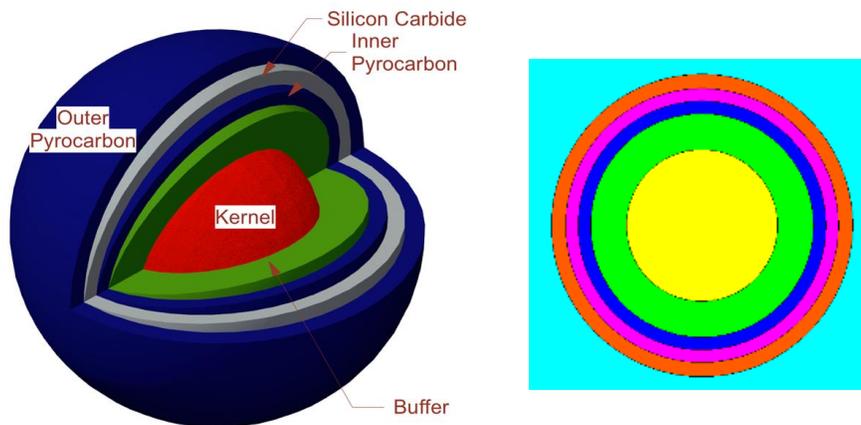


Figure 1.4. Left: TRISO fuel particle with all layers identified [3]. Right: AGR-2 TRISO particle modeled in SERPENT [11] with surrounding matrix material (teal) with average cubic pitch.

The original concepts of the AHTR using TRISO fuel particles proposed that the particles would be heterogeneously dispersed in fuel stripes. To start, this can be challenging (though still possible) for some transport codes to handle, as explicit heterogeneous particle placement needs to be either pre-generated before a transport simulation or done on-the-fly, which require a significant

amount of memory to read-in or additional computational overhead, respectively. One simplification to the modeling process would be to instead model the TRISO particles in a regular cuboidal lattice to better leverage the utilities readily available in reactor physics codes. A further justification for this is that recent fabrication research advancements involving additive manufacturing have introduced potentially novel ways to manufacture a new reactor fuel design. A new concept could be “3D printed”, so particles placement in regular intervals (i.e. a lattice) is entirely feasible. The Transformational Challenge Reactor (TCR) is already being investigated at ORNL and is planned to be completed in 2023 [12]. It will be fueled by TRISO particles and will be made using additive manufacturing techniques [13]. It is not unreasonable that if this technology proves successful and economical, a design such as AHTR could also be produced using additive manufacturing and have fuel particles in a lattice structure.

Table 1.2. TRISO particle layers and dimensions.

Parameter	Dimension
Fuel Kernel Radius	0.02135 cm
Buffer Outer Radius	0.03135 cm
Inner Pyrolytic Carbon Outer Radius	0.03485 cm
Silicon Carbide Outer Radius	0.03835 cm
Outer Pyrolytic Carbon Outer Radius	0.04235 cm
TRISO Particle Packing Fraction	40%
Average Particle Pitch	0.09266 cm

For this dissertation, a cuboidal TRISO particle lattice is used. The reasoning is that by starting with a lattice system independent in each dimension, future researchers could easily modify the design to suit their needs (including making the pitch cubic). While the average pitch for AGR-2 particles with a 40% packing fraction is 0.09266 cm, the assumed pitches in each direction are as follows: x-pitch 0.09406 cm, y-pitch 0.09128 cm, and z-pitch 0.09266 cm.

1.2.1.b Fuel Plank

TRISO particles are placed in two fuel stripes in each fuel plank of the AHTR (Figure 1.5). For this dissertation, it is assumed that fuel stripes are 202 particles wide in the x-direction and four

layers thick in the y-direction. This provides stripe dimensions of 19.00012 cm in the x-direction and 0.36512 cm in the y-direction. It is assumed that fuel stripe centers are centered along the coolant channel¹.



Figure 1.5. AHTR fuel plank with two fuel stripes.

Fuel planks are 23.1 cm long and 2.55 cm wide. There is a 0.1 cm graphite “sheath” or “sleeve” separating the fuel stripes from the coolant. Planks are angled at 60° and 120° at the corners. There are also burnable poison particles located at the centerline of the fuel plank, but those will be addressed separately in the section on reactivity control.

1.2.1.c Fuel Assembly

The AHTR fuel assembly is one-third rotationally symmetric, and many of its key features are highlighted in Figure 1.6. Each one-third assembly section contains six fuel planks, with a coolant channel on each side (seven in total). Fuel planks are separated from each other by large graphite spacers and from the wrapper and Y-structure by smaller graphite spacers (ten large, four small, and fourteen total spacers). Fuel planks are also embedded into the wrapper and Y-shaped structure of the assembly. These two features have indents to accommodate the fuel planks and notched elsewhere to keep the planks in place.

¹ Fuel planks are not exactly centered with respect to the coolant channel. They are embedded deeper in the central Y-shaped support member than in the assembly wrapper along the perimeter. This makes it so that the center of the fuel plank is shifted toward the central support member, which differs from the coolant channel center. This difference is quite small and results using either centering basis should be comparable.

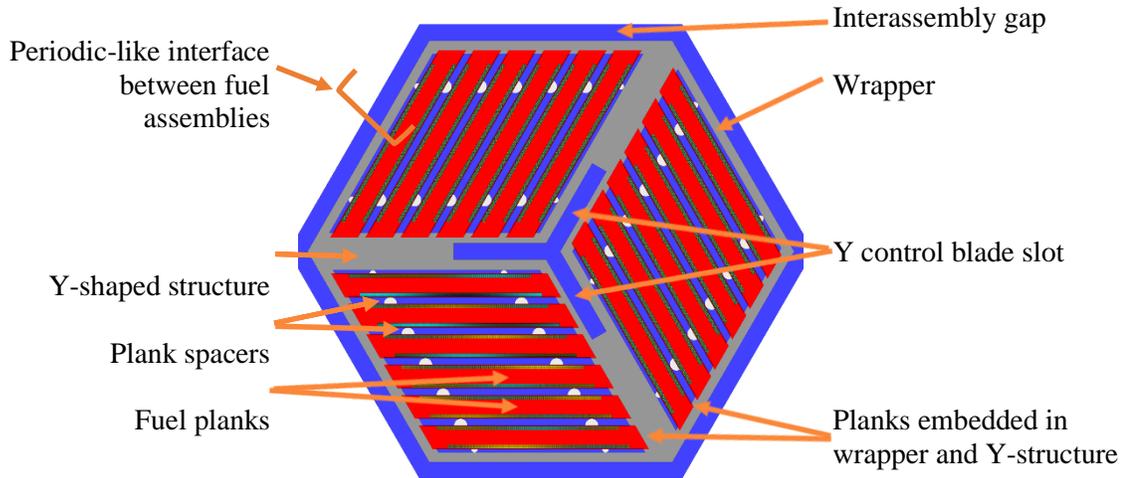


Figure 1.6. AHTR fuel assembly with key features identified.

Each assembly has a Y-shaped control blade (CB) slot at the center to accommodate insertion of a three-pronged cruciform control rod. The slot is larger than the CB to allow for insertion even under off-normal operating conditions. CBs will be addressed further in the section on reactivity control. A summary of the dimensions used in constructing the AHTR assembly geometry can be found in Table 1.3. The interassembly gap shown in Figure 1.6 is the full width of the gap. Each assembly would have its own gap half this size, but the full channel is shown here for illustrative purposes.

Table 1.3. Parameters and dimensions used to construct the AHTR fuel assembly.

Parameter	Dimension [cm]
Assembly apothem	22.5
Assembly pitch	46.8
Interassembly Gap	1.8
Large coolant channel width and spacer radius	0.7
Small coolant channel width and spacer radius	0.35
Spacer pitch	14
Wrapper thickness (plank indent)	1
Wrapper thickness (notch)	1.35
Y-shape thickness (notch)	4
Control blade arm length	10
Control blade arm width	1
Control blade channel length	10.38
Control blade channel width	1.76

1.2.2 Reactivity Control

Excess reactivity in the AHTR core is controlled by two means: passively by burnable poisons embedded in each fuel plank and actively by moving CBs. Each will be explored further below.

1.2.2.a Burnable Poison Spheres

Burnable poison (BP) spheres are composed of europia (Eu_2O_3) and are located at the center of the fuel plank. BP spheres are stacked in a few discrete columns separated by a set distance so that they extend to a length comparable to that of the fuel stripes. Several options were considered for the reference design [2], including the number of BP columns, size of particles, and other parameters. Due to issues with depletion simulations performed in that study involving the tracking of europium isotopes, the data used to inform reference BP loading underpredicted the reactivity penalty (especially over the cycle as the poisons burn out) of the europia. For this reason, as confirmed by the results of this dissertation, the reactivity penalty of the BP spheres is still very large at the discharge burnup. Improvements would thus be desirable and are considered in Chapter 6. The issue was acknowledged by the authors of the reference design and despite this, these reference values are carried through many of the studies conducted in the initial chapters of this document for consistency.

BP spheres have a cold radius of 350 μm and arranged in five columns. Spheres have a cold axial pitch of 0.09936 cm, which is slightly larger than the cold axial pitch of the TRISO fuel particles (0.09266 cm). Nominally, columns are separated by 4 cm in the reference document, but this dissertation assumes that this distance is 3.9744 cm cold (the separation distance is arbitrarily selected by the reference, and this separation allows for a single lattice of BP to be used for each plank (since $40 \cdot 0.09936 \text{ cm} = 3.9744 \text{ cm}$).

The europia material density for the BP spheres is assumed to be 5.0 g/cm^3 , which is 68% of the theoretical density for Eu_2O_3 [2]. It is already assumed to be feasible from prior design

documents that the BP spheres can be manufactured with a given level of porosity (or potentially graphite doping) to reduce their density, as this will be considered further in a later section.

1.2.2.b Control Blades

CBs are composed of molybdenum hafnium carbide (MHC). It is a fairly dense material (10.28 g/cm^3) which is advantageous for a system like AHTR with a relatively dense coolant salt (about 1.95 g/cm^3) and CBs being inserted from the top of the core. Each of the 252 fuel assemblies has a slot to accommodate a CB, which is illustrated in Figure 1.7. CBs are the primary means of active reactivity control and their movement is an important component of this dissertation. They are addressed in further detail in Chapter 6.

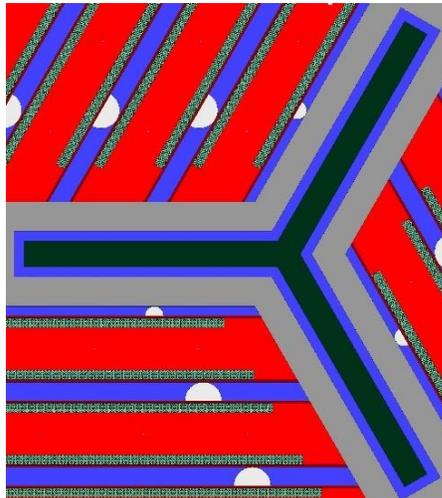


Figure 1.7. Control blade (dark green) inserted in the central location of an AHTR fuel assembly.

1.2.3 Materials and Compositions

A comprehensive list of materials used in the AHTR model is given in Table 1.4. Most material densities are assumed to be measured at room temperature (293 K), with the exception being FLiBe since the coolant density changes with temperature. At the inlet temperature ($650 \text{ }^\circ\text{C}$) the density is 1.963 g/cm^3 and decreases to about 1.938 g/cm^3 at the average outlet temperature (about $700 \text{ }^\circ\text{C}$). The average core density would correspond to the density at the average temperature (about $675 \text{ }^\circ\text{C}$), which is approximately 1.950 g/cm^3 . This is the only FLiBe material

listing in Table 1.4, but all other compositions are scaled from this composition relative to the local density of the coolant.

Table 1.4. List of materials and their compositions using in the reference AHTR model.

Material	Atomic Concentration [atom/b/cm]		Density [g/cm ³]
Fuel (9 w% enriched) (Approximate composition: UC _{0.392} O _{1.427})	92235	2.27325E-03	11.00
	92238	2.26948E-02	
	6012	9.79715E-03	
	8016	3.56187E-02	
	Total	6.91649E-02	
TRISO Buffer	6012	5.01845E-02	1.00
Inner Pyrolytic Carbon	6012	9.53506E-02	1.90
Outer Pyrolytic Carbon	6012	9.38450E-02	1.87
Graphite	6012	8.78229E-02	1.75
C-C Composite	6012	9.78598E-02	1.95
SiC	14028	2.14695E-02	3.10
	14029	1.09443E-03	
	14030	7.21861E-04	
	6012	2.32858E-02	
	Total	4.65717E-02	
FLiBe	3006	1.38301E-06	1.95
	3007	2.37132E-02	
	4009	1.18573E-02	
	9019	4.74291E-02	
	Total	8.30010E-02	
Burnable Poison (Eu ₂ O ₃)	63151	8.17979E-03	5.00
	63153	8.93275E-03	
	8016	2.56688E-02	
	Total	4.27814E-02	
Control Blade	6012	5.15897E-04	10.28
	72174	6.65953E-07	
	72176	2.18932E-05	
	72177	7.74170E-05	
	72178	1.13545E-04	
	72179	5.66892E-05	
	72180	1.46010E-04	
	42092	9.33048E-03	
	42094	5.85305E-03	
	42095	1.01075E-02	
	42096	1.06170E-02	

Table 1.4 (continued).

	42097	6.10143E-03	
	42098	1.54701E-02	
	<u>42100</u>	<u>6.20334E-03</u>	
	Total	6.46150E-02	
Boron Carbide	5010	2.05649E-02	2.37
	5011	8.27763E-02	
	<u>6012</u>	<u>2.58353E-02</u>	
	Total	1.29176E-01	
Pressure Vessel Liner (Hastelloy Alloy N)	28000	6.33036E-02	8.93
	24000	7.23985E-03	
	42000	8.96856E-03	
	26000	3.85193E-03	
	14000	1.91479E-03	
	25055	7.83104E-04	
	23000	5.27838E-04	
	6012	2.68889E-04	
	27059	1.82504E-04	
	29000	2.96198E-04	
	74000	1.46262E-04	
	13027	4.98283E-04	
	<u>22000</u>	<u>2.80871E-04</u>	
	Total	8.82627E-02	
Pressure Vessel (Alloy 800-H)	28000	2.64101E-02	7.92
	24000	1.92630E-02	
	6012	2.98096E-04	
	25055	1.30225E-03	
	16000	2.23119E-05	
	14000	1.69822E-03	
	29000	5.62923E-04	
	15031	6.92938E-05	
	13027	6.62890E-04	
	22000	3.73655E-04	
	<u>26000</u>	<u>3.61825E-02</u>	
Total	8.68453E-02		

1.3 Previous Studies

It should be noted here that the exact configuration of the AHTR assembly or core varied (sometimes significantly) for each study considered in this section. The most common difference is with respect to the fuel stripe: stripe dimension, heavy metal loading, and TRISO treatment

among heterogeneous, lattice, and various homogenization approaches. These differences are challenging to summarize for every case, but the reader should keep in mind that almost universally, the configurations used in the other works referenced in this section differ (from slightly to significantly) from the design detailed in the previous section.

1.3.1 Previous Work Conducted on Neutronics

Several reactor physics studies have been performed on the AHTR design in the past and have been summarized in previous reports [14]. A notable takeaway from analyses in the past is how the double heterogeneity of AHTR has challenged the modeling capabilities of the tools used. Previous depletion studies used various homogenization techniques such as reactivity-equivalent physical transformation [15] and the Dancoff correction method [16] to simplify the model to make simulations more tractable. Today, virtually all the major transport codes are capable of modeling double heterogeneous systems, but both long runtimes and large computational overhead remain as potential challenges.

While many studies have focused on specific core designs, ones relevant to this dissertation focused on sensitivity studies conducted on the AHTR system [17]. The importance of these sensitivity studies is to provide information on the expected impact from making changes to the core. These can include large design-based changes like dimensions of assembly components as well as materials, but also can provide estimates for relatively small changes like those expected from thermal expansion.

It will be noted here that the AHTR assembly configurations used to obtain the results of both Figure 1.8 and Figure 1.9 differ from the reference most significantly by having twice the heavy metal loading (CHM ratio closer to 200 instead of 400). Consequently, the fuel stripe is also larger than in the reference design but otherwise the assembly geometries are comparable. The expected impact of using lattice versus heterogeneously-placed TRISO particles is relevant to this dissertation, as discussed previously when selecting to use lattice TRISO particles for the reference

design. As can be seen in Figure 1.8, the difference between heterogeneous dispersion and lattice treatments can exceed 300 ± 43 pcm. The authors analyzed five different heterogeneous TRISO placements generated with different random number seeds (RAN 1, RAN 2, RAN 3, RAN4, and RAN 5) in each fuel stripe of an assembly. It does appear that the impact is largest at BOC and gradually decreases over cycle. This is important for quantifying the impact of the assumption made to use lattice TRISO particles in the reference model. The spectral differences observed between heterogeneous and lattice TRISO was very small [17].

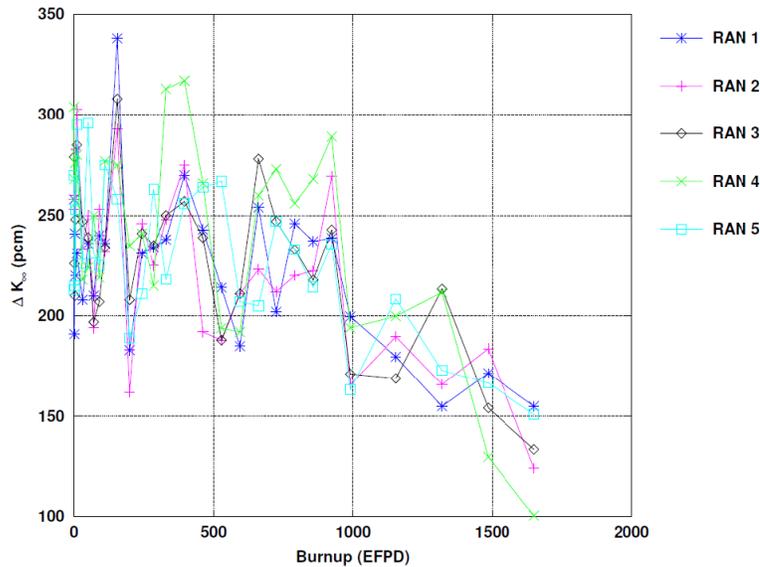


Figure 1.8. k_{∞} difference between lattice and five unique random heterogeneous dispersions of TRISO particles [17].

Another important study which will be referenced later in the document is the impact of lumping fuel stripes into single materials versus dividing them into segments in each fuel stripe. The results shown in Figure 1.9 are compared to a reference simulation of using nine equally-sized divisions in each of the 36 fuel stripes in a fuel assembly over the course of depletion schedule. “Lumped” in this study refers to two depletion regions used in the model: one for all stripes in the top of each fuel plank and one for all stripes in the bottom of each fuel plank. The researchers [17] concluded that five plank divisions gave comparable results to the reference of using nine divisions.

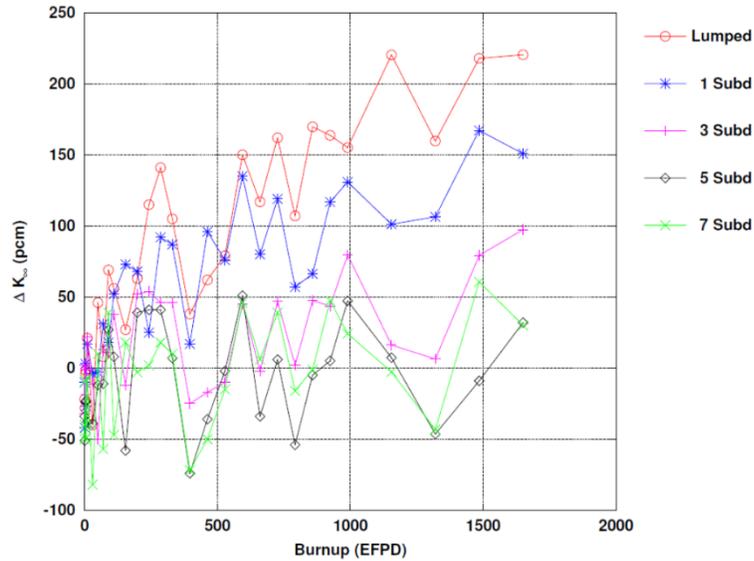


Figure 1.9. k_{∞} difference versus nine depletion subdivisions per fuel stripe [17].

The same referenced report also contains numerous sensitivity studies related to the assembly geometry. Some include plank number and thickness, thickness of assembly wrapper, and assembly size. Related work also included additional studies at the assembly level [18], a two-step procedure for coupling SERPENT with a nodal diffusion code [19], and refining the AHTR design through the use of machine learning [20].

1.3.2 Previous Work Conducted on Thermal Hydraulics

Previous TH modeling of AHTR has used specialized codes like RELAP5-3D [21] to capture the model channels within an AHTR fuel assembly. Fuel plate temperature profiles as well as results for convection to coolant channels were obtained [22]. Essentially, each component (plate, coolant channel, assembly wrapper, center structural Y-shape, etc.) was modeled in a nodal fashion with temperatures solved by RELAP. This process is slightly more robust than what is proposed for this work, mainly because the proposed model only aims to resolve temperatures at a one-third assembly level, not a per-channel basis. However, the results of this work can be compared to those of the past. It was not expected that there would be significant differences between the two sets of results.

An additional study also looked at the possibility of refueling the AHTR core while it is still operating on power [23]. While this is beyond the scope of this work, it is yet another benefit of FHRs. As they are operated at atmospheric pressure, there are no massive pressurized coverings with large bolts at the top of the vessel, so assemblies could be replaced from outside the pressure vessel without insurmountable complications.

1.3.3 Previous Work Conducted on Thermal Expansion

From what could be found in literature, no explicit work has been conducted on accounting for thermal expansion in AHTR. There have been related analyses, including by this author in results shown below, which considered sensitivity studies of the core geometry and materials [17] [24], but these are only able to provide inferential data. Later in this design process though, they can be used for verification of results. This work aims to account for thermal expansion in AHTR by making the corresponding geometry and material density changes for a given temperature. This is valuable because it has not been systematically addressed before and is shown to have a nontrivial impact on the results.

1.3.3.a Geometric Thermal Expansion of TRISO and TRISO Lattice

This section summarizes results of a series of sensitivity studies related to scoping the effects of geometric changes and geometric simplifications to a 2D assembly model. Three configurations are considered: no reactivity control or poison materials (“NP”), BP spheres located in the center of the fuel plates (“BP”), and both BP spheres and CB inserted (“BP & CB”).

To get an idea about the neutronic impact of modeling thermal expansion, simulations were run where only the TRISO layer thicknesses and TRISO lattice pitch were changed. The goal was to only investigate the *geometric* impact of TRISO lattice thermal expansion, so all material temperatures were held constant. Previous studies of the AHTR have assumed that the average fuel temperature under nominal operating conditions is about 1110 K [17]. Since all reference dimensions of the AHTR design are assumed to be at room temperature (293 K), this neglects an

average temperature difference of 817 K in the fuel. This sensitivity study assumed a rounded temperature change of 800 K for simplicity. A second expansion temperature difference of 1600 K was also considered out of interest of having another reference point for high-temperature, accident-like scenarios. Each simulated case completed with a statistical uncertainty of 5 pcm. Therefore, each difference listed in Table 1.5, Table 1.6, and Table 1.7 has a combined standard deviation of 7 pcm.

Table 1.5 shows the reactivity impact for the model with no reactivity control (i.e. no BP spheres and no CB inserted). For all three TRISO particle expansion temperatures considered, the lattice pitch expansion has very little impact on eigenvalue. All cases are within two standard deviations of the cases with the same TRISO expansion, so any trends observed for expanding the lattice are obfuscated by the statistical uncertainty. The same does not hold true when considering the TRISO particle expansion. Here, there is a definitive trend of thermal expansion in the TRISO particles resulting in a change of reactivity by about -104 ± 7 pcm in each 800 K increment considered (-0.13 ± 0.01 pcm/K).

Table 1.5. Reactivity impact of TRISO lattice geometric thermal expansion with no reactivity control. Differences have a statistical uncertainty of 7 pcm.

NP		Lattice ΔT [pcm]		
		0 K	800 K	1600 K
TRISO ΔT [K]	0 K	0	-8	-6
	800 K	-104	-120	-111
	1600 K	-202	-213	-215

Table 1.6 shows the reactivity impact for the model which uses BP spheres in the center of the fuel planks. Similar to the uncontrolled case in Table 1.5, unrestrained expansion of the TRISO particle layers results in a linear change in reactivity of about -67 ± 7 pcm for each 800 K increment. The effect is smaller in magnitude than the uncontrolled case, but the general trend is the same. The TRISO lattice expansion is more complex. While there was negligible effect for the uncontrolled model, this model has a nonlinear behavior where reactivity initially increases for

expansions from room temperature by 800 K then decreases for an additional thermal expansion by 800 K. The increase is about 54 ± 7 pcm from 0 to 800 K ΔT and the decrease is about -76 ± 7 pcm from 800 to 1600 K ΔT . Since the behavior is nonlinear, it is difficult to say which lattice expansion has the maximum reactivity but a guess of about 950 K from room temperature can be made from using only this limited data.

Table 1.6. Reactivity impact of TRISO lattice geometric thermal expansion with burnable poison spheres. Differences have a statistical uncertainty of 7 pcm.

BP		Lattice ΔT [pcm]		
		0 K	800 K	1600 K
TRISO ΔT [K]	0 K	0	48	-25
	800 K	-68	-11	-88
	1600 K	-144	-86	-164

Table 1.7 shows the reactivity impact for the model which uses both BP spheres along the centerline of the fuel planks and a CB in the center of the assembly. Qualitatively, the results are very similar to those of the model with just BP spheres. As the TRISO particles expand, reactivity decreases. The magnitude of the change is not as uniform as in the previous models, so interaction effects seem to be more important. As the lattice expands, the nonlinear behavior of reactivity increasing from 0 to 800 K ΔT and decreasing from 800 to 1600 K ΔT is again observed.

Table 1.7. Reactivity impact of TRISO lattice thermal expansion with both BP spheres and control blade. Differences have a statistical uncertainty of 7 pcm.

BP & CB		Lattice ΔT [pcm]		
		0 K	800 K	1600 K
TRISO ΔT [K]	0 K	0	68	-33
	800 K	-47	28	-61
	1600 K	-75	-25	-121

1.3.3.b Density Changes in Graphite Components

Thermal expansion will impact both the geometry and density of the expanding media. In the previous subsection, the effects of changing the TRISO particle and TRISO lattice geometry were considered. Here, attention is turned to carbonaceous material density changes. Recall in the

last section that a temperature change of about 800 K was considered average. Assuming a thermal expansion coefficient of $5 \times 10^{-6} \text{ K}^{-1}$ for graphite, one could expect density changes of about 1.2% from room temperature to operating conditions. Table 1.8 shows the reactivity impact of various carbonaceous structures changing density by $\pm 10\%$. It should be stated that solely modifying the density changes the amount of media present in the model, which differs from thermal expansion. The amount of media is preserved in thermal expansion due to the increase in dimension and corresponding decrease in density. The results of Table 1.8 are presented merely to illustrate that density changes can be inferred to be important, though overstated here due to both larger change in density and the combined effect of losing carbonaceous media which increases reactivity difference. Actual reactivity changes due to thermal expansion in the proposed model for this work are smaller than those presented in Table 1.8, yet still significant.

Table 1.8. Material-wise graphite density sensitivity [24].

Modified Material	k(-10%)*	Normal*	k(+10%)*	$\Delta k(-10\%)$ [pcm]	$\Delta k(+10\%)$ [pcm]
Matrix	1.39174	1.39327	1.39494	-153 ± 7	167 ± 7
Meat	1.38448	1.39327	1.4018	-879 ± 7	853 ± 7
Sleeve	1.39213	1.39327	1.39457	-114 ± 7	130 ± 7
Spacer	1.39301	1.39327	1.39367	-26 ± 7	40 ± 7
Structural	1.38875	1.39327	1.39776	-452 ± 7	449 ± 7
Wrapper	1.39084	1.39327	1.39577	-243 ± 7	250 ± 7
All (sum Δ)				-1867 ± 17	1889 ± 17
All (simulated)	1.3727	1.39286	1.41037	-2016 ± 7	1751 ± 7

*Each simulation had a reported statistical uncertainty of 5 pcm.

As has been discussed, numerous studies have been conducted on the AHTR core design. Other researchers and this author have done a variety of sensitivity studies at the assembly or simplified-core level. However, detailed full-core models as well as methodologies supporting future analyses are lacking and more could be done as is addressed in this dissertation.

CHAPTER 2

MOTIVATION AND OBJECTIVES

The driving purpose of this research is to create a practical tool to analyze an advanced reactor concept. Next generation nuclear power plant designs are currently at the forefront of the field, but some of the tools needed to properly study them simply do not exist. The generic tools available today have the necessary capabilities and features, but some are too cumbersome to use in realistic analyses. The overarching objective of this dissertation is to develop a full core modeling methodology and implement it for AHTR to provide accurate results on an acceptable timescale. Major efforts can be categorized as relating to the multiphysics phenomena of thermal expansion, TH, and neutronics. Thermal expansion focuses on capturing the dimensional changes of core components based on the temperature of those components. TH accounts for the heat transfer of fission energy produced in the fuel kernels to other assembly components until being removed by the coolant and the changing FLiBe coolant properties. Neutronics focuses on the reactor physics aspects of the design under normal operation, such as maintaining criticality and determining the power profile in the core, as well as accounting for depletion over cycle. Making improvements to these three areas should reduce uncertainties stemming from approximations used in previous studies. The key to this work is executing all three broad areas together to have a model, which can account for these behaviors in a coupled manner.

2.1 Motivation

Previous studies of the AHTR core primarily fall into two categories: detailed models of a small partition of the core (most commonly a reflected assembly section) or 3D models with significant simplifications (usually with respect to the fuel particles or fuel stripes being homogenized). The motivation of this work is to bridge the knowledge gap between these two areas by creating a novel 3D model of the AHTR core, which is still able to capture fine elements of the

geometry, incorporate multiphysics aspects, and remain trackable to simulate for typical users with the computational resources available to them. The 3D nature of the model should make the multiphysics results relevant to making system-level decisions while the fine details with limited simplifications should reduce the impact of approximations like those present in previous studies. This highlights the value of the conducted work and the importance of the resulting model for future studies of the AHTR core.

2.2 Objectives

Aligning with the motivation, this dissertation will create a tractable, detailed 3D model of the AHTR core, which employs multiphysics coupling. As stated before, “tractable” here refers to being able to run simulations using typical computational resources in a timely manner. This is imagined to be a computer cluster with a few nodes and between 100 to 1000 cores. “Timely manner” is highly case-dependent, but single state points should be able to complete on the order of hours with multipoint depletion cases running on the order of days given sufficient computational resources and statistical uncertainty tolerance. The bulk of this dissertation will follow the development of a “detailed 3D model of the AHTR core”. This can be broadly divided into individual components of multiphysics meeting three main objectives:

- Firstly, to account for thermal expansion of components in the core. When parts are fabricated, they are designed to a specific dimension and density for a given temperature. In this dissertation, it is assumed that the dimensions of components are specified at room temperature (293 K). As any medium changes in temperature, it undergoes thermal expansion by means of increasing in dimension and correspondingly decreasing in density with an increase in temperature. The AHTR inlet coolant temperature is 650 °C, or about 923 K. This means temperatures of structural materials are expected to change by at least 630 K but will be significantly higher in parts of the fuel assemblies near the fuel particles where heat is generated. These large temperature differences will cause expansions on the

order of a few tenths of a percent for most cases and warrant consideration in a detailed core model.

- Secondly, to conduct a TH analysis of the fuel assembly. As heat is generated in the TRISO fuel particles, it will ultimately be removed by the FLiBe coolant flowing through and around the fuel assemblies. Accurately modeling this heat transfer will provide temperature profiles for using the correct amount of thermal expansion as well as the corresponding densities of media in the core. Additionally, obtaining temperatures for components of the active core is essential for Doppler broadening cross sections, which is important to neutronics for obtaining the correct reaction rates.
- Thirdly, to accurately model the neutronic behavior within the core. A vital part of reactor operation is maintaining criticality over the fuel cycle. Since excess reactivity needs to be designed into a core loading pattern to account for the burnup of fuel over time, there also needs to be ways to control that excess reactivity. AHTR uses both BP spheres and movable CBs, which will need to be tracked over the fuel cycle. BP spheres will undergo significant isotopic changes over the fuel cycle and their reactivity penalty should diminish over the fuel cycle. CBs will need to be moved to maintain a critical configuration. Additionally, the neutron flux profile of a nuclear reactor highly influences how the core behavior evolves in time due to the depletion of fissionable material and the creation of new isotopes through capture events as well as those born as fission products. A converged neutron flux profile will also give a corresponding fission energy deposition profile, which describes how power is produced in the core.

As should be evident from these points, a detailed 3D model of the AHTR core will rely on these areas of multiphysics working together in an integrated fashion since their physics are coupled. Neutronics provides the power profile of the core, which when coupled together with TH, should provide the temperature distribution for every part of the reactor core. With these temperatures, it is possible to update the geometry and materials due to thermal expansion.

Changing the geometry and material compositions will affect the neutronics and power distribution, thus creating a coupled loop, which requires converging upon. This leads to an additional, and likely the most important objective, which is to ensure that these components of multiphysics vital to the 3D model provide consistent results especially when interfacing with each other.

2.3 Scope

The scope of this work is best realized by breaking individual tasks down into basic components. Each bullet below can be considered a success metric, as they should all be met to satisfy the objectives previously addressed. These individual areas of multiphysics will overlap and need to interface with each other, so points addressed in one area will in most cases need consideration to accommodate the intended capability in the other components of multiphysics.

Thermal Expansion

- Model the geometric thermal expansion of all in-core components of AHTR.
- Preserve the amount of media present in the core by correspondingly changing the density of structures as part of their expansion.
- Accommodate any conflicts in expansion by yielding to the more important neutronic structure.

Thermal Hydraulics

- Create temperature distribution for assembly section components dependent upon the 3D power distribution.
- Account for assembly-wise heating of coolant based on 3D power distribution.
- Have parameters and correlations be easy to modify so that future work can adjust to new physical parameters and multiple correlation options relevant for the system.

Neutronics

- Simulate the AHTR core with a high-fidelity model.

- Model should have variable input parameters, which can be controlled and changed by users.
- Iterate with the other areas to have a physically consistent model.
- Create a tractable tallying methodology capable of iterating on the power profile without significantly hindering conventional computing resources.
- Existing methods have been identified as being problematic since they can substantially increase simulation runtimes (by more than an order of magnitude).
- Create capability to control core reactivity by moving CBs.
- Active reactivity control maintains core criticality while also impacting the power profile in the core. Converging on this capability is relevant to neutronics but also affects TH and consequently thermal expansion and then back neutronics.
- Be able to deplete the core and adjust other areas as needed as core behavior evolves over cycle.

It should be noted that this tool is primarily intended for steady state operation of the core. This could potentially extend to long-timescale testing scenarios such as zero-power tests and steady high-temperature events such as loss of flow after an extended period post-shutdown (no natural convection capabilities but could model whole core as isothermal to estimate the eigenvalue). Beyond the scope of this tool are short-timescale scenarios such as transients and most accidents. There are no time-dependent considerations in the implementation and any event which evolves rapidly (i.e. cannot be reasonably approximated by a steady state assumption) should not be considered using this model.

2.4 Approach

This work aims to incorporate the multiphysics effects of thermal expansion and TH into high fidelity reactor physics simulations specifically created for AHTR. Thermal expansion and TH will be accounted for by in-house C++ scripts written specifically for the AHTR system.

Reactor physics simulations will be performed by using the Serpent 3D MC transport code [11]. A flowchart summarizing how the work is planned to come together is shown in Figure 2.1, which is comparable to a standard TH coupling scheme. While implemented in the past for other reactor designs, the goal of this work is to create a practical tool for the design of the AHTR. In the development process, temperature (TH) and dimension (thermal expansion) capabilities were developed and then coupled. Once both functioned as intended, criticality iteration was added by means of moving CBs. All these multiphysics capabilities together function to obtain an updated power solution. If the resulting distribution is within a given tolerance, convergence will be met.

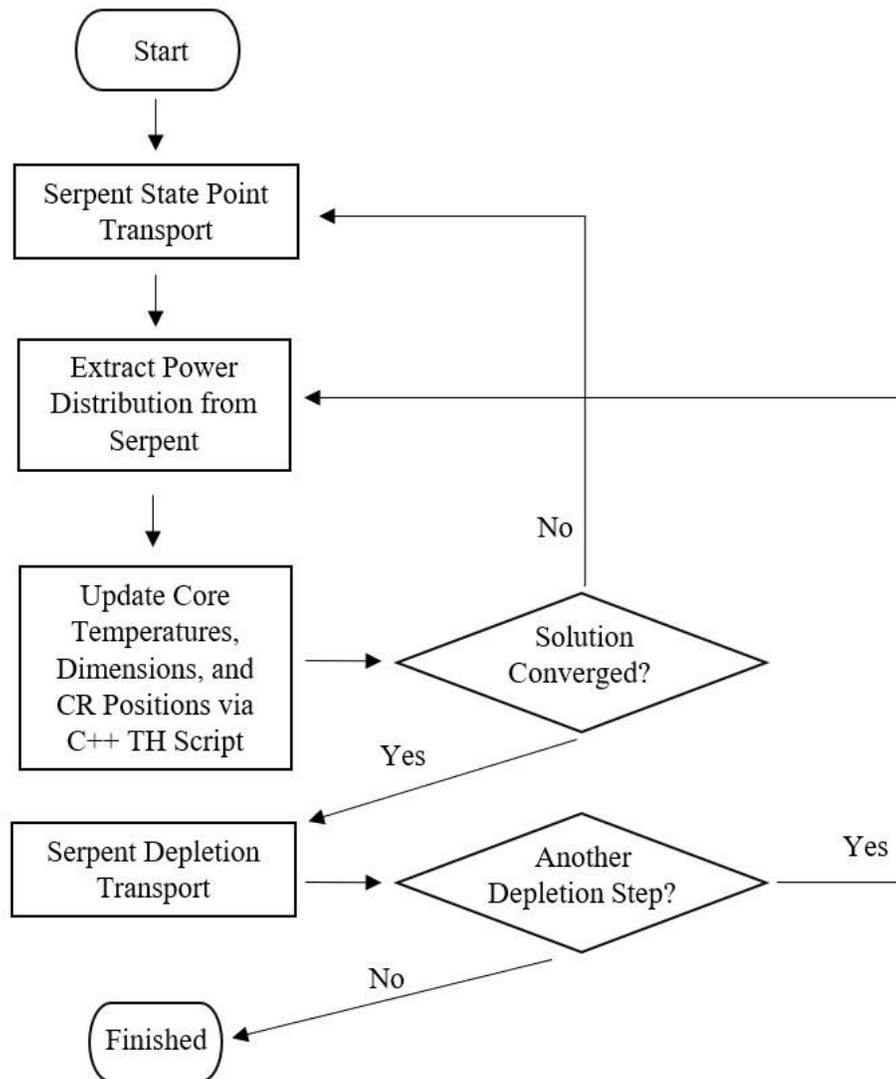


Figure 2.1. Conceptual flowchart of how the multiphysics effects will be coupled and implemented.

In the 3D full core geometry of the AHTR, there are 252 fuel assemblies which can be grouped into 84 one-third symmetric core locations. The codes used in this dissertation leverage these symmetric locations to reduce the number of uniquely tracked materials as part of depletion studies while still retaining one-third assembly resolution for results. Figure 2.2 shows how all the assembly locations are grouped. Generally, there are 84 groups with the “first” assembly of the group generally aligning in the positive-y direction from the center of the core and the clockwise one-third sector about that reference, and the corresponding symmetric group members progress clockwise around the core.

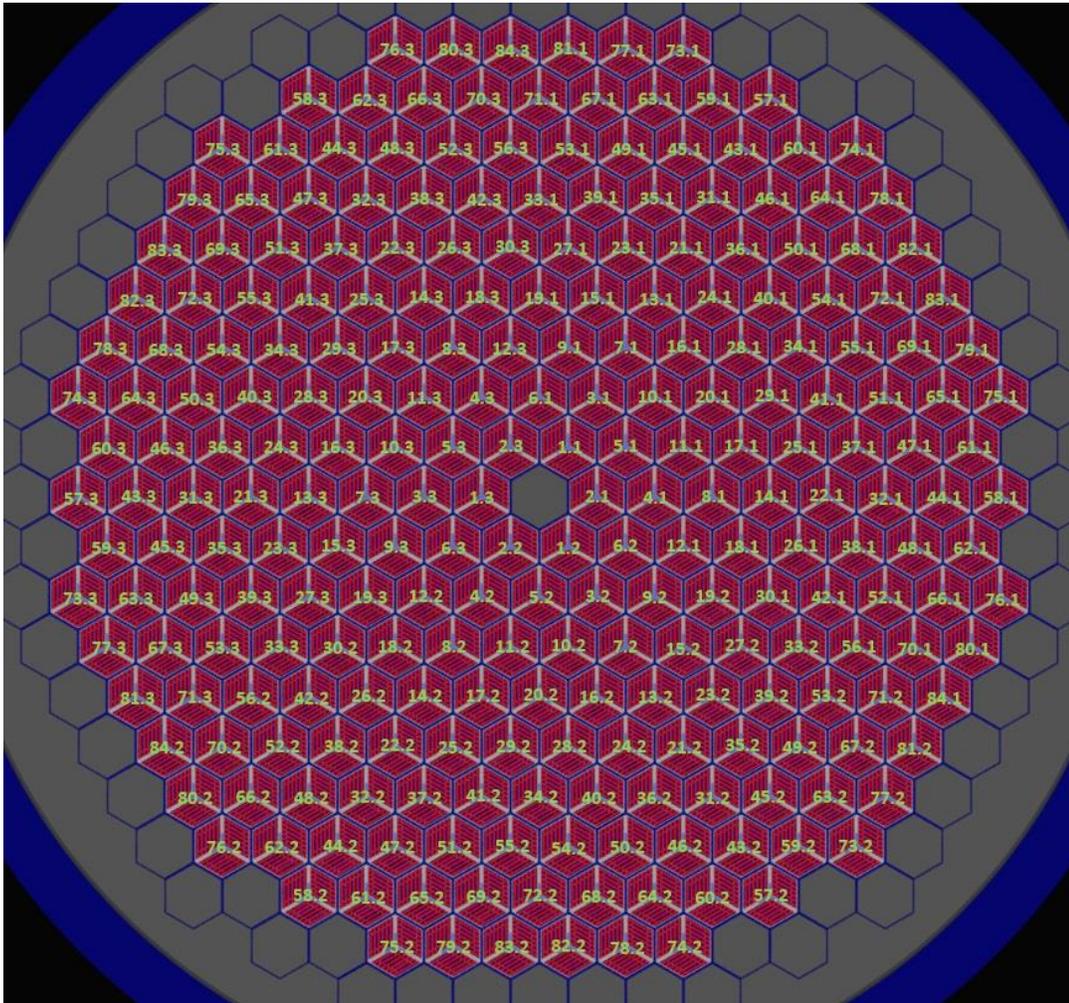


Figure 2.2. Radial assembly layout of 252 fuel assemblies. Assemblies are shown to be in 84 groups of three for corresponding one-third symmetric locations. Note that the geometry shown was obtained by rotating the reference geometry (shown in Figure 1.2) clockwise by 30° for ease of viewing.

Existing tallies in SERPENT which rely on physical region-specific partitioning (which includes material-, cell-, and universe-specific tallies) have been shown to drastically slowdown full core simulations of the AHTR core. However, geometric-specific tallies only slowdown simulations by a modest amount. Preliminary studies on the 3D AHTR model showed that tallying fission energy deposition in about 4,000 fuel regions increased transport calculation times by a factor of about twenty for the three physical tallies, whereas using a built-in hexagonal superimposed geometric tally only increased the runtime by about 0.2%; a trivial amount. Since a major objective of this dissertation is to remain tractable, one solution is to make a superimposed geometric tally capable of handling the one-third assembly regions of the AHTR core. This work creates a new triangular tallying scheme in the Serpent transport code, which increases resolution and efficiency compared to other options currently available to users. The new tally is based upon the existing hexagonal superimposed geometric mesh tally and modifies it by partitioning it into triangular elements. In Figure 2.1, this will serve as the means by which “Extract Power Distribution from Serpent” is achieved. Additionally, this new capability could improve studies of other advanced reactors using hexagonal elements in a similar way that moving from whole- to quarter-assembly granularity of tracking improved the spatial resolution of results in Light Water Reactor (LWR) analyses. All research objectives will be integrated at the end to conduct a fuel cycle analysis of the AHTR which will provide insight into the success of the work.

2.5 Overview of Dissertation

The proceeding chapters focus first on the development of multiphysics model components then summarize results found from using the model. Below is a short summary what to expect from the remainder of the dissertation:

- Chapter 3 - thermal expansion methodology, which will first address assumptions made and then the treatment of each individual component of the geometry.

- Chapter 4 - TH methodology details how heat is modeled to flow from the TRISO fuel kernels to the FLiBe coolant channels and all temperature distributions in between.
- Chapter 5 - development of a novel triangular mesh tallying feature within Serpent to increase the speed of transport calculations. Implementation helps achieve the objective of keeping the model tractable.
- Chapter 6 - implementation of a CB movement scheme within symmetric assembly groups. An example movement schedule is also presented which showcases the capability of the methodology. Will also include a discussion on depletion and solution stability, which can be challenging for large, loosely coupled systems such as AHTR since peripheral areas of the core are relatively far removed from each other.
- Chapter 7 – documentation of the C++ script. Summarizes the overall flow during execution and offers detailed documentation on how the user can interface with the code by means of a user input file.
- Chapter 8 - results by means of integrating (multiphysics coupling) the individual components of the dissertation as they function together as in Figure 2.1 to show how the AHTR core model performs over an example fuel cycle.
- Chapter 9 - conclusions to summarize the significance of results and give insights into how potential future works would best further the efforts put forth in this dissertation.

CHAPTER 3

THERMAL EXPANSION METHODOLOGY

This section details the logic and implementation of thermal expansion in the 3D full-core AHTR Serpent model. To begin, some model specifics and their room temperature (293 K) dimensions are shown in Table 3.1. While the reference design of the AHTR lists that the active core height is 550 cm, the active core height used in the Serpent AHTR model is slightly larger (550.02976 cm). This dimension spans 5,936 TRISO particle pitches axially and the difference is very small (about 0.005%, which is likely below the manufacturing tolerance). It is assumed that this will have a negligible impact on the results relative to other assumptions made with the model.

Table 3.1. Parameters of AHTR model (at 293 K).

Parameter	Value
Number of Axial Sections	16
TRISO Particles in Each Axial Section	371
TRISO Particle Vertical Pitch	0.09266 cm
Height of Each Axial Section	34.37686 cm
Total Height of Active Core	550.02976 cm

3.1 Assumptions and Justifications

There are adjustments and trade-offs which must be made in the geometry to accommodate complex expansions of multiple components in direct contact with each other. While most of these are small changes, they will be explicitly addressed and discussed here. The guiding reasonings will be summarized in this section.

3.1.1 Free Expansion

Assume free thermal expansion of components unless otherwise stated - many components are in direct contact with other components, which will lead to internal stresses due to thermal expansion. While unrestrained thermal expansion would result in free expansion in all directions, these internal stresses would lead to:

- Less expansion in directions experiencing compressive stress
- More expansion in directions experiencing tensile stress
- More expansion in directions perpendicular to compressive stresses
- Less expansion in directions perpendicular to tensile stresses

Accurately capturing these additional (i.e. higher order) effects could be considered the “next step” in accounting for geometric changes, but the goal of this work is to capture the first-order effects of thermal expansion. Stress-related geometric changes due to thermal expansion are smaller than the changes due thermal expansion itself, so they will be largely ignored.

3.1.2 Integrity of Shapes

Maintain the integrity of geometric shapes – as addressed in the previous point, thermal expansion will induce internal stresses on heated components. Some of these stresses will be compressive. Stress mechanics dictate that bodies (especially thin bodies, such as the outer assembly wrapper) under compressive stress will experience buckling. If buckling were to occur in a component, a rectangular shape would instead need to be modeled as a curvilinear surface, requiring a more complicated modeling approach. However, the buckling for structures of the AHTR assembly is expected to be minimal, with the impact being small to negligible, and is therefore not modeled.

3.1.3 Priority in Expanding Components

If conflict arises between thermal expansion of adjoining components, priority should be given to those closest to fuel - since the fissile material drives the neutron economy, the fuel is given the highest priority over other expanding materials. Most of the conflicts resolved by this assumption are between graphite/graphite or graphite/carbon-carbon structures. If one form of carbonaceous material displaces another, this should have a trivial impact on the results unless the amount of media present would no longer match due to density differences. Even in this case, the

local difference is likely small so long as the total amount of material is conserved, as will be addressed in a future assumption. Previous studies [24] and Table 1.8 have shown that carbonaceous structures far away from the fuel are less impactful to the neutronics than those to the fuel, supporting this assumption.

3.1.4 Axial Expansion

Uniform axial expansion in each modeled axial partition – due to the nature of MC modeling, this is required to keep the memory requirements at a manageable level. Hotter assemblies will have more thermal expansion (including in the axial direction). However, all assemblies are constrained axially by the support plates, which will force the same expansion over the entire assembly height. The average expansion should still provide acceptable accuracy: room temperature is 293 K with average fuel particle matrix temperature rising to 1110 K (change of 817 K) is large relative to variances in local temperatures due to local power peaking. This assumption will be compensated in radial expansion for most assembly components where this would apply, as will be addressed further in the next assumption.

3.1.5 Radial Expansion

One-third assembly-wise radial expansion – despite having uniform axial thermal expansion, it is still desired to capture thermal expansion in the radial direction due to local temperature changes. The total effect will need to be accounted both for what is expected from free expansion as well as to compensate for the fixed axial expansion addressed in the previous point. This will be accomplished by using two separate methods: additional temperature-compensating expansion in the radial directions and density corrections. Most assembly features will employ density corrections, but the TRISO fuel particle lattice will use temperature-compensating expansion.

By using the average axial expansion in each axial core section, local thermal expansion effects are not entirely captured since the z-direction expansion is not assembly-specific. Axial

expansion would be larger for higher temperature regions and lower for reduced temperature regions, relative to the average axial temperature assumed for the average axial expansion. To compensate for this, the radial x- and y-direction free expansion needs to be modified accordingly. This means using a correspondingly larger radial expansion for hotter assemblies and similarly a smaller radial expansion for cooler assemblies, relative to the average assembly temperature and expansion for the axial section. This methodology will only be implemented for the TRISO fuel lattice since its expansion effects need to be captured most accurately.

For structures beyond the TRISO fuel lattice, they will use the fixed axial expansion of the axial section and the radial expansion corresponding to free, unrestrained expansion. This means that unlike in the fuel lattice treatment, there will be no additional compensating expansion in the radial directions. This assumption captures the free expansion prediction of components in the radial direction with the simplification of using the average axial expansion. The net expansion will differ slightly from expectation, but the effect will be compensated instead by a density modification. Hotter sections relative to the axial section average which would require additional radial expansion will instead use a slightly higher density than dictated by free expansion, and similarly cooler sections will use a lower density than dictated by free expansion.

Some small geometric modeling errors will be expected when coupled with the axial treatments above, but the correct general thermal expansion behavior will be retained. The model will use an adequate number of axial sections to remain tractable (since axially uniform in each section) and still capture the effective thermal expansion with respect to all three dimensions as well as density change.

3.1.6 Conservation of Media

Conserve material with corresponding density changes – since the amount of material does not change, corresponding densities must also change accordingly. This requires careful attention especially with respect to components having fixed axial expansion and compensated radial

expansion. Accounting for geometric thermal expansion without conserving material would be worse for results than not accounting for expansion whatsoever, so this assumption is essential for model consistency.

3.1.7 Fuel Assembly Pitch

Axial layer-wise radial inter-assembly thermal expansion can be independent - each axial core layer is composed independently, so there is no constraint for the inter-assembly pitch between assemblies to be axially uniform. This will allow for larger inter-assembly spacing at the top of the core than at the bottom of the core, or “flowering”. A visualization of the axial staggering treatment used in the Serpent model can be seen in Figure 3.1. This staggering (Figure 3.1.c) differs from how axial flowering would physically occur, which is to have the assemblies remain in-line but vertically offset (Figure 3.1.b). Both account for axial differences in radial fuel assembly pitch which are not present in the case of assuming average/uniform radial assembly pitch (Figure 3.1.a) in all axial sections. Using axial staggering is more conducive for a MC code since each axial section is still a vertical prism and each section can simply be “stacked” on top of the other sections. One can likely infer that axial section interfaces using the staggering methodology will have disjoint structures due to the shifting. For example, the fuel plates of one section will not align with the fuel plates in the sections above and below it. The differences are expected to be small, and although unphysical, still closely match reality.

This behavior is bound at the top and bottom by the thermal expansion of the axial support plates according to inlet and outlet conditions. During operation, this is nominally set to 650 °C and 700 °C, respectively. These temperatures are easy to modify in the script in order to conduct zero power isothermal tests (outlet same as inlet temperature), model accident conditions (elevated temperatures above the nominal conditions), and any other profile. It should be noted that core flowering is an important behavior to capture for fast reactor designs where geometric changes and leakage are more relevant than in large thermal systems such as LWRs. While the impact is not

expected to be large when applied to AHTR since it is a large thermal reactor, the methodology would be relevant if applied to a different reactor design in the future. The option to toggle axial staggering is implemented in the code, which allows the user to select the average fuel assembly pitch expansion (Figure 3.1.a) given by the average of the top and bottom axial support plate expansion over the entire fuel assembly height to retain the prismatic nature of the fuel assemblies.

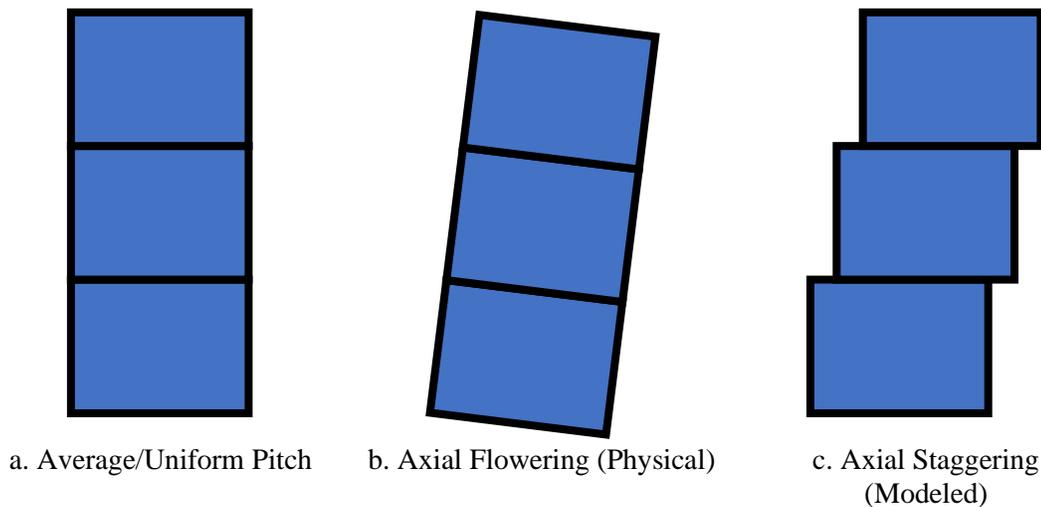


Figure 3.1. Visualization of different fuel assembly radial pitch treatments in axial sections. Deviations away from vertical in b. and c. are exaggerated for visual effect and would be relatively less in an actual implementation.

3.1.8 Components Radially Beyond the Active Core

Geometric features radially beyond those of individual fuel assemblies, removable reflector assemblies, and the permanent reflector will use inlet thermal parameters for thermal expansion - components this includes: boron carbide layer, core barrel, downcomer channel, vessel liner, and pressure vessel wall. These components experience less axial heating (from conduction through the permanent reflector and gamma heat from the fuel) than the active core as well as receive cooling via the downcomer with inlet temperature coolant, so axial differences in thermal expansion are expected to be small. Additionally, these components are outside of the active core and would have a lesser impact on the results. Thermal expansion for these components is

considered at inlet conditions for all axial partitions. For zero-power studies “inlet”, “average”, and “outlet” conditions are identical, and their naming conventions are interchangeable.

3.2 Treatment of Particles

This subsection focuses on thermal expansion of particle (fuel and poison) components of the active core. Physical properties assumed in the model are provided along with derived equations.

3.2.1 TRISO Particles

The TRISO particles used in AHTR are AGR-2 UCO type [10], which were first used for advanced gas-cooled reactors. The particles consist of five regions: fuel, buffer, inner pyrolytic carbon, silicon carbide, and outer pyrolytic carbon. Particles are placed in a graphite matrix material (teal), as can be seen in Figure 3.2. Properties of the regions are shown in Table 3.2.

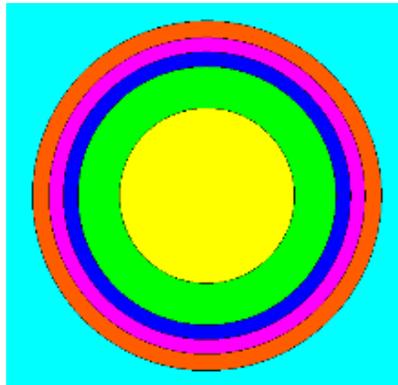


Figure 3.2. TRISO particle used in AHTR.

Table 3.2. Dimensions and properties of TRISO particle layers.

Region	Color	Radius [cm]	Density [g/cc]	α ($\times 10^{-6}$)	Thermal Conductivity ² [W/(m · K)]
Fuel	Yellow	0.02135	10.9	7.6 [25]	3.7
Buffer	Green	0.03135	1	5.5 [26]	0.5
IPyC	Blue	0.03485	1.9	5.5 [26]	4
SiC	Purple	0.03835	3.2	5 [26]	16
OPyC	Orange	0.04235	1.87	5.5 [26]	4

² Thermal conductivities are listed for nonirradiated media. Over the course of core residency, the thermal conductivities of components (perhaps most notably graphite) are expected to decrease. Density change induced by irradiation is a similar phenomenon. There is an experimental capability to apply a fluence-dependent functionality to change the graphite thermal conductivity and density over the fuel cycle, but it did not receive extensive testing as part of this work.

TRISO particle fuel kernels are dense and surrounded by a low-density, porous buffer region. Therefore, fuel will expand according to free expansion. This means that the fuel radius is expanded simply by Equation 3.1. Fuel density is also adjusted to preserve the original amount of fuel material in Equation 3.2.

$$R_{Fuel}^{Hot} = R_{Fuel}^{Cold} (1 + \alpha_{Fuel} \cdot \Delta T_{Fuel}) \quad (3.1)$$

$$\rho_{Fuel}^{Hot} = \frac{\rho_{Fuel}^{Cold}}{(1 + \alpha_{Fuel} \cdot \Delta T_{Fuel})^3} \quad (3.2)$$

The next TRISO layer after the fuel is a carbonaceous buffer region of porous graphite. Since the buffer is low in density, it is assumed that it will accept all thermal expansion inside and outside of it. For this reason, after the free expansion of the fuel kernel, all other TRISO layer thermal expansion and density changes will work inward from the graphite matrix to the buffer. Dimensions will be dictated by corresponding free expansion densities (similar to Equation 3.2), with the exception of the buffer itself.

TRISO particles are embedded in a graphite matrix material. As the matrix expands, it is expected that the “hole” which a TRISO particle occupies will grow according to free expansion of the graphite matrix. Hence, this will dictate the outer boundary of particle (Equation 3.3).

$$R_{OPyC}^{Hot} = R_{OPyC}^{Cold} (1 + \alpha_{Matrix} \cdot \Delta T_{Matrix}) \quad (3.3)$$

Interior TRISO layers will need to account for the expansion of the previous layer before conducting their own. All layers surrounding the kernel are done in the same way, and the SiC outer radius (OPyC inner radius) is given as an example below in Equation 3.4.

$$R_{SiC}^{Hot} = \sqrt[3]{(1 + \alpha_{OPyC} \cdot \Delta T_{OPyC})^{-3} (R_{OPyC}^{Hot})^3 - ((R_{OPyC}^{Cold})^3 - (R_{SiC}^{Cold})^3)} \quad (3.4)$$

Essentially, the inner layer's thickness is thermally expanded by starting where the previous region expanded, preserving the amount of initial material, and assuming the density changes as predicted by free expansion for non-buffer materials. Since the buffer region has its inner and outer radii forced by the expansion of other TRISO layers, its density will need to be adjusted accordingly to preserve the initial amount of material (Equation 3.5).

$$\rho_{Buffer}^{Hot} = \rho_{Buffer}^{Cold} \left(\frac{R_{Cold,Buffer}^3 - R_{Cold,Fuel}^3}{R_{Hot,Buffer}^3 - R_{Hot,Fuel}^3} \right) \quad (3.5)$$

While the other TRISO layers will decrease in density (as expected from free thermal expansion), the buffer will become denser (compress) since it will accommodate the thermal expansion of the other layers.

3.2.2 Europium Oxide Spheres (Burnable Poison)

Europium oxide spheres are placed 3.9744 cm (cold dimension) apart in the center of the fuel planks. The distance used here differs slightly from that specified in the original design documents which somewhat arbitrarily choose 4 cm. The important behaviors to capture are: firstly to expand the particles (thermal expansion coefficient assumed to be 7.5×10^{-6}) [27] and secondly to ensure the particles remain in the center of the fuel planks. As the planks will also be expanding, attention will need to be paid to making sure the europium spheres are not skewed toward either fuel stripe. One consequence which will be addressed here is the amount of graphite displaced due to the thermal expansion of the BP spheres (since poisoned particle expands more than graphite). The density of the meat graphite at the center of the plank could be adjusted to accommodate this change as well as others. However, this density correction will be relatively small since the BP spheres only occupy about 2.4×10^{-4} of the volume in the large meat graphite region. As an example, for the assumed average fuel matrix temperature of 1110 K, the relative density increase required from free expansion is about 1.5×10^{-6} for the meat graphite to account for the lost volume. Even

for the hottest assembly section, this difference will likely not exceed 3×10^{-6} . For such a minuscule loss of material, ignoring the loss will have a negligible impact on the results.

3.3 Axial and Radial Treatment for All Remaining Regions

As covered in the assumptions, each axial section of the Serpent model needs to have a uniform height over the entire section. While this height can vary between axial sections (hotter sections will expand more), thermal expansion in each individual assembly still needs to be accounted for. The basis for determining how much axial thermal expansion is experienced in a section is dictated by the average expansion of the TRISO fuel matrix. The logic for selecting the fuel particle matrix is that it is the most impactful region for neutronics to the model. Two methods are used for having a fixed height for assembly components. For the TRISO particles, the TRISO lattice has additional radial expansion to account for the fixed axial expansion. This should still provide assembly-specific expansion despite the fixed section height. For all other geometric features, they expand freely in the radial directions and their densities compensate for the forced axial uniformity.

3.3.1 TRISO Fuel Stripe Matrix

The TRISO fuel stripe matrix is the region determining axial thermal expansion. While the fuel matrix size will be variable and can be specified by the user as part of the script, the default sizes and dimensions are given as: the TRISO particle lattice width is four particles with pitch 0.09128 cm cold (total cold width: 0.36512 cm) and the TRISO particle lattice length is 202 particles with pitch 0.09406 cm cold (total cold length: 19.00012 cm). For thermal expansion purposes, the TRISO particle pitch and their graphite lattice are considered to expand as if entirely made of graphite. The carbonaceous non-fuel layers expand similarly to graphite, so this is likely an acceptable assumption. The thermal expansion coefficient of graphite is assumed to be $5 \times 10^{-6} \text{ K}^{-1}$.

The proposed method for accounting for assembly-specific thermal expansion while also fixing the axial expansion is to have additional compensating expansion in the radial directions. Hotter sections will see slightly additional radial expansion and colder sections will see less radial expansion. The reasoning behind this is to hold the axial expansion constant (within an axial section) while still capturing the correct density change and maintaining initial cold loading mass balance. The assemblies below the average temperature will see less lattice change in the x- and y- directions since the average (hotter) axial pitch is used, and vice versa for the hotter assemblies. Logically, the volume of the expanded TRISO unit cell needs to be the same between the free and fixed-height methods. In order to achieve this, the thermal expansion in the non-vertical directions needs to be corrected (say, with a factor β). For the three pitches p_x , p_y , and p_z :

$$p_x \beta * p_y \beta * p_z (1 + \alpha \overline{\Delta T}) = p_x p_y p_z (1 + \alpha \Delta T)^3 \quad (3.6)$$

$$\beta = \sqrt{\frac{(1 + \alpha \Delta T)^3}{(1 + \alpha \overline{\Delta T})}} \quad (3.7)$$

The relative impact in each radial dimension is small. Using an average height for all assemblies in an axial section is possible by mildly perturbing the radial expansions from their expected free expansion results. This allows for ease of modeling with MC while maintaining the ability to account for thermal expansion on an individual assembly basis, while preserving both the individual expansion behavior as well as the total amount of material.

3.3.2 Fuel Planks

Fuel planks (Figure 3.3) are 2.55 cm (cold) wide and 23.1 cm (cold) long. However, fuel planks are embedded into the structural components on both ends (the Y-shape at the center and the wrapper on the periphery). The actual coolant channel width for the planks is about 22.1125 cm (cold). The logic is that the planks will expand freely in the width of the channel (plank short dimension) according to an average plank temperature (fuel stripe and central meat) but will expand

in length according to the slightly cooler expansion dictated by coolant. The reasoning is that in the channel, planks will expand freely, but the length expansions will be limited by the structural components of the assembly. Essentially, the fuel plank will facilitate the expansion of hotter (near-fuel) geometric features and cooler (structural) ones. Great attention will need to be paid to preserve the initial mass of graphite in the fuel planks to account for: TRISO lattice expansion, BP particle expansion, and conflicting expansions stemming from different temperatures. Since the plank is graphite, the thermal expansion coefficient is assumed to be $5 \times 10^{-6} \text{ K}^{-1}$.

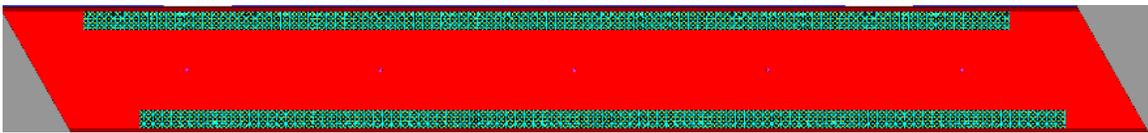


Figure 3.3. Fuel plank (red) embedded in structural graphite (gray).

3.3.3 Fuel Plank Spacers

The fuel plank spacers (white structures in Figure 3.4) are graphite, so they have a thermal expansion coefficient of $5 \times 10^{-6} \text{ K}^{-1}$. Structurally, they ensure that fuel planks remain apart from each other, the assembly wrapper, and the central Y-structure. From a thermal expansion point of view, they will be used as a catch-all to reconcile different temperatures within the AHTR fuel assembly. A fuel assembly will want to expand at the temperature of the structural graphite components, which will be assumed to be at the coolant temperature due to the relatively high thermal conductivity of both graphite and FLiBe. However, fuel planks will be at a higher temperature (due to heat generation in the fuel stripes) and will expand more in the channel direction (but restricted laterally by the structural wrapper and Y-shaped central region). This means that while the width of each fuel section will expand (due to structural component expansion), the coolant channels will have relatively less expansion (due to larger degree of expansion in fuel planks). Since the fuel plank spacers are in these expansion-restricted channels, they too will expand less than what would be expected from free expansion.

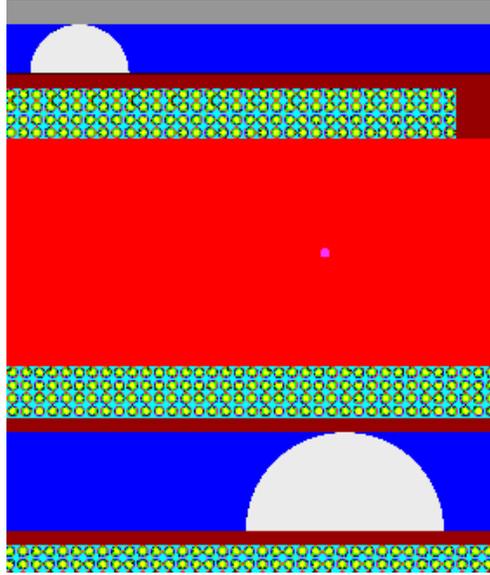


Figure 3.4. Two sizes of plank spacer are used in the model, depending on fabricated channel width.

There are three considered options for spacer expansion:

- Maintain integrity of shape, have spacer fit the channel width despite restriction. Use corresponding expansion density instead of free expansion density.
- Purposefully undersize spacers as part of the cold dimension definition so that they fit well under normal operating conditions.
- Have spacers expand normally in size but have them truncated at the tangent interface so that there is a small flat “edge” due to deformation from stress. Correspondingly increase the density from the free expansion prediction to account for volume lost to this deformation.

The first bullet aligns well with assumption 3.1.2 in maintaining the integrity of shapes, but the density would be farther from free expansion (assumption 3.1.1) than the third bullet point. The second bullet point would satisfy both of these assumptions, but from a design standpoint would require the fuel assemblies to be “loose” relative to each other at temperatures below normal operation. Further, this process is not easily extendable over a 3D system (axially dependent temperature profile, so axially dependent expansion within an assembly) and not feasible over a

depletion cycle (power profile changes over cycle, so local expansion also changes). Essentially, the second option would likely default to the first option for ease of implementation. The third bullet defies assumption 3.1.2 for shape integrity, even though the spacers could still be modeled as half cylinders but just have a small region cut away. The density modification needed to capture this treatment will be closer to free expansion than that resulting from the first option. Logically, this implementation makes sense from the standpoint of that if large (greater than 20 cm long) fuel planks and their relatively small spacers (less than 1 cm) experience compressive stress due to thermal expansion of components at different temperatures, it would be the spacers that would have to accommodate this. This stress would concentrate at the spacers, and specifically at the smallest interface (the tangent one versus the one with a 1.4 cm cold dimension for the large spacers, 0.7 cm cold dimension for the small spacers). This stress would cause deformation in the spacer, resulting in the behavior suggested by the third option. This work implements the treatment described in the first bullet for all simulations going forward, although any of these methods would have likely yielded comparable results.

3.3.4 Control Blade

The CB is Y-shaped and is essentially formed by three rectangular sections rotated by 120°. Each section is 1 cm wide and 10 cm long. The blade itself is not in direct contact with any other solid assembly component, so it is free to expand into the surrounding FLiBe coolant. Since it is free to expand, the expansion process is straightforward. The CB is composed of MHC. It is possible that there could be some thin cladding around the CB, but no previous reports have referenced one so this work will also assume the CBs to be unclad. The thermal expansion coefficient of MHC is assumed to be $7.5 \times 10^{-6} \text{ K}^{-1}$ [28].

3.3.5 Structural Y-Shape

Unlike the graphite used in other carbonaceous structures of the assembly, the structural Y-shape (Figure 3.5) is made of carbon-carbon (C-C) composite. Literature suggests that the

thermal expansion coefficient of some C-C composites are larger ($8 \times 10^{-6} \text{ K}^{-1}$) [29] than that of graphite ($5 \times 10^{-6} \text{ K}^{-1}$), but there is a spectrum of possible values based on the exact composition selected. Depending on the specific type of either C-C composite or graphite, the thermal expansion coefficient can obviously vary. For simplicity and consistency within the model, $5 \times 10^{-6} \text{ K}^{-1}$ will be used over the larger value suggested by literature. This will simplify the thermal expansion process and should not be considered too large of a deviation due to the amount of variance related to graphite expansion, C-C composite expansion, and exact compositions of carbonaceous materials used in the AHTR system.

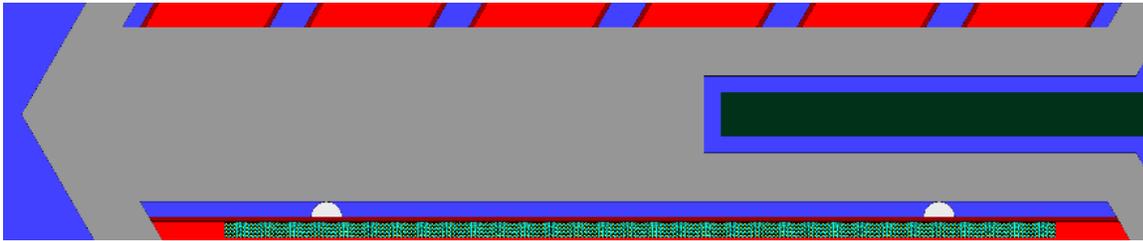


Figure 3.5. Structural Y-shape (gray). Fuel planks not shown as embedded to better highlight this structural feature.

3.3.6 Assembly Pitch

Axially, assemblies are bound above and below by the support plates. Two options exist here (so both are coded and can be toggled using an option in the model) and are driven by the expansion of the axial support plates. The first option allows for axial “staggering” of assemblies (Figure 3.1.c) where the cooler bottom plate expands less than the hotter top plate. This results in a larger assembly pitch near the top of the core and means that each axial partition’s pitch can be modeled as a linearization between the two depending on location. The axial tilt of the assemblies (Figure 3.1.b) would be neglected in this treatment and each axial partition would be modeled as being vertical but radially staggered from the same assembly sections above and below it in other axial partitions. The second option uses the average (upper and lower) support plate expansion to change the assembly pitch of all axial partitions by the same amount (Figure 3.1.a). This method assumes that assemblies remain vertical in the core.

3.4 Extra-Assembly Components

This section will focus on geometric features beyond those found within an assembly. It is expected that changes to these features have much smaller impact on core physics than the previous ones in the active core, so the treatment might be slightly more lenient/approximate if needed.

3.4.1 Radial Layers Just Beyond the Core

These layers include the permanent radial reflector, boron carbide layer, and core barrel (Figure 3.6). Both the permanent radial reflector and the core barrel are composed of C-C composite (thermal expansion coefficient of $5 \times 10^{-6} \text{ K}^{-1}$). Boron carbide has an expansion coefficient of $5.65 \times 10^{-6} \text{ K}^{-1}$ over the range 285-1213 K [30], which should serve the model adequately.

Since the assembly lattice of the core expands radially into where the cold permanent radial reflector would be, it will be assumed that the permanent radial reflector will expand in a similar manner to the assembly lattice to accommodate the expansion of the lattice. The peripheral removable reflector assemblies will have the same interassembly gap width with the permanent radial reflector as used for all the other assemblies in the axial section. This will establish the inner boundary for the permanent radial reflector. It is assumed that the permanent radial reflector will expand freely at the average coolant temperature for the axial section. The outer radius of the permanent radial reflector will then be prescribed by material conservation based on these two constraints.

The two outer layers (boron carbide and core barrel) are thin and relatively close to the downcomer coolant (at inlet temperature). These will expand at inlet temperature. In a process similar to the expansion of TRISO particle layers between the buffer and the graphite matrix, the material density of these two layers will be assumed to be of free expansion density and the radius of each region will be changed (similar to Equation 3.4) to conserve material. This process will start with the outer radius of the permanent reflector and continue forward for both the boron carbide then core barrel regions.

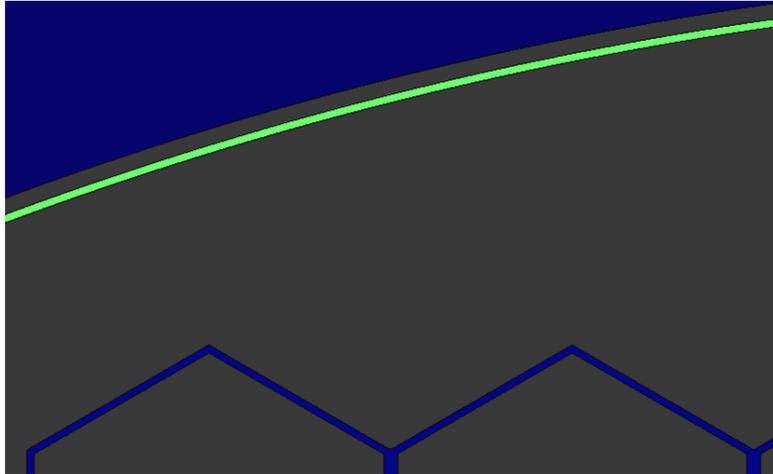


Figure 3.6. Portions of removable reflector assemblies, permanent reflector region, boron carbide layer, and core barrel.

3.4.2 Radial Layers Beyond the Coolant Downcomer

These layers include the pressure vessel inner liner and the pressure vessel wall (Figure 3.7). The vessel liner is composed of Hastelloy N alloy, which was found to be compatible with FLiBe during the MSRE. The reactor pressure vessel is composed of Incoloy alloy 800H. Both layers are considered to be at the inlet temperature.

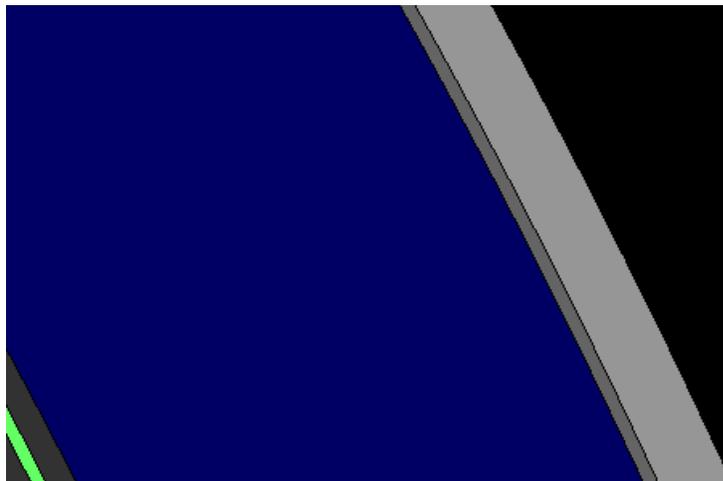


Figure 3.7. Downcomer, pressure vessel inner liner, and pressure vessel wall.

At 650 °C, the thermal expansion coefficients are $13.6 \times 10^{-6} \text{ K}^{-1}$ for alloy N and $17.3 \times 10^{-6} \text{ K}^{-1}$ for alloy 800H. Since these expansion coefficients differ but the two materials are bonded

together, it will be assumed that the thicker layer (alloy 800H) will expand freely, and that the liner will have its correct free expansion density with inner radius forced to achieve this. Since the reactor would want to be thermally isolated, heat loss from the pressure vessel will not be considered.

3.4.3 Lower and Upper Axial Reflectors

Directly above and below the active core, there are 25 cm tall axial reflector sections. Geometrically, these are identical to the active core but there are no fuel stripes nor BP particles. The TRISO fuel stripes and BP spheres are simply replaced with graphite. For the lower axial reflector, all assemblies have inlet conditions (nominally 923 K for normal operation). Thermal expansion for this layer should be very straightforward as it has a similar process to the active fuel regions, except there will be no need for radial assembly-specific expansion (i.e. radially uniform, all of them at inlet temperature). For the upper axial reflector, assembly-specific thermal expansion will be used which assumes outlet conditions from each top axial active core section.

3.4.4 Bottom and Top Support Plates

The bottom support plate is a 35 cm thick section at the bottom of the model. It is modeled as a homogenous mixture of graphite (21.1% by volume) and FLiBe (78.9% by volume). Since it is at the inlet, it is assumed that the temperature is simply the inlet temperature (nominally 923 K for operating conditions). Due to the bottom support plate being homogenized, there are no physical structures requiring thermal expansion. At 923 K, liquid FLiBe has a density of 1.9628 g/cm³ and the graphite is assumed to be structural C-C composite (density 1.95 g/cm³ at 293 K, 1.9317 g/cm³ at 923 K). The density of the homogenized bottom support plate is then 1.9562 g/cm³.

The top support plate is a 35 cm thick section at the top of the model. It is assumed to be a homogenous mixture of SiC (57% by volume) and FLiBe (43% by volume). Since it is at the outlet and there is no space for mixing of assembly coolant before this point, there would be a temperature profile within the top support plate. However, it is assumed that the temperature is simply the average outlet temperature (nominally 973 K for operating conditions). This assumption should be

acceptable since the region is away from the active core and thus less impactful. Due to the top support plate being homogenized, there are no physical structures requiring thermal expansion. At 973 K, liquid FLiBe has a density of 1.9384 g/cm³ and the SiC has a density of 3.0686 g/cm³ (3.1 g/cm³ at 293 K). The density of the homogenized bottom support plate is then 2.5826 g/cm³.

CHAPTER 4

HEAT TRANSFER AND THERMAL HYDRAULICS METHODOLOGY

FLiBe at operating conditions is a single-phase liquid, which makes TH more straightforward than systems like LWRs, which must deal with bulk or subcooled boiling. The FLiBe of AHTR will not boil under any normal conditions since the boiling point is 1430 °C (higher than the peak plate temperature) at atmospheric pressure. Therefore, single phase correlations are used with FLiBe physical data to obtain coolant properties within the core.

In the AHTR assembly design, there are two 0.35 cm and five 0.7 cm coolant channels in each one-third assembly section. Each 0.35 cm channel cools one fuel stripe whereas each 0.7 cm channel cools two fuel stripes. Effectively, if a line of symmetry is drawn down the middle of each 0.7 cm channel, every fuel stripe sees a 0.35 cm coolant channel. Since this work only resolves temperatures and thermal performance down to the one-third assembly level, there is no need to uniquely track each coolant channel and fuel stripe. Accordingly, all channels are evaluated using their respective one-third assembly section average values. The difference between two halves of a 0.7 cm channel and a single 0.35 cm channel, which has an interface with the structural material will be assumed to be small.

Heat conduction through graphite structures is based on analytic one-dimensional (1D) heat transfer equations. The 1D slab under consideration can be seen in

Figure 4.1. The fuel plates of the AHTR are composed mostly of graphite with TRISO particles dispersed in fuel stripes. A half-width channel consists of a 0.35 cm FLiBe coolant channel, a 0.1 cm section of graphite, and a 1.175 cm section containing both the TRISO fuel stripe layers in graphite matrix as well as central plank graphite.

In order to develop a heat transfer model for multiphysics capabilities, it is necessary to find a temperature profile for the plank and coolant channel. The general process implemented for doing this can be summarized by the following steps executed for each assembly section:

- Convert fission energy deposition from neutronics results to a volumetric heat rate.
- Obtain coolant temperature (from using the heat rate and the temperature of coolant axially below the section of interest) as the radial profile boundary condition.
- Homogenize fuel plank to obtain an effective thermal conductivity coefficient through the fuel stripe to model as 1D with uniform volumetric heat rate.
- Solve 1D heat transfer for the plank with the homogenized fuel stripe.
- Obtain average temperatures for coolant and graphite (not fuel stripe) structures.
- Recover discrete nature of TRISO particles by superimposing a heat rate modulation function over the homogenized solution.
- Obtain an average effective fuel kernel temperature.
- Update temperature-dependent thermal conductivities and iterate through prior steps.
- Implement temperatures in thermal expansion model to update dimensions and densities.
- Use converged temperatures, dimensions, and densities in the next iteration of neutronics.

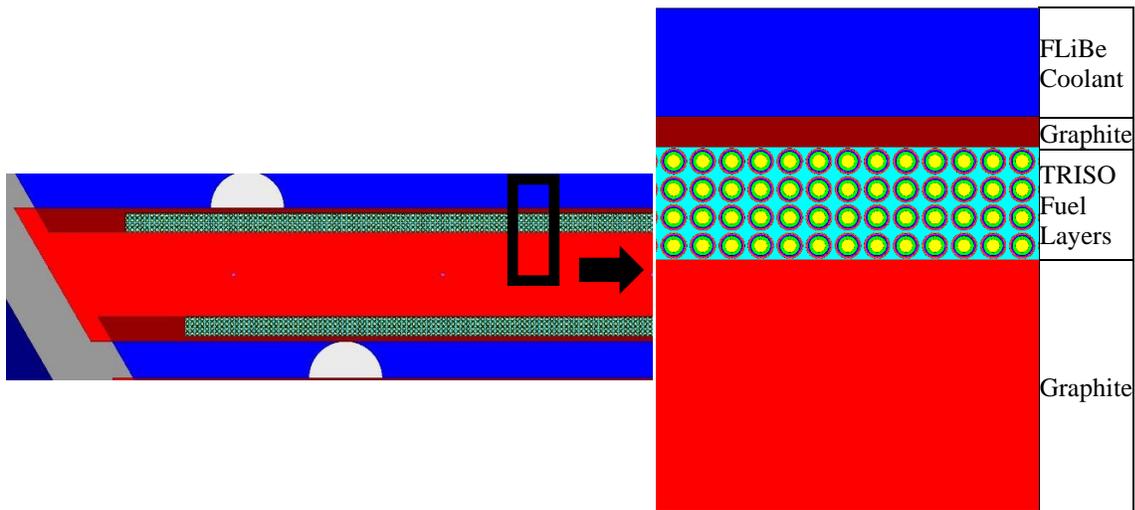


Figure 4.1 AHTR half-plank channel with individual features identified.

4.1 Fuel Stripe Homogenization

A TRISO particle and fuel stripe matrix homogenization process was developed and applied to heat transfer considerations for conduction through the fuel plank. In contrast, all neutronic models use explicit TRISO particles. TH will later return to a heterogeneous treatment after finding the average temperature profile.

Analytically computing the heat transfer across the fuel plank is challenging due to the 3D spherical shapes of the TRISO particles (Figure 3.2) in an otherwise simple geometry which can be approximated as a 1D slab. One simplification to remedy this issue is to homogenize the TRISO particles into the stripe matrix graphite. This reduces the plank heat transfer problem to a 1D Cartesian system, which has well-known analytical solutions.

The TRISO homogenization method implemented in this work is Maxwell's method [31]. It finds the thermal conductivity of a mixture (k_m) by averaging the thermal conductivities of spherical discontinuous particles (k_d) suspended in a continuous medium (k_c) based on the phase volume fraction of the discontinuous particles (P_d), as expressed in Equation 4.1.

$$k_m = \frac{k_c[k_d + 2k_c - 2P_d(k_c - k_d)]}{k_d + 2k_c + P_d(k_c - k_d)} \quad (4.1)$$

When the above equation is applied to the TRISO particles from an in-out approach working from the fuel kernel to the graphite matrix, a core-average fuel stripe thermal conductivity of 9.967 W/(m·K) is obtained, as summarized in Table 4.1. Column “ k_m ” is for the iterative layer-wise results from using the in-out approach, with the last row showing the value of the whole fuel stripe homogenization. The fuel and matrix graphite thermal conductivities shown here are conservatively low values for higher temperatures or high fluence, respectively. These values can be modified by the user and can even change with fluence within the code implementation, but sample values are presented for an illustrative calculation of the fuel stripe homogenization process.

Table 4.1. Parameters used to thermally homogenize TRISO particles in fuel stripe for fuel assembly under core-average operating conditions with conservative fuel and graphite thermal conductivities.

Material	Thermal Conductivity [W/(m·K)]	Diameter [cm]	P_d [-]	k_m [W/(m·K)]
Fuel	3.7	0.0427	0.316	3.700
Buffer	0.5	0.0627	0.728	0.911
IPyC	4	0.0697	0.750	1.582
SiC	16	0.0767	0.743	4.303
OPyC	4	0.0847	0.400	4.224
Matrix Graphite	15	0.11496 [†]		9.967

[†]Equivalent radius of a cuboidal TRISO lattice element assuming a packing fraction of 0.4.

4.2 1D Heat Transfer Across Fuel Plate and Coolant Channel

With the fuel particles homogenized into the fuel stripe graphite matrix, it is possible to conduct a 1D analysis on the fuel plank to obtain the average temperature profile. The heat transfer solution steps include:

1. Convection from FLiBe coolant (boundary condition) to the plank wall.
2. Conduction through the 0.1 cm section of graphite separating coolant from fuel stripe.
3. Conduction through the fuel stripe, assumed to have a constant volumetric heat generation rate (due to homogenization process).
4. Conduction through central graphite region (assumed adiabatic due to symmetry).

The solution process proceeds in the opposite direction than heat is expected to flow (cooler coolant to hotter plank). This is because the coolant temperature is the boundary condition for the solution. At the inlet, the coolant temperature is well-known (since inlet conditions are prescribed). As coolant flows up through the core, it is heated in each fuel assembly section which has its own specific heating rate. Under steady-state conditions, all this heat must be rejected to the coolant if axial conduction in structural components is neglected. Prior AHTR sensitivity studies have shown that the gradient in the axial direction is about 1000 times lower than that in the radial direction with the peak temperature difference in structural components being 0.003 °C and on the order of 10^{-4} °C for coolant temperature differences [23], which will be considered as ignorable for this work. If all the heat produced in a steady-state section is assumed to move to the coolant, the bulk

heating of the coolant in each assembly section is well-characterized and can be analytically solved. Moving backward with this known coolant temperature and the same specific heating rate, the average temperature profile of the rest of the assembly section can be found.

Within each coolant channel, the channel spacers separate the coolant into three non-connected volumes. This work will assume that the coolant in each is at the same temperature. Due to the high thermal conductivity of graphite and FLiBe, heat would transfer well between the two sides of a channel spacer. Additionally, is it not unreasonable to anticipate that future design modifications might add structures comparable to mixing veins to the spacers to promote cross-mixing between the separated coolant volumes. This would allow for enhanced cooling within the assembly while still fully benefiting from the structural integrity gained from using the channel spacers. While the coolant volume displaced by the channel spacers still needs to be accounted for when calculating channel parameters for heat transfer purposes, spacers will otherwise be neglected in 1D heat transfer calculations.

So far, the TH process has only discussed the coolant inside the intra-plank channels. There is additional coolant in the CB slot as well as between assemblies in the interassembly gap. The volumetric proportions of each coolant region are shown in Table 4.2, depending on whether the CB is withdrawn or inserted (since this would displace coolant volume in the CB channel). One can see that there is a significant (almost half of the total) volume of coolant beyond that contained in the intra-plank channels. According to reference documents from ORNL, 95% of the power produced in the fuel assembly is carried away by coolant flow within the fuel assembly [3]. It will be assumed that the fuel assembly orificing will be done in such a way to achieve this despite the differences in coolant volume, essentially allowing for a higher flow rate through the intra-plank channels and lower elsewhere. This assumption, combined with the high thermal conductivity of assembly components (principally FLiBe and graphite), provide support of the approximation to use the same temperature coolant within each one-third assembly axial segment.

The last outstanding consideration is to account for the remaining 5% of power, which will be assumed to be covered by a correspondingly lower enthalpy change in the coolant due the excess volume beyond that of the intra-plank channels. To clarify, the 1D heat transfer method models complete heat transfer to the intra-plank coolant. However, it is assumed that only 95% of this heat moves to the intra-plank coolant and the balance to the remaining coolant volume in the assembly section. Further, it will be assumed that the remaining coolant will be at the same temperature as the bulk intra-plank channel coolant. This final assumption is expected to have a very small impact on results. Previous analyses [32] have shown that the coolant reactivity coefficient is about -1.3 pcm/K, so even if the non-intra-plank coolant were entirely at either inlet or outlet conditions, the maximum system-wide impact of this assumption would be at most 15 pcm. However, the expectation is that the assumption is reasonable, and the true reactivity impact would be even less than this.

Table 4.2. Coolant volume fraction (as percent) by region for cold reference design.

	Intra-Plank Channels	CB Channel	Interassembly Channel
CB Withdrawn	56.30	11.89	31.82
CB Inserted	60.26	5.69	34.05

4.2.1 Convection to Coolant

Heat convection from the fuel plank wall to the bulk coolant has a temperature change given by Equation 4.2.

$$\Delta T = \frac{0.95q''}{h} \quad (4.2)$$

Where q'' is the total heat flux from all TRISO layers in a stripe divided by the area of the plank wall in contact with the coolant. The assumption of only 95% of the heat being produced in the section moves to the intra-plank coolant is implemented here. h is the convection heat transfer coefficient, given by Equation 4.3.

$$h = \frac{Nu \cdot k}{D_h} \quad (4.3)$$

k is the thermal conductivity of FLiBe, D_h is the hydraulic diameter of the coolant channel, and Nu is the Nusselt number. Multiple different correlations are available for computing Nu , but the one used here is the Gnielinski correlation (Equation 4.4).

$$Nu = \frac{(f/8)(Re-1000)Pr}{1+12.7(f/8)^{0.5}(Pr^{2/3}-1)} \quad (4.4)$$

Re is the Reynolds number, Pr is the Prandtl number, and f is the friction factor. The Gnielinski correlation is valid for $0.5 < Pr < 2000$ and $3000 < Re < 5000000$, which is satisfied for the conditions experienced in the AHTR ($Pr=13.525$ and $Re \approx 9000$). The friction factor can be obtained from the Petukhov correlation (Equation 4.5).

$$f = (0.79 \ln(Re) - 1.64)^{-2} \quad (4.5)$$

The Petukhov correlation is applicable over the same Reynolds number region as the Gnielinski correlation ($3000 < Re < 5000000$). The Reynold's number is given by:

$$Re = \frac{\rho u_m D_h}{\mu} \quad (4.6)$$

Where ρ is the coolant density, u_m is the mean flow velocity, and μ is the viscosity. Relevant data required for evaluating parameters in this subsection can be found in Table 4.3, with many of the values driven by the assembly geometry. If the user chooses to change any geometric parameters, the values will correspondingly change. This will allow for future design work to modify the AHTR as desired to obtain updated results. These values are from the reference design assuming room-temperature dimensions, with the option for temperature-dependence implemented within the code.

Table 4.3. Parameters of FLiBe coolant channel.

Parameter	Value
Thermal Conductivity (k)	1 W/(m·K)
Coolant Channel Length	22.112515 cm
Coolant Channel Width*	0.35 cm
Hydraulic Diameter (D_h)	1.35063 cm
Prandtl Number (Pr)	13.525
Viscosity (μ)	0.0056 Pa·s
Density (ρ)	1.95 g/cm ³

*Each fuel stripe effectively sees a 0.35 cm coolant channel as discussed previously.

4.2.2 Conduction through Graphite Plank Sleeve

Since it is assumed that no heat is generated in the graphite (power generated in graphite via neutron scattering is neglected in this work), the heat equation is simple (Equation 4.7).

$$\frac{d^2T}{dy^2} = 0 \quad (4.7)$$

This produces a linear temperature profile (assuming constant thermal conductivity) across the graphite. The temperature change is given by Equation 4.8.

$$\Delta T = \frac{q'' \Delta y}{k} \quad (4.8)$$

4.2.3 Conduction through Fuel Stripes

TRISO particles are homogenized into the fuel stripe graphite for obtaining an average TH profile, and it is assumed that the heat produced by the TRISO particles is uniformly generated in the fuel stripes post-homogenization. This produces a heat equation given in Equation 4.9.

$$\frac{d^2T}{dy^2} + \frac{q'''}{k} = 0 \quad (4.9)$$

This produces a quadratic temperature profile across the fuel stripe. The temperature change is given by Equation 4.10 (for Δy taken with respect for stripe-central graphite interface).

$$\Delta T = \frac{q''' (\Delta y)^2}{2k} \quad (4.10)$$

4.2.4 Central Graphite Region (Adiabatic Boundary)

Since it is assumed that no heat is generated in the central graphite region (power generated in graphite via neutron scattering is neglected in this work), its heat equation is the same as that of the graphite in contact with the coolant (Equation 4.7). This means that the temperature profile in this region is linear. However, there is an additional boundary condition due to the adiabatic condition stemming from symmetry which states that there is no temperature gradient at the boundary (fuel plank centerline). This forces the temperature profile in this region to be constant.

4.2.5 Dependence on the Number of TRISO Layers

This section includes studies considering the impacts of varying the number of TRISO fuel layers used within the fuel stripes. This work assumes that four layers of particles are used in each fuel stripe, with a width of 202 particles. Some of the following cases use a lattice size slightly different from the reference design for ease of modeling for comparison purposes, but the modifications are minor and the results and conclusions should still be applicable.

4.2.5.a Reactivity Impact

One of the implemented capabilities of the C++ script is the ability to easily change the size of the TRISO particle lattice. To demonstrate this capability and show the impact of TRISO fuel stripe size on eigenvalue, four different lattice sizes were considered which conserved heavy metal loading: 4 by 204, 6 by 136, 8 by 102, and 12 by 68. Table 4.4 has the results for both uncontrolled (no BP spheres nor CBs inserted) and BP sphere cases for each of the four TRISO lattice sizes. The number of TRISO layers considered include the physical limit possible for a lattice TRISO arrangement. Along the length of the fuel plank, the number of particles needed to maintain the same heavy metal loading for three (and fewer) layers would be too long to fit within the fuel plank. Along the width of the plank, thirteen layers would overlap into the other fuel stripe and the BP spheres. Thus, the feasible TRISO lattice size is bound between four and twelve layers. This

comment is a moot point for heterogeneously dispersed TRISO particles ,which are not restrained by discrete pitch size, but heterogeneously placed particles are beyond the scope of this work.

Table 4.4. Eigenvalues of cases varying the number of TRISO layers (while conserving total fuel particle loading).

Layers	Uncontrolled		BP Spheres	
	k_{eff}	σ [pcm]	k_{eff}	σ [pcm]
4	1.36751	7	1.05725	8
6	1.39733	7	1.07062	9
8	1.4154	7	1.07709	9
12	1.43808	6	1.08429	9

Both sets of eigenvalues are plotted in Figure 4.2 to show the behavior. As can be observed, the trends are similar: eigenvalue increases with the number of TRISO layers but at a diminishing rate. Note that this study does not account for TH feedback. More TRISO layers would result in higher fuel plank temperatures, which would decrease reactivity.

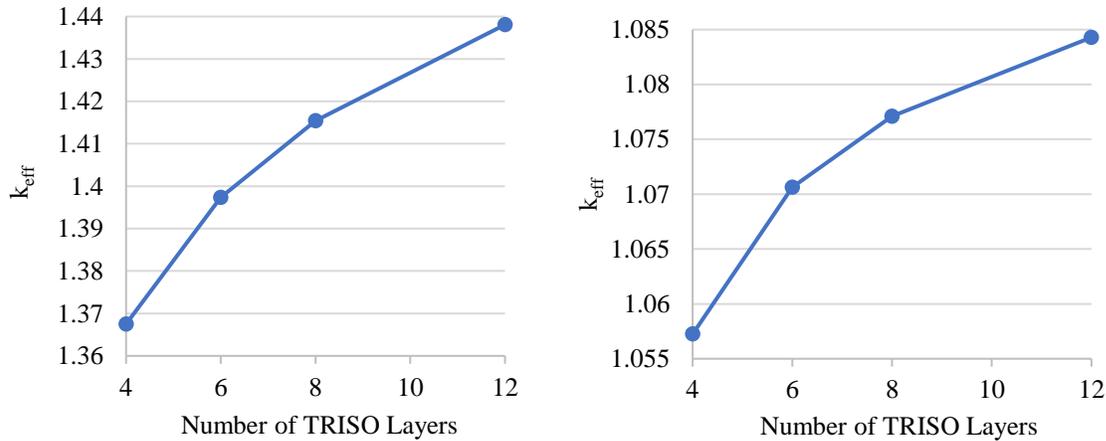


Figure 4.2. Eigenvalue behavior with variable TRISO layer configuration. Left: uncontrolled. Right: controlled with europa BP spheres at the center of the plank.

4.2.5.b Thermal Impact on Fuel Plate Temperature Distribution

As the number of TRISO layers increases, the temperature at the center of plank also increases (Figure 4.3). This makes physical sense because if heat is generated farther from the coolant boundary, then there is more media to conduct through and thus more thermal resistance to

reject heat. Having additional TRISO layers means that the heat from deeper layers must additionally traverse the preceding layers before reaching the coolant boundary.

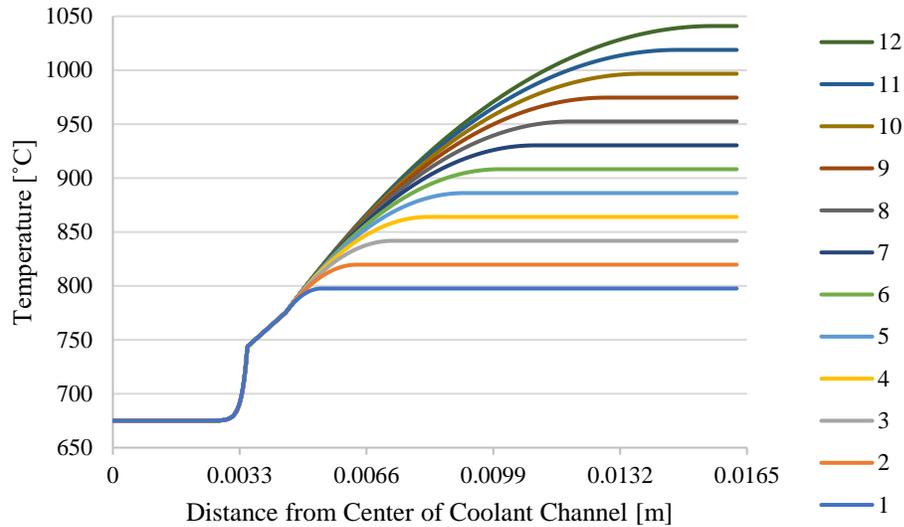


Figure 4.3. Temperature distribution for homogenized fuel stripe depending on the number of TRISO layers. Assumes average core power and coolant temperature.

Figure 4.3 was obtained assuming that the power generation in the fuel stripe is uniform. This assumption of uniformity can be corrected by assuming that the true power profile ((a) in Figure 4.4) can be decomposed to the sum of the now-known constant profile and their difference (“modulation”, (b) in Figure 4.4) [33]. If the difference profile temperature distribution is summed with the average temperature distribution seen in Figure 4.3, then the temperature profile through TRISO particle centers can be seen in Figure 4.5. As expected, in Figure 4.5 the temperature in fuel kernels is higher and the temperature in the surrounding graphite matrix is lower than the temperature profile from the homogenized fuel stripe. For a single TRISO layer, the temperature in the fuel kernel is about 80 °C higher than that expected from the uniform power profile. For the case of four layers (matching the reference AHTR design in this work), the difference is about 20 °C. In all cases the total heat produced in the fuel stripe is assumed constant, so with more layers added the heat produced in each layer is proportionally reduced.

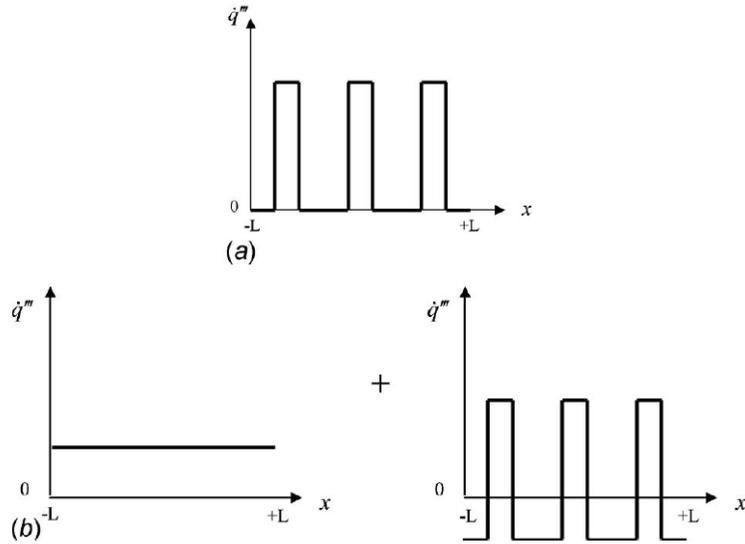


Figure 4.4. Power distribution (a) and decomposed power distribution (b). [33]

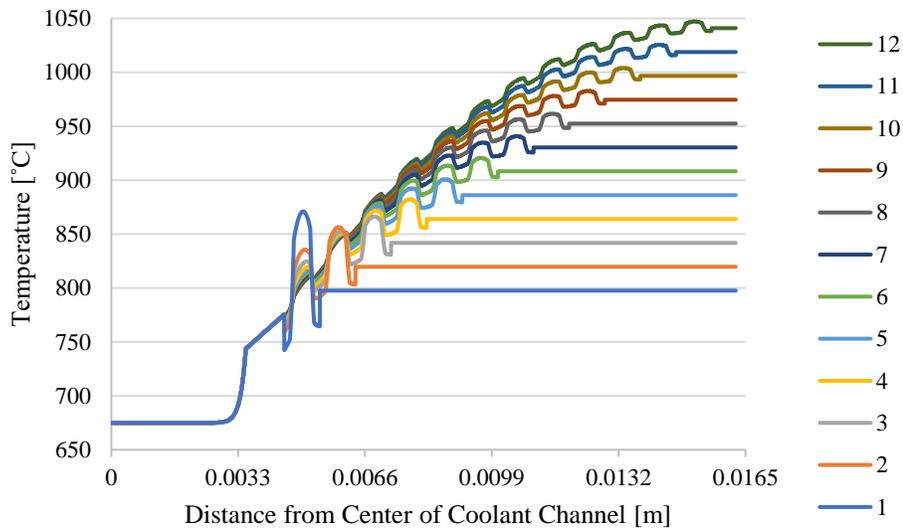


Figure 4.5. Temperature distribution for heterogeneous model through the center of TRISO particles based on the number of TRISO layers. Assumes average core power and coolant temperature.

Tabulated values of the results seen in Figure 4.3 and Figure 4.5 can be seen in Table 4.5. Since the total power generated in all layers is held constant, the temperature profiles in the coolant and sheath (sleeve) graphite regions are constant for all cases. As can be seen in Figure 4.3, the minimum fuel stripe temperature for all homogenized fuel stripe cases is 775.5 °C. Additionally,

both the maximum (22.1 °C/layer) and average (11.1 °C/layer) homogenized fuel stripe temperatures scale linearly with the number of fuel layers used.

Table 4.5. Average and maximum temperatures for both homogenized fuel stripe and explicit fuel kernels. All cases assume core-average specific power and average coolant density and temperature. The cases using four TRISO fuel layers are highlighted as they correspond to the reference design used in this dissertation.

TRISO Layers	Average Homogenized Stripe Temperature [°C]	Maximum Homogenized Stripe Temperature [°C]	Average Fuel Kernel Temperature [°C]	Maximum Fuel Kernel Temperature [°C]
1	786.5	797.6	858.2	871.0
2	797.6	819.7	838.9	856.4
3	808.6	841.8	842.3	866.3
4	819.7	864.0	851.4	882.3
5	830.8	886.1	862.8	900.8
6	841.8	908.2	875.3	920.4
7	852.9	930.3	888.4	940.8
8	864.0	952.5	901.9	961.6
9	875.0	974.6	915.7	982.8
10	886.1	996.7	929.7	1004.1
11	897.2	1018.9	943.9	1025.5
12	908.2	1041.0	958.1	1047.1

4.2.6 Dependence on Axial Location

The amount of power produced in a certain site is dependent upon the axial location in the core. Figure 4.6 shows how the temperature profile changes from the coolant interface to the center of the plank for a radially average assembly for different axial locations in the active core. A height of 0 m corresponds to the bottom of the active core. The height of the active core is 5.5 m, so the peak difference occurs at 2.75 m (if symmetric). A truncated cosine axial profile with a peaking factor of 1.3076 (found from previous studies) and an extrapolation distance of 0.75 m was assumed for the axial power corresponding to Figure 4.6. Note that the multiphysics model used later in the dissertation obtains its power profile directly from fission rate tallies of the previous Serpent simulation. Figure 4.6 shows the temperature profile of the fuel plank from the coolant-plank interface (left) to the farthest part of the fuel stripe (right). Due to symmetry, the temperature at the center of the fuel plank is assumed to be constant, so the temperature distribution is cut at the fuel stripe / center of plank interface to show more detail.

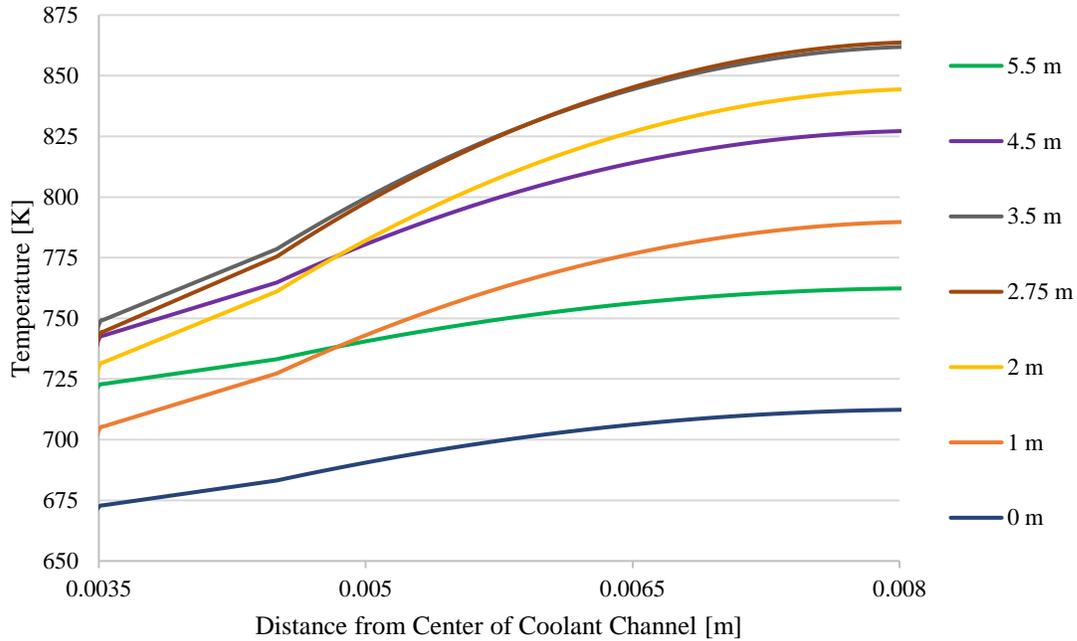


Figure 4.6. Average plank temperature distribution by axial location for four TRISO layers.

For the assumed power profile, one can infer that there is about a 23 °C jump between the bulk coolant temperature and the plank surface at the inlet/outlet and about a 70 °C jump at the center of the core (assumed peak power location). As for temperature changes across the planks themselves, there is an increase of about 40 °C at the inlet/outlet and about 122 °C at the center of the core.

4.3 Neutronic Impacts of Temperature Variation

4.3.1 Use of the Average Fuel Temperature within Fuel Stripes

A study of the fuel temperature coefficient was conducted on a reflected assembly model. This case uses BP spheres with the CB fully withdrawn. As can be seen from Figure 4.7, the behavior is well-described as linear, which has an R^2 -value of 0.9982. Error bars lie within the data points shown in Figure 4.7, which is why they are not visually discernable. The resulting fuel temperature coefficient is -4.26 ± 0.03 pcm/°C. This value is comparable to those typically seen in LWR systems, another thermal spectrum design.

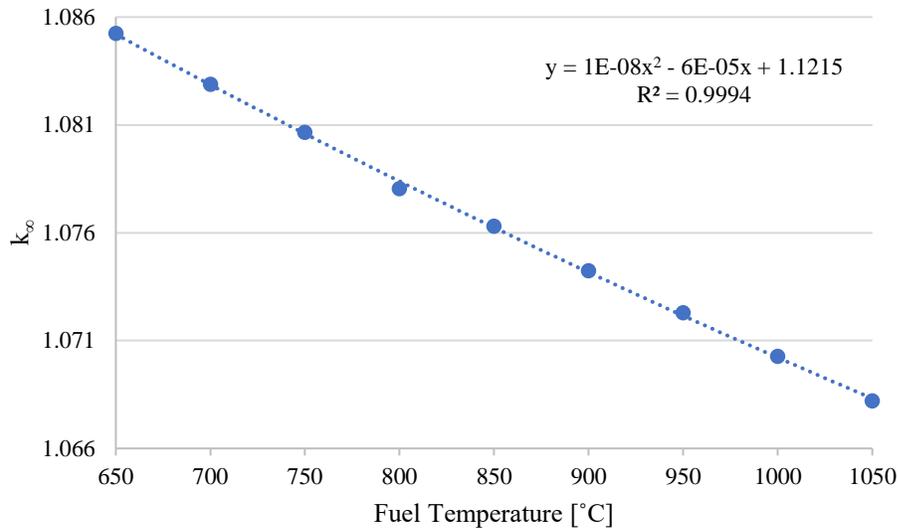


Figure 4.7. Fuel temperature coefficient for a reflected assembly.

The results of Figure 4.7 suggest that the Doppler temperature coefficient for an AHTR assembly to be well-characterized by a linear relationship. Further, it supports the assumption to use the average fuel temperature for the effective Doppler temperature of the fuel since the reactivity impact from using the average fuel temperature should be comparable to the average reactivity impact from using layer-specific temperatures. This assumption requires that the flux depression inside fuel kernels to be small, which is in fact the case. Recall from Table 1.4 that the density of ^{235}U in the fuel kernel is about 2.27×10^{-3} atom/barn/cm. Given that the thermal absorption cross section at 2200 m/s for ^{235}U is about 678 barns [34], this gives a macroscopic absorption cross section of about 1.55 1/cm and a mean free path of approximately 0.65 cm when neglecting other reaction channels. Recall from Table 3.2 that the radius of a fuel kernel is 0.02135 cm, which is only about 3.3% of the mean free path. Thus, one would expect that the thermal flux would only decrease by about the same amount, providing confirmation that the flux depression is small inside the TRISO fuel kernels and further that using temperature averaging should be an adequate estimate for the effective Doppler temperature.

As was observed in Table 4.5, the amount of temperature variation seen in the fuel kernels of a single section of the AHTR reference design (four TRISO particle layers) is much narrower

than the range of temperatures considered in Figure 4.7. This means that the averaging process used falls well within this range and is appropriate for both fuel particles within fuel layers and fuel layers within fuel stripes.

4.3.2 TRISO-Wise Fuel Temperature Impact on Power Distribution

It is important to characterize the magnitude of impact to be expected from temperature changes to the fuel since multiphysics coupling and the resulting variation of neutronic and thermal parameters are integral to this dissertation. To address this, cases using average power density (about 78.14 mW per particle, corresponding to 3400 MW_t for full core system) were compared with cases using twice (double) the average power density (about 156.28 mW per particle). These two power densities should provide a representative basis for expected variation in thermal and neutronic performance due to effects like spatial peaking. For both power densities, it is assumed that the coolant is at average temperature (675 °C) and only the reference case of four TRISO layers is considered. The resulting temperature profiles from the center of the coolant channel to the center of the fuel plate for these power densities can be seen in Figure 4.8.

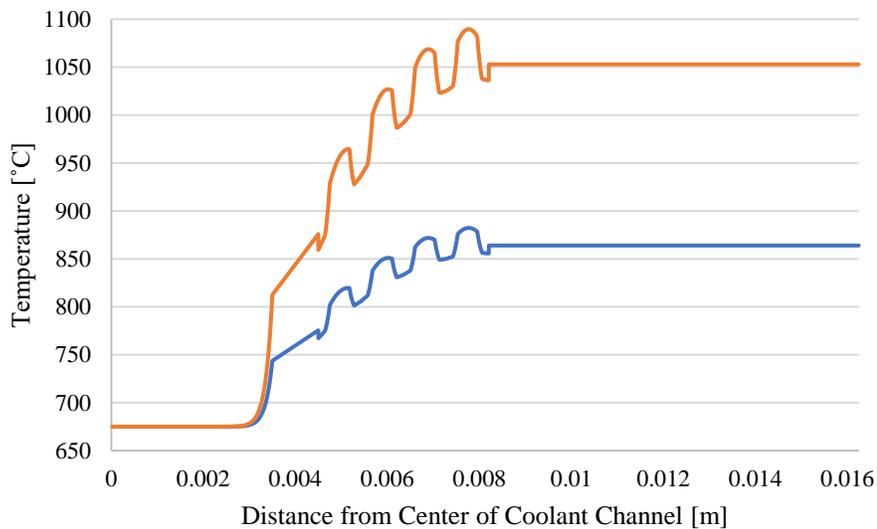


Figure 4.8. Temperature distribution across coolant channel and fuel plate for average power density (blue) and double average power density (orange) of a fuel stripe with four layers.

The temperature variation in the four TRISO layers is evident in Table 4.6. For each individual layer, the volume-weighted average temperature for the fuel kernel was calculated. Layer 1 is closest to the coolant channel and Layer 4 is closest to the center of the fuel plank. Additionally, the average temperature of all four layers is also found and is shown in the row called “Average”.

Table 4.6. Layer-wise and average fuel kernel temperatures for average and double power density.

TRISO Layer	Fuel Kernel Temperature [K]	
	Average Power	Double Power
1	1086	1223
2	1119	1290
3	1141	1334
4	1152	1356
Average	1125	1301

Based on the fuel kernel temperatures shown in Table 4.6, four cases were run in Serpent to track TRISO-wise power distributions for a 2D assembly model. The cases use either average or explicit TRISO layer temperatures as well as either average or double power density. The two most relevant comparisons are discussed below but all sets showed similar results. The Serpent models used in this dissertation will most closely resemble the case with average TRISO layer temperature and average power density. For this reason, this case is compared to explicit TRISO layer temperatures and average power density as well as average TRISO layer temperature and double power density. The comparisons respectively address the questions: 1. “what is the impact to the fission distribution from using explicit temperatures for each TRISO layer versus the average?”, and 2. “what is the impact to the fission distribution from varying the fuel kernel temperature?”. Ideally, to show that the differences are not significant, it needs to be demonstrated that the differences are negligibly small and lie within the statistical uncertainties of the considered cases. The results supporting this are shown in Figure 4.9 and Figure 4.10. These figures show TRISO-wise power comparisons in each of the twelve fuel stripes per assembly section. Fuel stripe dimensions are 202 particles in length and four in width, totaling 9696 particles shown. All cases

were run using a large number of particle histories in an attempt to keep TRISO-wise uncertainties small: 2.5×10^6 particles per generation, with 5000 active cycles and 100 inactive cycles. Simulations each took about 151 h wallclock on 24 processors.

Figure 4.9 shows the results for comparing explicit versus average TRISO layer temperature. As is evident, there does not appear to be an observable trend in the differences between the two cases. 95.99% of individual TRISO values fall within $\pm 2\sigma$, which agrees closely with the classically expected value of 95.45% for a normal distribution of uncertainties.

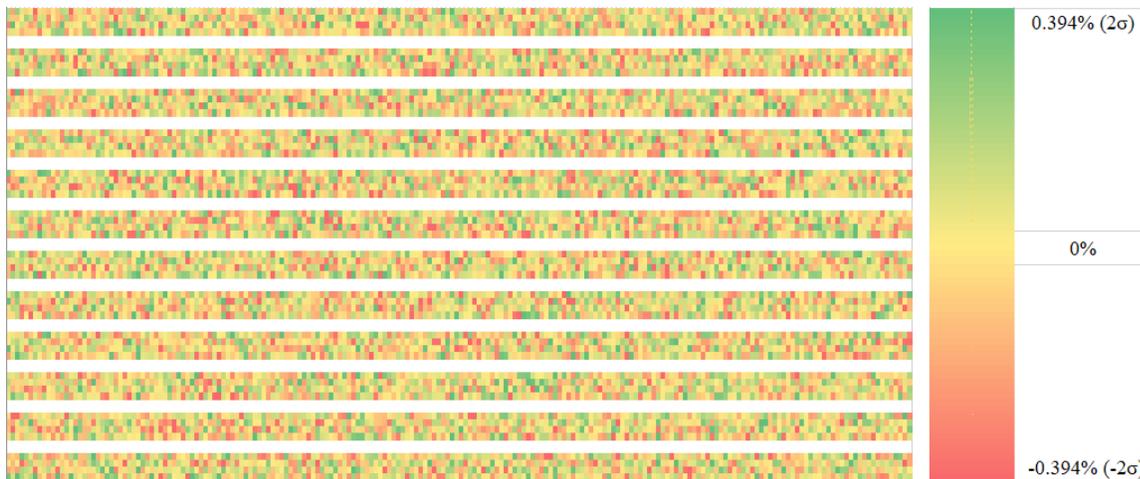


Figure 4.9. Relative differences from Explicit Layer Temperature Average Power case to Average Layer Temperature Average Power case. Maximum: 0.741%. Minimum: -0.673%.

Figure 4.10 shows the results for comparing average layer temperatures due to average and double power density. Similar to the results seen in Figure 4.9, Figure 4.10 does not appear to have an observable trend in the differences between the two cases. 95.64% of cell values fall within $\pm 2\sigma$, which agrees well with the classically expected value of 95.45% for a normal distribution of uncertainties.

Both Figure 4.9 and Figure 4.10 support the conclusion that there is no significant impact on the fission rate distribution in fuel stripes due to either explicit temperature treatment of the fuel kernels in TRISO layers or expected differences in fuel kernel temperature (due to effects such as spatial peaking) for average TRISO layer temperature treatment. This means that the average fuel

kernel temperature for an entire fuel stripe obtains comparable fission rate results to using explicit temperatures for each TRISO layer, so the assumption of using the average is shown to be as good as the layer-wise treatment.

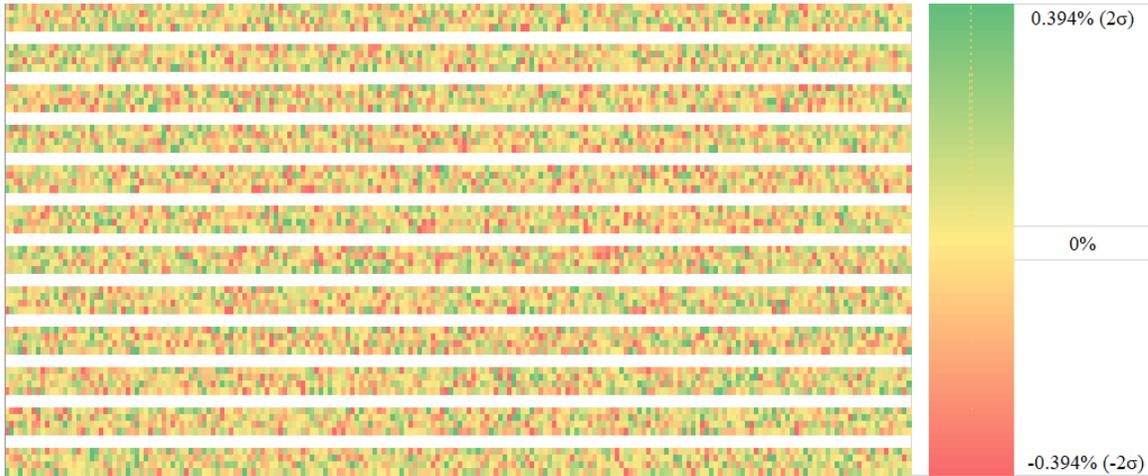


Figure 4.10. Relative differences from Average Layer Temperature Double Power case to Average Layer Temperature Average Power case. Maximum: 0.679%. Minimum: -0.694%.

4.4 Temperature-Dependence of Fuel Thermal Conductivity

The previous results and figures in this section were generated using an assumed average fuel thermal conductivity of 3.7 W/(m·K). However, the thermal conductivity of the fuel used in AHTR is not constant. Driven mainly by differences in local power in the core, the fuel kernels will be at different temperatures in the core. A feature was developed to capture this dependence of fuel thermal conductivity on fuel temperature. As can be seen in Figure 4.11 [3], fuel thermal conductivity decreases as fuel temperature increases. This means that fuel sections producing more power will not only see larger gradients due to linear scaling from using the average power solution but will also have even hotter fuel kernels due to increased thermal resistance from a lower thermal conductivity.

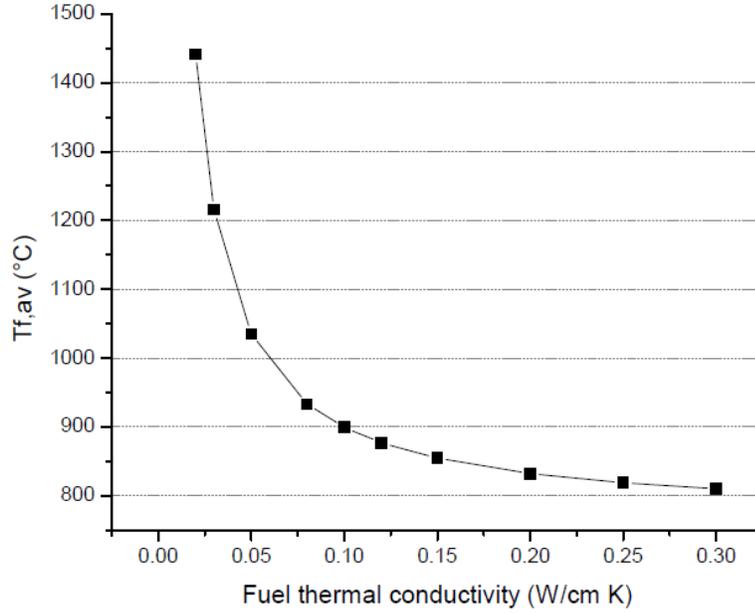


Figure 4.11. Dependence of fuel thermal conductivity on fuel temperature. [3]

In Figure 4.11, thermal conductivity is only provided for temperatures ranging from 805 to 1450 °C. Since the behavior outside of this range is not characterized, the thermal conductivity will be assumed to be constant both above and below the range by using the boundary value. The best fit relationship over this range and the assumed values beyond it is given by Equation 4.11.

$$k_f(T) = \begin{cases} 31 & T < 805 \\ 805.855 \cdot (T - 768)^{-0.902311} & 805 < T < 1450 \\ 2.25 & T > 1450 \end{cases} \quad (4.11)$$

k_f is the thermal conductivity of the fuel in W/(m·K) and T is the temperature of the fuel in °C. The power function described over the applicable temperature range of Equation 4.11 has an R^2 value of 0.9949, which is an adequate fit for the datapoints visually obtained from Figure 4.11.

4.5 Impact of Neutron Fluence on Graphite Thermophysical Properties

Thermophysical properties of graphite tend to be both temperature and irradiation dependent and are different for each grade of graphite [35]. Some temperature-dependent behaviors of other materials were addressed in the previous section, and fluence effects will be addressed

here. The user will need to provide the unirradiated value for the specific material (graphite) assumed in their analysis. For illustration, results from a study considering two grades of graphite which underwent irradiation will be analyzed below.

This section presents results from another research group [36] which considered the effects of fast neutron fluence on graphite thermophysical properties. Two grades of graphite were considered as part of this work: G347A and G458A, both from Tokai Carbon Co., Ltd. The pre-irradiated density of both of these grades is about 1.85 g/cm^3 , which is higher than that assumed for graphite used in AHTR (1.75 g/cm^3) but lower than that assumed for carbonaceous structures used in AHTR (1.95 g/cm^3). The results of this prior work are likely extendable to both AHTR materials due to similarities.

Studies were conducted by irradiating samples in a high-flux facility and analyzing sample properties once withdrawn. Results are typically calculated in both the transverse (TR) and axial (AX) directions, though the two are quite similar in the select figures presented here.

To get a feel for the magnitude of the fluence values used for this graphite study, previous work with the AHTR [32] estimates that the fast flux ($E > 0.1 \text{ MeV}$) in the average assembly section is about $6 \times 10^{13} \text{ n/cm}^2\text{s}$. Over a six-month period (length of a single fuel cycle), the average fast fluence within the assembly should be about $9.46 \times 10^{20} \text{ n/cm}^2$. Roughly, this is approximately $1 \times 10^{21} \text{ n/cm}^2$ for each fuel cycle or $3 \times 10^{21} \text{ n/cm}^2$ total fluence at core discharge for a three-batch refueling schedule.

4.5.1 Volumetric Change

One property of graphite is that in addition to normal thermal expansion, it experiences a roughly parabolic swelling effect with respect to fast neutron fluence. Initially, for low fluence levels, the graphite contracts. At higher fluence levels, the behavior changes and the graphite expands; ultimately beyond that of its starting volume. Figure 4.12 shows how samples of graphite G347A fared under irradiation at different temperatures. One will notice that the purple dataset

corresponding to 684 °C lies within the thermal operating range of AHTR. This dataset has a maximum negative volume change (shrinkage) of about 4% at a fast fluence of about 10.5×10^{21} n/cm².

This work is assumes that graphite has a thermal expansion coefficient of 5×10^{-6} . For a temperature increase of 700 °C (approximate change from room to operating temperature for graphite), this only corresponds to a volumetric change of about 1%. One immediately sees that graphite swelling due to fluence can be a stronger behavior at higher fluence levels and will need to be tracked. Ultimately, graphite swelling can lead to a limiting condition inside the core since the rate of volume increase is quite large once the starting volume is passed. This is only applicable for permanent graphite structures in the core (such as the permanent radial reflector) since removable features (like fuel assemblies) of the core will be replaced well before this amount of fluence is reached. The purple (684 °C) dataset has a fit given by Equation 4.12.

$$\frac{\Delta V}{V_0}(F) = \frac{4}{10.5^2}(F - 10.5)^2 - 4 = 0.0362812F^2 - 0.761905F \quad (4.12)$$

F is the neutron fluence of the graphite (in 1×10^{21} n/cm²) and the fraction change in volume $\Delta V/V_0$ is given as a percent. For sample calculations, assuming the average core fluence, this would cause about -0.76% relative change after one fuel cycle and about -1.96% relative change at core discharge for a fuel assembly.

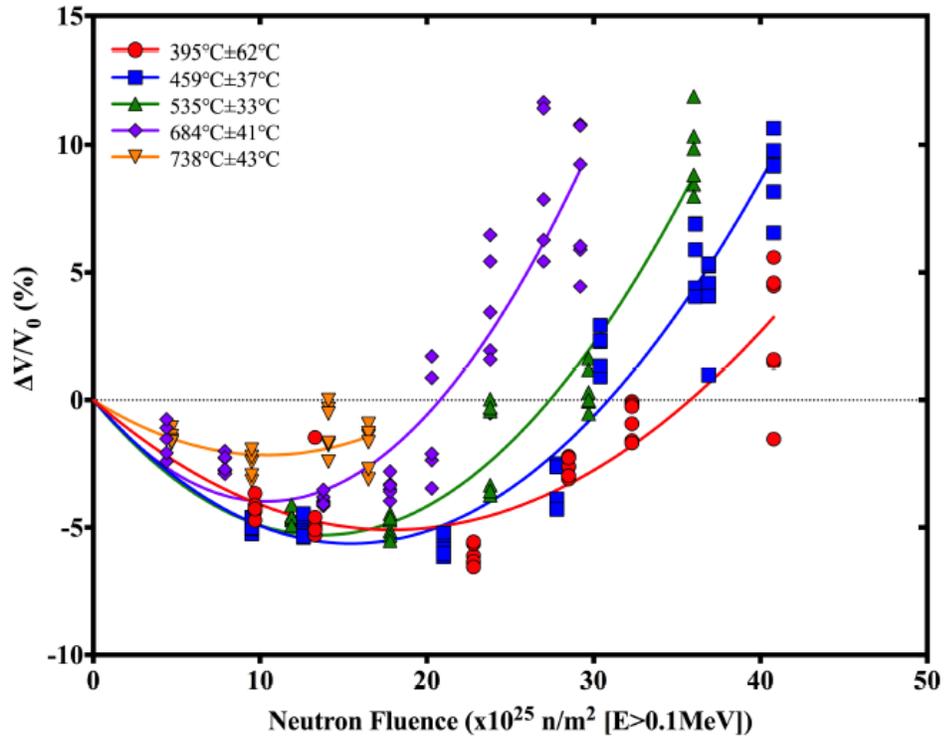


Figure 4.12. Volume change versus neutron fluence for specimens irradiated at different temperatures. Each data point is the volume change of an individual specimen. The lines were fitted to the data using a second-order polynomial [36].

4.5.2 Thermal Expansion Coefficient

Fluence-dependent change to the thermal expansion coefficient can be seen in Figure 4.13. Once again, the purple (684 °C) plot will be used as a reference for the behavior of materials used in AHTR. On top of the temperature-dependence of the thermal expansion coefficient of graphite, there appears to be significant change to the thermal expansion coefficient due to fast fluence. For low irradiation levels, the thermal expansion coefficient increases. At higher irradiation levels, it decreases. For very high irradiation levels (fast fluence greater than 20×10^{21} n/cm²), the normalized change to the thermal expansion coefficient plateaus for all irradiation temperatures. This last behavior was thought to be attributable to radiation-induced defects being more numerous/larger for higher temperature samples, which result in the same mean thermal expansion coefficient over the considered temperature range. The fluence of core components is expected to stay below 10×10^{21} n/cm², so this behavior should not be relevant.

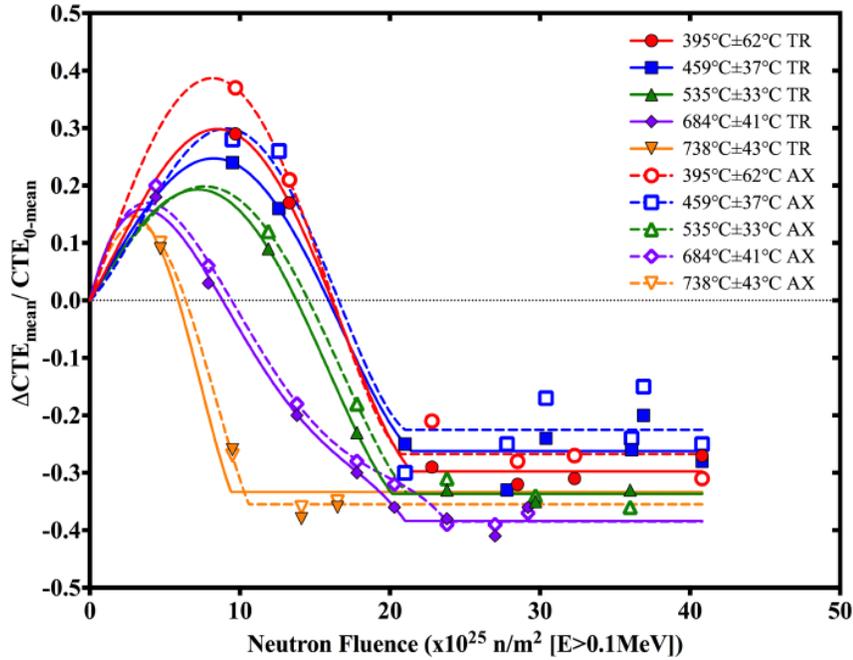


Figure 4.13. Normalized change of mean thermal expansion coefficient (referenced to 25 °C) measured at irradiation temperatures, plotted versus neutron fluence [36].

The expected average core discharge fluence is about $2\text{-}3 \times 10^{21}$ n/cm², depending on the number of fuel cycles residing in the core. Even at the center of the fuel assembly, the fluence should not exceed 10×10^{21} n/cm² (more than three times the average value). Therefore, a functional fit will only be made over this range. The behavior is approximately parabolic with zeroes at 0 and 10×10^{21} n/cm² and a maximum value of about 16%. This produces Equation 4.13.

$$\frac{\Delta CTE_{mean}}{CTE_{0-mean}}(F) = -\frac{16}{5^2}(F - 5)^2 + 16 = -0.64F^2 + 6.4F \quad 0 < F < 10 \quad (4.13)$$

F is the neutron fluence (in 1×10^{25} n/m²). For sample calculations, assuming the average core fluence, this would cause about 5.76% relative change after one fuel cycle and about 13.44% relative change at core discharge (three cycles) for a fuel assembly.

4.5.3 Thermal Conductivity

Another relevant graphite parameter to consider is thermal conductivity. This will greatly impact the temperatures of components in the core, the most important likely being the fuel Doppler coefficient. As can be seen in Figure 4.14, thermal conductivity changes significantly for the fluence amounts presented. The exception to this is the set of values shown for a neutron fluence of 9.5×10^{21} n/cm² which is omitted from further consideration and believed to be inconsistent due to the similarity to the results for 21×10^{21} n/cm².

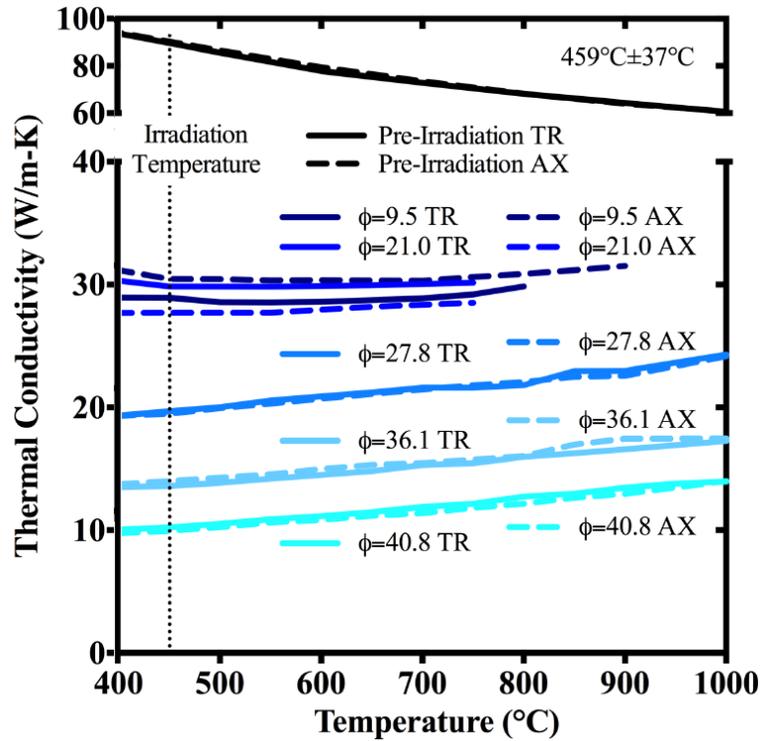


Figure 4.14. Measured thermal conductivity versus measurement temperature for as-received and specimens irradiated at 459 ± 37 °C. The vertical dashed line indicates the irradiation temperature. The labels indicate the specimen orientation and total neutron fluence ($\times 10^{25}$ n/m² [$E > 0.1$ MeV]) [36].

A two-parameter (fluence F in 1×10^{25} n/m² and temperature T in °C) functional fit of the thermal conductivity (k) results shown in Figure 4.14 can be expressed as Equation 4.14.

$$\frac{k(F,T)}{k_{0,459}} = e^{-0.053364F} \left(1 - \frac{T-459}{1938} (3 - 2e^{0.019364F}) \right) \quad \begin{array}{l} 0 < F < 40.8 \\ 459 < T < 1000 \end{array} \quad (4.14)$$

In Equation 4.14, $k_{0,459}$ refers to the thermal conductivity of the graphite with zero fluence at 459 °C. It is useful to place the behavior in terms of a single thermal conductivity reference because while the thermal conductivity may vary between graphite grades, it will be assumed that the functional trends are comparable to these functional changes in both temperature and fluence. A derivation of Equation 4.14 can be found in APPENDIX A (page 240). A heat map of Equation 4.14 can be seen in Figure 4.15.

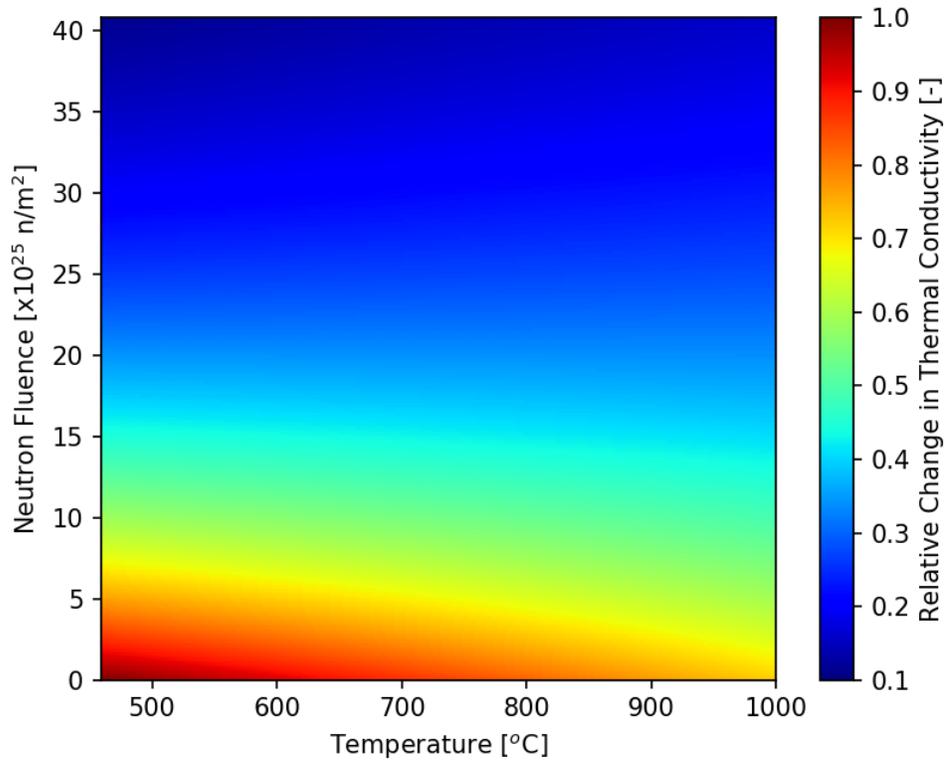


Figure 4.15. Heatmap of the relative change in thermal conductivity of grade G347A graphite when varying both temperature and neutron fluence ($E > 0.1$ MeV). Change is relative to value for nonirradiated graphite at 459 °C.

In this chapter, a temperature profile was found for the coolant channel and fuel plate using the average volumetric heat rate. The profile was obtained by homogenizing the fuel stripe, conducting a 1D heat transfer analysis on the result homogenized system, and then superimposing a power modulation function over the homogenized solution to obtain a profile accounting for

explicit TRISO particles. Some confirming analyses were conducted to show that the fuel Doppler coefficient is linear (allowing for simple temperature averaging), and to show that significant heat rate changes did not meaningfully change the power distribution (allowing for that same power distribution to be assumed for any temperature profile). The effects of fast neutron fluence on graphite were also considered and shown to be quite significant for some thermophysical properties. For those, representative functional fits were developed to be used in multiphysics simulations.

CHAPTER 5

TRIANGULAR MESH TALLYING IN SERPENT

Any analysis using 3D MC ultimately aims to have an accurate distribution of results with an adequate level of resolution. In this case, an important component of coupled multiphysics modeling is having some means of determining how much heat (in this case, fission power) is locally produced during operation. While deterministic methods automatically obtain these results as an inherent part of the solution process, MC codes require additional tallying efforts, which can be computationally costly.

The 3D AHTR Serpent model requires resolution of hexagonal assembly-wise power profiles at the one-third assembly level. Tallying fission power results with fuel-specific tallies is a possibility, but in application is often time prohibitive. As will be shown, traditional MC tallies which use combinatorial-based binning (i.e. physical properties of the model such as cell, material, universe, etc.) can be very slow and significantly hinder the neutron transport portion of a MC simulation. Typically, MC codes also have superimposed mesh tallies, which perform faster based on how they search the phase space. For this reason, it is desirable to have all fission power tallying done using a superimposed geometric mesh, which has a relatively small computation overhead with respect to transport runtime.

Currently in Serpent, various geometric mesh tallies exist but do not offer a straightforward way to obtain one-third assembly power distributions in a core with hexagonal assemblies. A Cartesian mesh is the most commonly used and could work fine in the axial direction, but the triangular pitch of assemblies makes binning in radial directions challenging. Curvilinear meshing is infeasible and using a hexagonal mesh can only resolve the power distribution at a whole assembly level. An option, which could work but would be needlessly cumbersome would be to use an unstructured mesh in Serpent [37]. This process requires Serpent to read-in multiple data files to function properly: a points file to establish boundary locations, a faces file to establish

boundary surfaces, and a neighbor file to establish what other cells are next to each created cell. For an arbitrary geometry this would be necessary, but for a regular lattice geometry this would be similarly computationally prohibitive to using a physical tally.

Upon inspection of the AHTR geometry, one can see that each one-third assembly region is a parallelogram, but further that it can be thought of as the union of two equilateral triangle-shaped sections. The sum of the power from the two corresponding triangular regions would give the power for the one-third assembly section. This work creates a new mesh type for Serpent: a triangular mesh. This mesh type makes equilateral repeating triangles in 2D (x- and y-directions) and equilateral triangular prisms in 3D (normal Cartesian partitioning in the axial z-direction). It performs similarly to the hexagonal mesh, which also needs to establish a triangular pitch structure but further partitions into six triangles instead of a whole hexagon. The capability is based upon dividing the existing hexagonal tally regions into sixths. This is achieved by expanding the current code from binning at the whole hexagon level and further dividing the hexagon into sixths. Both this application as well as other general cases including reactor designs utilizing hexagonal fuel elements would benefit from having this capability within Serpent. The two major objectives of this proposed feature are to create a superimposed mesh-based tally which is able to obtain one-third assembly fission power profiles for the 3D AHTR Serpent model and at the same time not meaningfully impact the transport simulation runtime.

5.1 Current Tallying Efficiencies in Serpent

MC codes like Serpent have multiple options when it comes to tallying results. The runtime of the transport simulation can be impacted, sometimes significantly, by the type of tally used. The reason for this is the amount of logical and numerical operations needed to check the phase space associated with the tally. When the number of bins in a tally grows large, the effort required to perform these checks can make the simulation cumbersome. For region-specific physical tallies (cell, material, universe, etc.), bins within the tally are systematically searched to determine if and

where the score should be recorded. For several thousand bins, this results in many checks with a vast majority returning no score. This extensive checking process can be avoided by using a geometric mesh tally instead. The mesh is superimposed over the model geometry and uses uniform spacing in each direction. The benefit is that when an event is scored, the uniform regular mesh allows for a much faster recording process where the single applicable bin can be quickly found. This ability to efficiently record the event and proceed to the next in the transport simulation translates to a much smaller computational penalty over a region-specific physical tally and for most simulations will only trivially increase the runtime.

Another MC transport code commonly used in the field is MCNP (Monte Carlo N-Particle) [38]. In work done by van Veen and Hoogenboom [39] [40] using MCNP, they detail computational challenges involving tallying in many regions. They reported that the increase in runtime has an approximately linear relationship with the number of tallying regions, moreover, that it quickly exceeded the transport simulation time. To test this claim, a similar study was conducted on the Serpent 3D AHTR model using the material tally for a different number of tallying bins. As can be seen in Figure 5.1, the results verify what is also observed with MCNP. Further, it highlights how computationally taxing the tallying process can be when using a non-spatial tally.

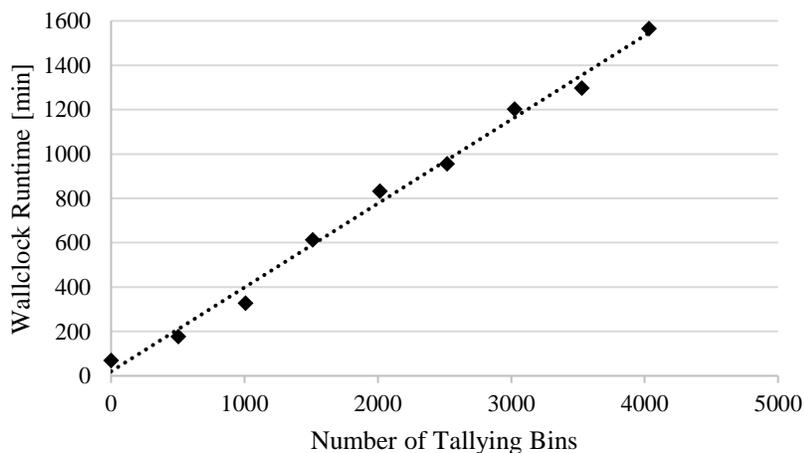


Figure 5.1. Dependence of transport runtime on the number of material tally bins in Serpent.

In Figure 5.1, the set of material tally simulations used 4032 unique fuel materials (252 fuel assemblies, 16 axial segments) in the model geometry. Figure 5.1 used a material tally, which is a type of physical tally since it corresponds to a physical feature in the combinatorial-based geometry. From an input and output perspective, these tallies are the most straightforward. The user provides a unique identifier that is already used in the model geometry, and additionally the output uses the same designation. From a simulation efficiency perspective, physical tallies (when a large quantity of bins are present) use a considerable amount of time searching for the correct bin to score each event.

As mentioned before, geometric mesh tallies eliminate this search process by quickly mapping an event site to the corresponding tally bin. The only search that needs to be performed is a relatively simple one establishing where in the uniform mesh the event took place. Table 5.1 shows a comparison of runtimes for different fission power tallies using the 3D Serpent AHTR model. Note that radially, the hexagonal mesh is tallied over each individual whole assembly while the three region-specific tallies score results in each one-third core symmetric, one-third assembly location. While the two sets of results use different meshes, the two meshes have a similar number of bins and should be comparable for runtime testing purposes.

Table 5.1. Tallying efficiency of fission power in Serpent for the 3D AHTR core.

Tally Type	Number of Bins	Transport Time [min]	Relative Slowdown Factor
None	0	70	1 (reference)
Hexagonal Mesh	5776	70	1.0
Cell	4032	1876	26.8
Material	4032	1566	22.4
Universe	4032	1785	25.5

From Table 5.1, one can see that for a model of this size, using a superimposed geometric hexagonal mesh has virtually no impact on the speed of the transport calculation relative to using no power tallying whatsoever. However, the region-specific physical tallies experience a significant slowdown: cell, material, and universe binning slowdown the transport calculation by more than a

factor of 20. This level of computational penalty is not acceptable for a model of this size and efforts should be made to use a geometric mesh if possible.

5.2 Introduction of Triangular Mesh Tally

As explored in the previous section, using a superimposed geometric mesh is the preferred method for fission power tallying in a model the size of the 3D Serpent AHTR core because using physical tallies is likely to make the simulation prohibitive due to the increase in runtime. Recall that the aim in this work is to resolve power at a one-third assembly level. As it stands now, there is not a geometric mesh in Serpent that would allow for such tallying. Therefore, modifications have been made to the existing capabilities in Serpent to make a new mesh type, which has spatial partitioning within a hexagon to capture the behavior at the section-level. This work splits each hexagonal mesh region into sixths to divide it into equilateral triangular bins. Each one-third AHTR assembly section is then simply the sum of the two corresponding one-sixth hexagon tallies.

Creating a new geometric mesh type based upon the existing hexagonal mesh with subdivisions within the hexagon serves the needs of the AHTR geometry as well as maintains a level of generality applicable to other reactor designs which frequently use hexagonal assemblies such as liquid metal-cooled fast reactors and high temperature gas-cooled reactors. This generality makes this type of tally useful for other researchers beyond the development for analyzing AHTR.

Currently in Serpent, there are two hexagonal orientations available for creating both model surfaces and superimposed geometric meshes. This includes one with one side perpendicular to the x-axis (type "hexx") and another with one side perpendicular to the y-axis (type "hexy"). These can be visualized in Figure 5.2. Having both options available greatly helps with modeling flexibility and most reactors, which use hexagonal assembly design features can leverage one of the two orientations.

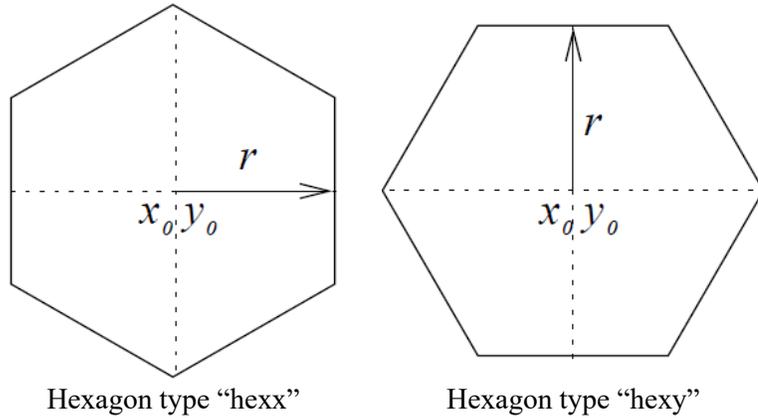


Figure 5.2. Orientations of the types of hexagons used in SERPENT.

The implemented triangular mesh has the capability to use either of these orientations as a basis for the new tally. Divisions are made within the hexagon to allow for separate bin partitions. A visualization of this partitioning can be seen in Figure 5.3.

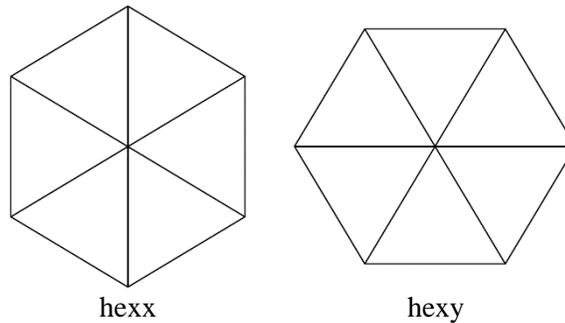


Figure 5.3. Hexagons divided into six equilateral triangles.

5.3 Implementation of Triangular Mesh Tally

5.3.1 Changes Made to the Serpent Source Code

Within Serpent, the superimposed hexagonal geometric mesh tally uses the designation “dh” within the tally (referred to as a “detector” in Serpent) declaration of an input file. An example of this declaration is given in Equation 5.1.

$$\text{det } \langle \text{name} \rangle \text{ dh } \langle \text{type} \rangle \langle \text{x0} \rangle \langle \text{y0} \rangle \langle \text{pitch} \rangle \langle \text{Nx} \rangle \langle \text{Ny} \rangle \langle \text{zmin} \rangle \langle \text{zmax} \rangle \langle \text{Nz} \rangle \quad (5.1)$$

Parameters in angle brackets (< >) are user-specified inputs. The <type> can be hexx (<type>=2) or hexy (<type>=3). The center of the tallying region is given by <x0> and <y0>, with pitch between hexagons as <pitch>. The size of the mesh is prescribed by the number of hexagons in both the x- (<Nx>) and y- (<Ny>) directions. Axial binning is Cartesian-like with lower and upper bounds given by <zmin> and <zmax> with uniform divisions into <Nz> bins.

The “dh” tally was modified to allow for triangular region binning and given the keyword “dht”. The syntax is otherwise the same compared to the “dh” tally but is given in Equation 5.2 for completeness.

$$\text{det } \langle \text{name} \rangle \text{ dht } \langle \text{type} \rangle \langle \text{x0} \rangle \langle \text{y0} \rangle \langle \text{pitch} \rangle \langle \text{Nx} \rangle \langle \text{Ny} \rangle \langle \text{zmin} \rangle \langle \text{zmax} \rangle \langle \text{Nz} \rangle \quad (5.2)$$

The reason why no new parameters are needed is that it is understood by using the “dht” tally that six partitions will be used when tallying. This does not need to be conveyed to Serpent as the partitioning is hard-coded. Additionally, all the other parameters are the same as in a superimposed hexagonal mesh since the two tallies share the same geometric basis other than the divisions in the “dht” tally.

To create the “dht” tally, eighteen subroutine or header files in the SERPENT 2.1.31 source code required modification. The changes are summarized in Table 5.2 with additional discussion below. Complete source code modifications can be found in Appendix B. Changes to Serpent source code files can be summarized in six broad areas, specified by the first column of Table 5.2. The first area introduces new reserved words and named parameters in the code related to the new tally. The second area initializes the tally when the code is reading an input file. Small changes were also made to preexisting mesh types where the triangular-related treatments would need to be trivially set to one (no additional indexing) so that the other tallying options still work with the modified code. The third area allocates the memory for where scores will be recorded. Certain tallying functions needed to be broadened to account for the increase in the size of the phase space (adding the triangular indexing). The fourth area deals with scoring within the tally during a transport simulation. This is where the most significant changes and additions were

made since this is where the logic and math for scoring in specific sections is contained. The fifth area is for formatting and printing the results to an output file once the transport simulation is complete. The sixth and final area has no significant changes to the listed files, but they contain function calls, which have arguments expanded to include the triangular phase space.

Table 5.2. List of Serpent source code files modified to create the triangular mesh tally.

Function of Modified File	Filename	Description of Modifications
Recognition of input and source code keywords	header.h	Initialize new mesh type, modify affected functions
	locations.h	Declare new triangle-related variables
Initialization of new tally	readinput.c	Tell code how to construct “dht” tally
Create tallying structure for recording scores	createmesh.c	Allocate memory of triangular bins
	processdetectors.c	Make additional partition for triangular regions
Find where a score should be recorded	collectdet.c	Cumulative scoring due to triangular addition
	detbin.c	Addition of triangular index
	detidx.c	Addition of triangular indexing
	getlatticeindexes.c	Do math/logic check to find scoring index
	meshindex.c	Obtain the index to score
Write out results	detectoroutput.c	Allows for triangular bin results to be printed
No direct impact, just needed to adjust functional calls to match the number of arguments in modified functions	boundaryconditions.c	Only needed to modify the structure of a function in the file (no major change otherwise)
	dfpos.c	
	findlatticeregion.c	
	icmidx.c	
	nearestmeshboundary.c	
	scoreufs.c	
ufsfactor.c		

5.3.2 Division Scheme for Triangular Bins

The bin numbering structure used within the “dht” tally can be seen in Figure 5.4. While any arbitrary convention could have been selected, the rationale behind the choice is addressed here. First, bins should proceed in a rotational order. Second, the positive corresponding axis should form the boundary between bin 1 and bin 6. This is the y-axis for hexx-type geometry and the x-axis for hexy-type geometry. This provides the numbering scheme seen below. Note that for type hexy, numbering starts at an angle of 0° and proceeds in the positive rotational direction. This is likely the most intuitive scheme to select and was the basis for the convention. For type hexx, numbering starts at an angle of 90° and proceeds in the negative rotational direction. While the rotational directions change between them, the two schemes are rotated mirrors of each other along

the symmetry line $y = x$. This symmetry is leveraged by both types in the source code for the bin-searching functions and other features.

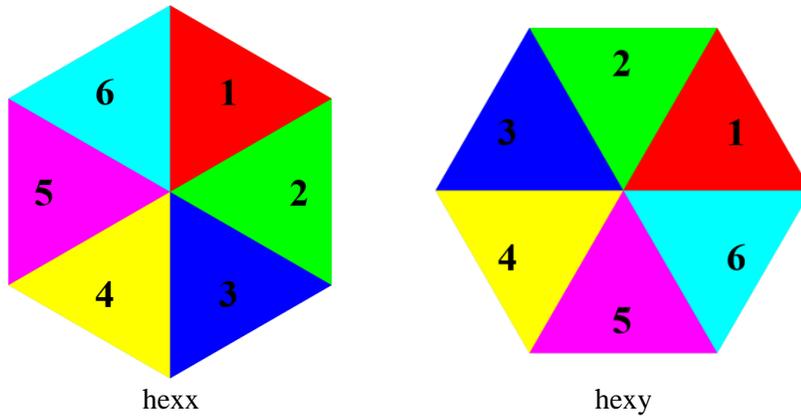


Figure 5.4. Partitioning scheme used to divide “dht” tally bins.

5.3.3 Reporting of the Triangular Mesh Bin Location

Beyond reporting the scores for a given tally, an additional component of the output is to convey where the corresponding tally is located. For the “dh” hexagonal tally this work is based on, the center of each hexagonal region is output with the scores to verify that results match the physical location in the model geometry. For an equilateral triangle, the centroid is located at the intersection of altitudes, i.e., two-thirds of the length of an altitude away from a vertex (for example in Figure 5.5, segment AO with respect to segment AD) [41].

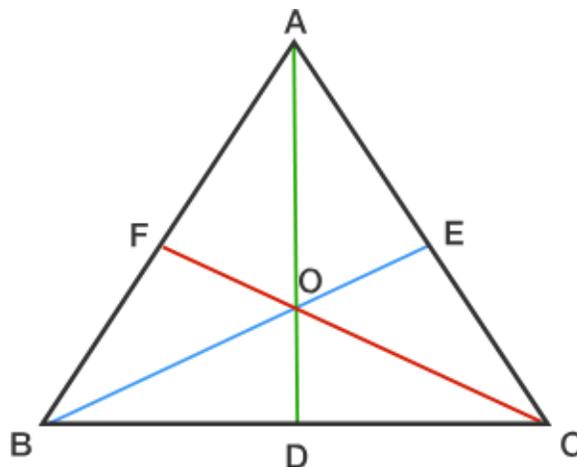


Figure 5.5. Geometry and centroid (point O) of an equilateral triangle [41].

In a regular hexagon, the altitude of one of the six equilateral triangles contained within is the hexagon's apothem, which is half of the pitch in the corresponding hexagonal lattice. Therefore, the centroid of each triangular region in the "dht" tally is at a distance one-third of the pitch from the center of the hexagon. In the Serpent source file "detectorout.c", the corresponding absolute geometric location of each triangular tally region centroid is also written to the output file with the tally results. This allows for ease of verifying where scores correspond to in the model.

5.4 Results and Testing from Using Triangular Mesh Tally

This section will address how the results of the "dht" tally are printed in a Serpent output and present test cases to demonstrate that the tally performs as intended. As for the test cases, six were considered to test common implementations of what may be encountered when using the "dht" tally. These include both hexx- and hexy-type geometry cases for both odd and even lattice sizes as well as cases to demonstrate that axial binning works in conjunction with the tally. Keep in mind that not all tally combinations were tested with other binning options. As was the case before implementing the "dht" tally, it is possible to define tallies which will never return a score (non-overlapping meshes, physical tallies of materials absent in the superimposed mesh, etc.) and care should be made when defining any mesh to avoid this. Since based upon the "dh" tally, the "dht" tally has similar compatibility criteria and should perform as expected with other tallying options when making a compound detector tally definition in Serpent.

5.4.1 Interpreting the Output of the Triangular Mesh Tally

The results from the "dht" tally can be found in the corresponding Serpent detector output file. This file will be named <InputFileName>_det<BurnupStepNumber>, with <BurnupStepNumber> being zero for any statepoint (no depletion) simulation. An example of the output from using the "dht" tally can be found in Figure 5.6. Note that the dimensions do not correspond to those of an AHTR assembly (different size and homogenous composition used for testing only).

```

DEExample = [
  1  1  1  1  1  1  1  1  1  1  1  4.04896E-02 0.00033
  2  1  1  1  1  1  1  1  1  1  1  3.94996E-02 0.00029
  3  1  1  1  1  1  1  1  1  1  1  4.04682E-02 0.00033
  4  1  1  1  1  1  1  1  1  1  1  4.05113E-02 0.00038
  5  1  1  1  1  1  1  1  1  1  1  3.95896E-02 0.00033
  6  1  1  1  1  1  1  1  1  1  1  4.05674E-02 0.00034
];

DEExampleCOORD = [
1.560000E+01 2.701999E+01
3.120000E+01 0.000000E+00
1.560000E+01 -2.701999E+01
-1.560000E+01 -2.701999E+01
-3.120000E+01 0.000000E+00
-1.560000E+01 2.701999E+01
];

DEExampleZ = [
0.00000E+00 1.00000E+02 5.00000E+01
];

```

Figure 5.6. “dht” tally results from SERPENT detector output file.

The first of the three output blocks in Figure 5.6 shows the results for each of the tally bins. From left to right, the first eleven columns present the bin identifier and binning phase space indices. The last two columns contain the tally value and relative statistical uncertainty. A comprehensive list of the significance of each column value can be found below.

1. Value index (unique bin number within the specific detector tally)
2. Energy bin index
3. Universe bin index
4. Cell bin index
5. Material bin index
6. Lattice bin index
7. Reaction bin index
8. Z-mesh bin index
9. Y-mesh bin index
10. X-mesh bin index
11. Triangular mesh bin index (new and only used for the “dht” tally)

12. Mean value (tally bin result)

13. Relative statistical uncertainty

The eleventh column for the triangular mesh bin index is only present in the output file for tallies using the “dht” feature in the tally declaration. This column was added as part of this work and is not printed for other tally types. For all other outputs, the code produces the traditional twelve-column output with the first ten used for the same indices followed by the bin result and its statistical uncertainty.

The second of the three output blocks in Figure 5.6 shows the x- and y-coordinates for the centroid of each mesh bin. As discussed previously in Reporting of the Triangular Mesh Bin (section 5.3.3), the centroids are located one-third of the hexagon pitch away from the center of the corresponding hexagon for the bin. For the hexx-type tally used in Figure 5.6, this results in coordinates at a distance 31.2 cm away from the origin at rotational locations 60°, 0°, 300°, 240°, 180°, and 120° (following the partitioning scheme shown in Figure 5.4).

The third and final output block of Figure 5.6 shows the axial binning values. Specifically, the three columns correspond to the lower boundary of the bin, the upper boundary of the bin, and the centroid of the bin, respectively. Note that these values apply to each of the x- and y-bins though only listed once for the entire axial slice.

5.4.2 Radial Testing of the Triangular Mesh Tally

The entire reason for creating the “dht” tally was to allow for power tracking using a superimposed geometric mesh at a finer level than what is currently offered in Serpent. The following series of tests confirm that the results found using the new tally match those from conventional physical tallies currently implemented in Serpent. As previously discussed, hexagons in Serpent come in hexx and hexy varieties. In addition, lattice sizes can be even or odd. Odd lattices are the most straightforward for hexagonal tracking because the center of the lattice also corresponds with the center of a hexagonal element. All other elements are simply some integer

number of lattice pitches away. For the even lattice size case, the center of the lattice falls between lattice elements, which requires an extra half-pitch spacing to achieve the correct alignment. Since the spacing characteristics between even and odd lattice sizes use different treatments, it is important to test both to confirm that both perform as intended. Verify the functionality, 2x2 (group of three assemblies) and 3x3 (ring of seven assemblies) test geometries were modeled in Serpent using both a physical and “dht” tally. The two sets of results should be identical (in both mean value and reported statistical uncertainty) and are shown to be so. The test assemblies used in this section are homogeneously filled with a fuel-bearing material and are sized twice as large as AHTR assemblies (assembly apothem 45 cm, lattice pitch 93.8 cm).

5.4.2.a Even Lattice Size

A cluster of three assemblies was modeled in Serpent with a lattice size of 2×2^3 using both the hexx and hexy orientations. A visualization of the geometries can be seen in Figure 5.7. Note that the centers of the models do not coincide with the center of one of the hexagonal elements, a characteristic feature of using an even lattice size.

For reference, the numbering scheme of triangular elements in Figure 5.7 matches that shown in Figure 5.4. Assembly numbering is arbitrary but also proceeds in a similar fashion: the upper right assemblies is the starting point (Assembly 1) with progression rotating clockwise for the hexx geometry and counterclockwise for the hexy geometry. As evidence that the “dht” tally performs as intended, its results for the hexx geometry case are compared with those from a cell-

³Due to a known issue with the output of mesh indexing [62] in the distributed version 2.1.31 of Serpent, these results were skewed by one lattice index. The error likely stems from either excessive usage of a half-pitch positioning modifier or an incorrect starting index in a loop in the binning process (either score searching or output writing). This can be resolved by either shifting the center of the “dht” tally to accommodate this error or by simply using a larger superimposed lattice and being aware of the incorrect indexing. This work chose the latter option and actually used a 4x4 “dht” tally with indexes shifted down by one from what would normally be expected. As an additional note, this error was only observed to impact even-sized mesh lattices as part of this work. Others have observed issues with odd-sized meshes as well [62], but that was not the case with this work as the odd indexing performed as expected. There are two reasons this work did not correct this error as part of the implementation. First, an objective was to implement a modification of the existing “dh” tally. This error affects the “dh” tally as well, so for the sake of consistency between the existing feature and the new feature, no modification was made. Second, since this error is known due to being addressed in the developers’ discussion forum, it is likely a fix will be implemented in the next publicly available version of Serpent. Correcting the issue now might introduce the possibility of a double-correction in the future; whereby the indexing error is addressed twice (once by the developers and once through the implementation of the “dht” tally) and thus results in the same erroneous indexing shift but in the other direction. Taking no mitigation action now is the most likely path for the “dht” tally to be functionally compatible with a future release of Serpent.

based physical tally, shown in Table 5.3. As can be seen, the results are identical, as expected, and demonstrate that the “dht” tally functions properly. The results for the hexy geometry are not shown in the interest of saving space but also matched exactly

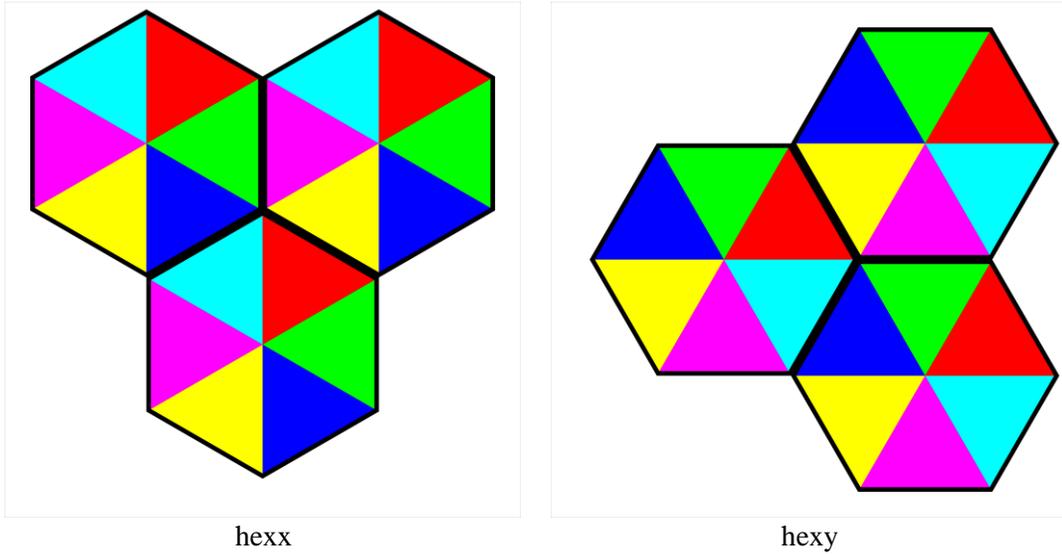


Figure 5.7. Test geometries for even lattice size.

Table 5.3. Tally comparison between using a cell-based physical tally and a superimposed triangular mesh tally for hexx type geometry test case using an even lattice size.

Assembly Triangular Index	Cell Physical Tally		Superimposed “dht” Tally	
	Tally Value	Rel. Uncertainty	Tally Value	Rel. Uncertainty
1-1	0.0402885	0.00064	0.0402885	0.00064
1-2	0.0400410	0.00060	0.0400410	0.00060
1-3	0.0402192	0.00057	0.0402192	0.00057
1-4	0.0401100	0.00061	0.0401100	0.00061
1-5	0.0401289	0.00058	0.0401289	0.00058
1-6	0.0403092	0.00064	0.0403092	0.00064
2-1	0.0402107	0.00069	0.0402107	0.00069
2-2	0.0401158	0.00059	0.0401158	0.00059
2-3	0.0401875	0.00069	0.0401875	0.00069
2-4	0.0402635	0.00059	0.0402635	0.00059
2-5	0.0400771	0.00071	0.0400771	0.00071
2-6	0.0402216	0.00067	0.0402216	0.00067
3-1	0.0401250	0.00061	0.0401250	0.00061
3-2	0.0402077	0.00064	0.0402077	0.00064
3-3	0.0402129	0.00063	0.0402129	0.00063
3-4	0.0402569	0.00062	0.0402569	0.00062
3-5	0.0402261	0.00064	0.0402261	0.00064
3-6	0.0401755	0.00065	0.0401755	0.00065

5.4.2.b Odd Lattice Size

A cluster of seven assemblies was modeled in Serpent with a lattice size of 3x3 using both the hexx and hexy orientations. A visualization of the geometries can be seen in Figure 5.8. Note that the centers of the models coincide with the center of one of the hexagonal elements, a characteristic feature of using an odd lattice size.

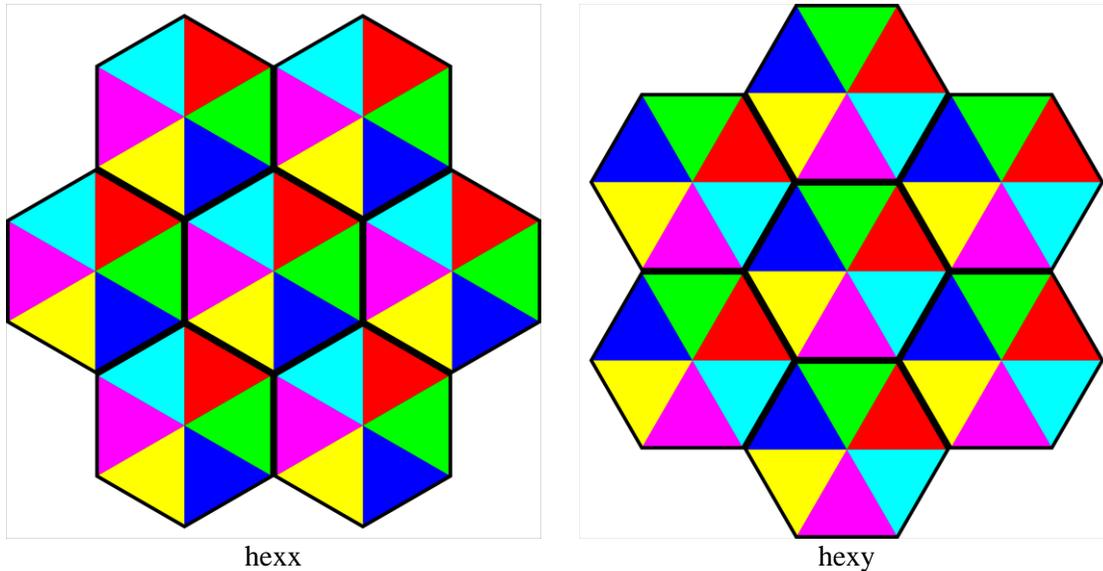


Figure 5.8. Test geometries for odd lattice size.

As with the even test cases, the “dht” tally was compared with the results of a cell-based physical tally. All 42 triangular elements matched exactly in both comparisons and in the interest of saving space, the results are not shown. This demonstrates that the “dht” tally functions properly now for both possible orientations.

5.4.3 Axial Testing of the Triangular Mesh Tally

An important component of the 3D AHTR Serpent model is capturing the axial behavior of the core in addition to the one-third assembly resolution made possible by the “dht” tally. Therefore, it is important to confirm that Cartesian axial binning in the axial direction works with the new tally. Testing was done on a single assembly divided into two axial segments for both hexx

and hexy orientations. Results were tallied using both the “dht” tally and a cell-based physical tally. The results for the hexx test case are shown in Table 5.4 for comparison. As can be seen, the two tallies are identical and demonstrate that the radial hexagonal “dht” tally and axial Cartesian “dz” tally are compatible. The hexy test also had identical agreement but results are not shown.

Table 5.4. Tally comparison between using a cell-based physical tally and a superimposed triangular mesh tally for hexx type geometry test case also using axial partitioning.

Axial-Triangular Index	Cell Physical Tally		Superimposed “dht” Tally	
	Tally Value	Rel. Uncertainty	Tally Value	Rel. Uncertainty
1-1	0.0406028	0.00056	0.0406028	0.00056
1-2	0.0396266	0.00050	0.0396266	0.00050
1-3	0.0405772	0.00057	0.0405772	0.00057
1-4	0.0405455	0.00052	0.0405455	0.00052
1-5	0.0396185	0.00054	0.0396185	0.00054
1-6	0.0406178	0.00054	0.0406178	0.00054
2-1	0.0404195	0.00057	0.0404195	0.00057
2-2	0.0395512	0.00059	0.0395512	0.00059
2-3	0.0404656	0.00063	0.0404656	0.00063
2-4	0.0403838	0.00050	0.0403838	0.00050
2-5	0.0394410	0.00057	0.0394410	0.00057
2-6	0.0404018	0.00048	0.0404018	0.00048

5.4.4 Testing Compatibility with the 3D AHTR Serpent Model

The 3D AHTR Serpent model uses hexy-type geometry for its assembly placement (Figure 5.3). In the TH feedback naming convention, assemblies are assigned numerical identifiers as shown in Figure 5.4. Note that Figure 2.2 may appear as hexx-type, but that is because it was rotated clockwise by 30° for ease of viewing since hexy-type lattice construction proceeds along a 30° incline. Consider a single axial slice of the AHTR core like that shown in Figure 1.2. In the 252 fuel assemblies, a triangular mesh tally would yield 1512 bins with results. Since results should be resolved at the one-third assembly level, this means there are 756 total sections after summing the corresponding triangular regions. If one-third core symmetry is imposed in addition to that, three symmetric sections will combine to produce once again 252 one-third core symmetric, one-third assembly regions of interest which are the net result of six bins from the triangular mesh implementation.

Testing was carried out with the 3D AHTR Serpent model using both a universe/material hybrid physical tally for and the new “dht” tally. For the physical tally, binning occurred at the assembly universe and fuel material level, which allowed for independent one-third assembly power tracking without one-third core symmetric contributions. The “dht” tally used a 19x19 radial partitioning, producing 2166 bins with 1512 returning scores. These were pair-wise summed to correspond with the 756 one-third assembly, core symmetry independent physical tally regions. Results in all 756 physical tally regions matched those from summing two adjoining triangular tally results. Further discussion of the cross-tally one-to-one mapping of results is addressed in Appendix C (page 258).

Through this series of tests, results for the newly implemented hexx- and hexy-type superimposed triangular mesh were compared with those obtained using existing cell-based physical meshes for even-sized, odd-sized, and axially partitioned cases. For all six cases, each region of interest exactly matched between the “dht” and cell-based results for both the reported value and its relative statistical uncertainty. The new tally implementation had a negligible impact on the runtime of test simulations, as was a major objective of its creation. It was also tested with the 3D AHTR Serpent model and successfully tallied one-third assembly fission powers to satisfy the other major objective of the work. These all demonstrate that for simple detector definitions, the newly created “dht” tally performs as intended. Its existence will enable one-third assembly fission power tallying of the 3D AHTR Serpent model using a superimposed geometric mesh. Additionally, its generality should make it applicable to other reactor designs which also use hexagonal assemblies or similar geometric features.

CHAPTER 6

NEUTRONICS METHODOLOGY AND CORE REACTIVITY CONTROL

Addressed here is the final topic of neutronics for discussion in a multiphysics system, though several parameters used for neutronic simulations have already been discussed as part of the previous sections. The geometry is dictated by the results of the thermal expansion. The material temperatures are given by the TH results. Material densities are obtained from the TH for liquids and thermal expansion for solids. Topics left to discuss include the sensitivity studies considered in making design decisions for AHTR analysis, specifying relevant data which will need to be obtained from neutronic simulations, and the types of simulations which will be used in the testing of this work.

The most relevant results coming from MC simulations for TH iteration are fission density and power distribution tallies (specifically for this work, total fission energy deposition). These tallies inform TH on the amount of thermal power (heat) generated at each location in the core which will then inform thermal expansion calculations to adjust structural dimensions. From the previous chapter, the type of tally used in a MC code can have an enormous impact on simulation run time. Tallying the fission rate at the one-third assembly level, in 84 one-third core assembly groups (252 fuel assemblies total), and sixteen axial zones (4,032 total partitions) showed that using physical tallies (such as material, cell, and universe) each slowed down the transport calculation by a factor of at least twenty versus using a superimposed geometric mesh. This penalty was too large to accept, and the newly developed triangular mesh tally will be used extensively in the remainder of this work.

While spherical europia burnable poisons have been discussed as a passive means of excess reactivity control, criticality is actively maintained by the movement of CBs. A CB movement scheme is established with the assumption that CBs are only moved in groups corresponding to

three symmetric assemblies. At BOC, it is established which groups of CBs are inserted to maintain criticality, and which remain withdrawn reserved for shutdown capabilities. Over the course of a fuel cycle, a single assembly grouping of three symmetric assemblies will be withdrawn at a time. Once full withdrawn, a new group will be moved to continue the process until the End of Cycle (EOC) is reached. The objective of this implementation is the capability to model the movement of CBs in support of the other components of this research endeavor. It is not intended for these control movement schemes to be fully optimized (by any definition, including: cycle length, fuel utilization, peaking factor minimization, or others) since it is not the focus of this work. Nevertheless, the schemes tested in this work are chosen in such a way that beneficially shape the power profile in the core over the burnup cycle. The focus lies in the capability itself with testing carried out by selecting favorable (but not necessarily optimized) insertion schemes.

6.1 Reactivity Sensitivity Studies

Several studies were conducted to get a better idea of the impact of moving control rod banks and how the reactivity changes over cycle. These studies inform design decisions from both a practical standpoint as well as being aware of configurations which may best showcase the capabilities of the developed procedure and the script that implements it.

6.1.1 Control Blade Reactivity Impact

This set of studies assesses the reactivity impact of CB insertion for various locations in the core. Moving radially out from the center of the core, one symmetric assembly group from each of the first eight “rings” of assemblies are considered. The list of assembly groups is 2, 4, 8, 14, 22, 32, 44, and 58, as indicated by the orange box in Figure 6.1. The corresponding symmetric core locations for each of the indicated assemblies highlighted in Figure 6.1 are also used in the study. The relative power generated by these eight assembly groups is shown in Table 6.1.

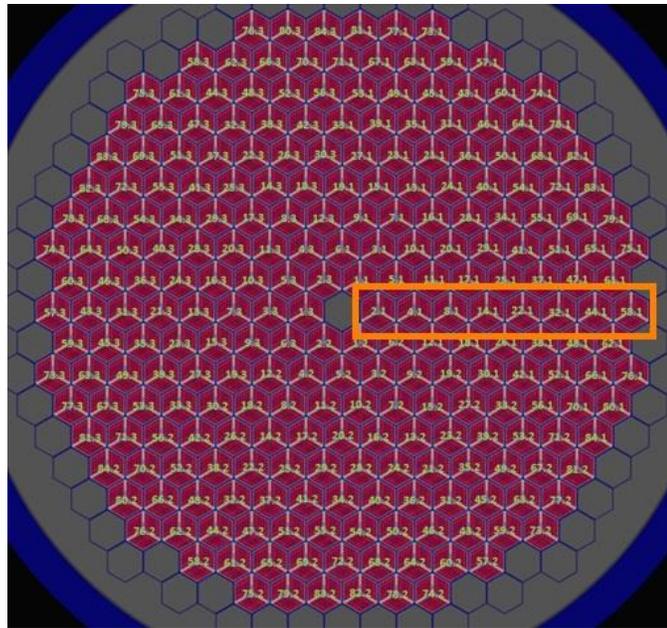


Figure 6.1. Eight assembly groups being assessed for control blade insertion.

Table 6.1. Normalized power (via fission energy deposition) generated by the assembly groups indicated in Figure 6.1 from each of the first eight radial rings in the AHTR core.

Radial Ring	Assembly Group	Normalized Power [-]
1	2	1.880
2	4	1.656
3	8	1.515
4	14	1.344
5	22	1.141
6	32	0.937
7	44	0.714
8	58	0.533

The reactivity impact from completely inserting the CBs into each radial ring location (CB bank worth) can be seen in Figure 6.2. An immediate observation is that CB insertion near the center of the core is more impactful than insertion near the periphery, which is what would be expected from neutron importance considerations. Another observation is that the second-ring assembly group experiences the largest reactivity change due to CB insertion. This may be due to a variety of reasons, which may include self-shielding for the central group since they are essentially lumped with the central moderating assembly. It could also be due to the first ring of assemblies having fewer fueled assemblies as neighbors, so there are fewer impacted fission sites in the immediate vicinity. Additionally, the difference at the center of the core is relatively small

and this observed difference in behavior could be due to statistics. In any case, the first three radial rings have comparable CB bank worth with decreasing worth when moving radially outward.

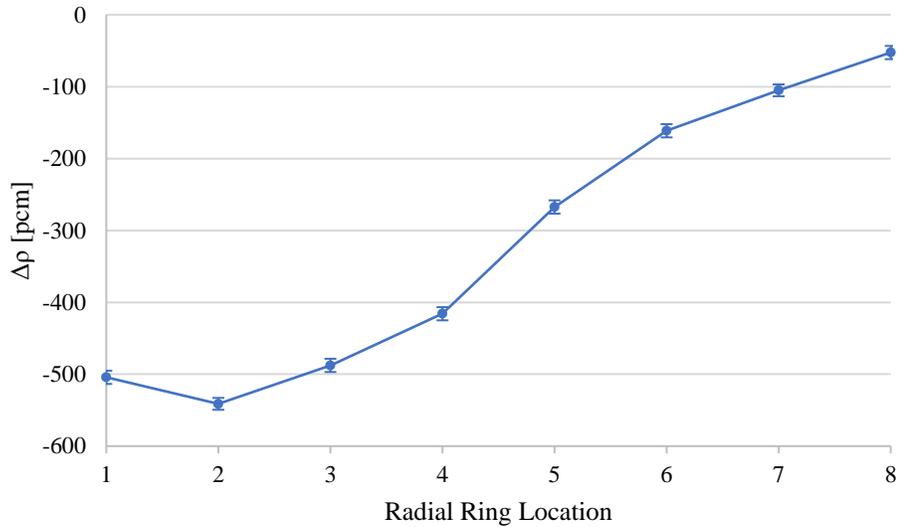


Figure 6.2. Control blade bank worth of a three-assembly group within radial rings of the first eight rings.

Figure 6.3 shows the integral rod worth of gradually inserting CB banks in increments of one-eighth of the active core height in each of the first eight radial rings of the active core. Like in Figure 6.2, the largest reactivity worth is seen in the first three radial rings of assemblies which have similar integral profiles. The trend in other radial ring groups is that rod worth decreases while moving toward the periphery.

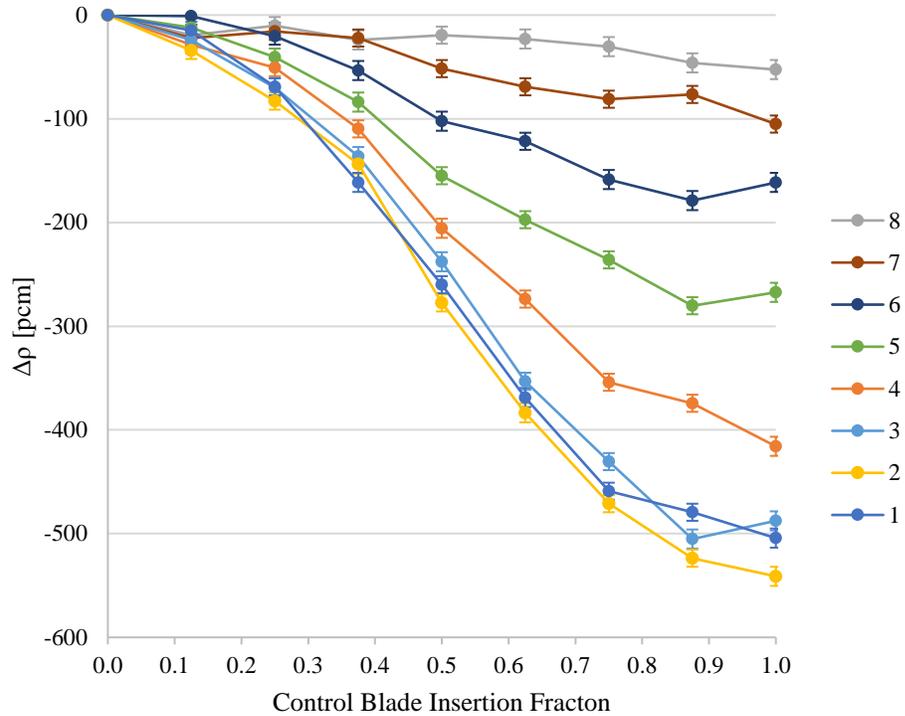


Figure 6.3. Integral rod worth of a representative three-assembly group from each of the first eight radial rings of assemblies of the active core.

6.1.2 Control Blade Axial Power Impact

The same eight radial ring groupings considered in the previous section were also studied for axial power effects. Figure 6.4 shows the axial offset (AO) of the core for CB movement. As expected, the three radial rings closest to the center of the core experience the largest AO change. The largest magnitude value always occurs, as expected, when the CBs are half-inserted. The equation used for AO can be seen in Equation 6.1. P_T and P_B refer to the thermal power produced in the top and bottom halves of the active core, respectively.

$$AO = \frac{P_T - P_B}{P_T + P_B} \quad (6.1)$$

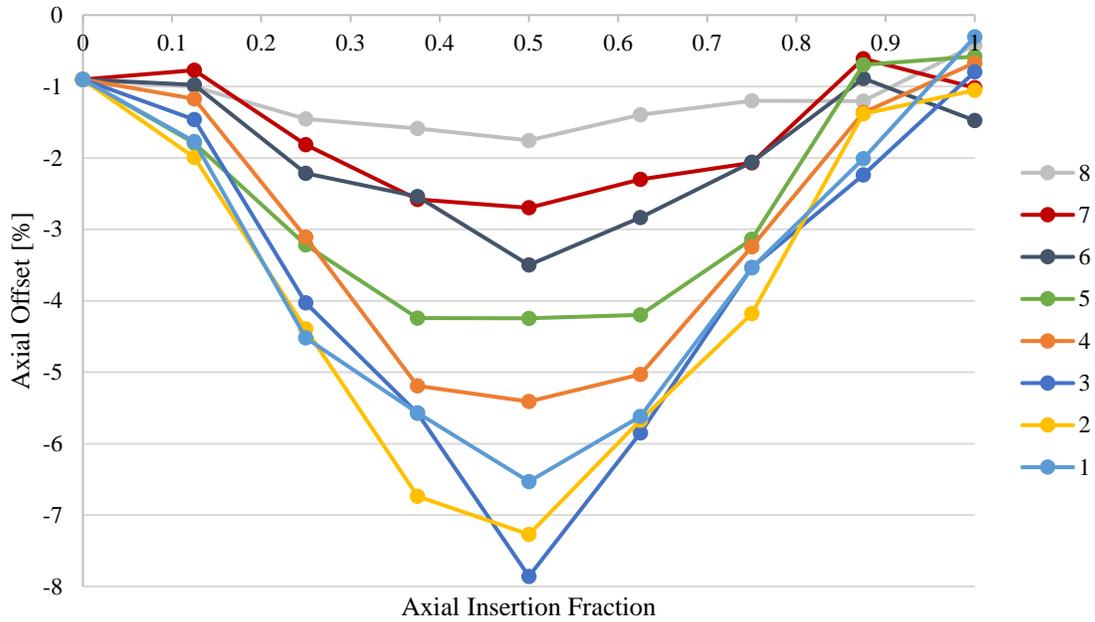


Figure 6.4. AO caused by one-eighth active core height increment insertion of control blade groups from specified radial rings.

Figure 6.5 shows the axial power profiles for the incremental insertion of the CB bank into the third radial ring group considered (assembly group 8). This bank was selected since it causes the largest observed AO in Figure 6.4. Other banks have similar behaviors but with smaller differences between the insertion fractions (due to lower AO). The AO changes from about -1% for fully withdrawn to about -8% for half inserted. The axial peaking is not 0% (symmetric) for the fully withdrawn and inserted cases due to the difference in the top and bottom core support plates composition (silicon carbide at the top and graphite at the bottom). Additionally, for the inserted cases, the CBs are present in the top axial reflector (due to descending from the top). This further contributes to axial asymmetry in the case of full insertion in which the active core region is axially symmetric. Axially varying temperature and coolant density were not considered as part of these simulations but would be expected to further impact the AO; further shifting the results toward the bottom half of the active core. There is very little change to the axial peaking factor however, as there is only about 2% increase to the axial peaking factor from about 1.39 to about 1.41 in this

worst-case from fully withdrawn to half inserted. Other assembly groups experience even smaller differences to the axial peaking factor with respect to CB movement.

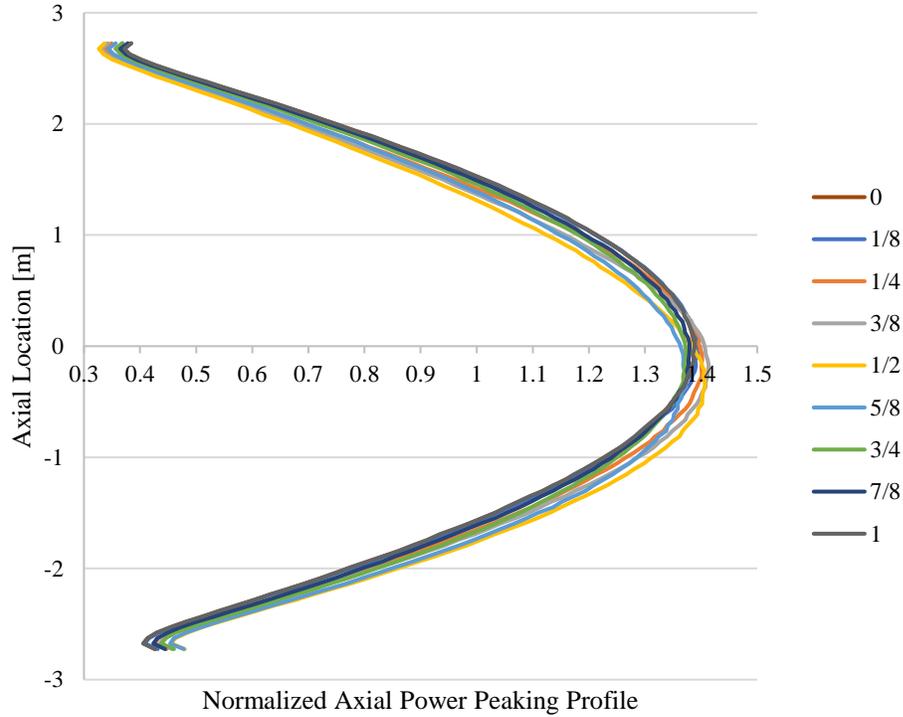


Figure 6.5. Axial power peaking profiles for one-eighth active core height increment insertion of control blade group in the third radial ring (assembly group 8).

6.1.3 Depletion Using Various Burnable Poison Densities

In most design specifications of the AHTR, the burnable poison spheres are composed of porous europia (Eu_2O_3) with density 5.00 g/cm^3 (68% of theoretical density of 7.42 g/cm^3). However, since it is assumed that the europia spheres can be manufactured with some level of porosity, it would be of interest to investigate additional densities as well. In Figure 6.6, four fresh fuel loading depletion cases are considered. They each use no, quarter, half, or full burnable poison density ($0.00, 1.25, 2.50,$ and 5.00 g/cm^3 respectively) with the reference model density as being “full” and not the theoretical density.

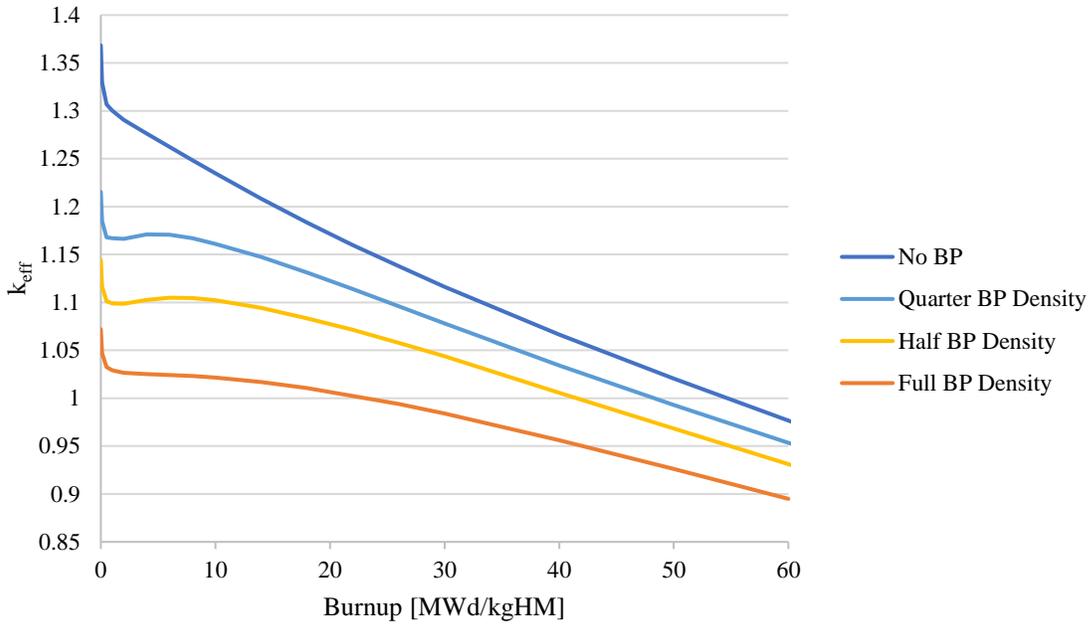


Figure 6.6. Change in eigenvalue with respect to burnup for various europia densities. Densities considered are 0.00, 1.25, 2.50, and 5.00 g/cm³, respectively.

One concern when using an integral burnable absorber in a reactor design is that since the absorbing material cannot be physically removed during operation, any material which is not burned up (i.e., transmuted to another nuclide with low absorption cross section) during the cycle becomes a residual reactivity penalty at EOC which needs to be compensated by having either a higher fuel enrichment or shorter cycle length (or some combination of the two). This directly translates to either a high fuel cost or more time offline due to more frequent refuel outages. These factors worsen the economic viability of a reactor design and should be avoided by engineered means. In Figure 6.6, one can see that the excess reactivity penalty at 60 MWd/kgHM due to the burnable poison spheres is quite significant for the poisoned cases considered (quarter, half, and full BP density) relative to the unpoisoned case (about 2300 pcm, 4600 pcm, and 8200 pcm respectively). Initial scoping of the AHTR system showed that for a reflected assembly, having the CB inserted as the only means of reactivity control was insufficient to suppress excess reactivity, with $k_{\infty} = 1.03204 \pm 0.00003$ [32] [42]. BPs are necessary for controlled operation but using the full reference density (5.00 g/cm³) europia might be too much. Possible solutions include using a lower-

density porous europia or selecting a different BP material altogether. One of the reasons why selecting a different material may result in a more favorable EOC residual reactivity penalty has to do with the neutron absorption chain of europium. Europium’s two naturally occurring isotopes (^{151}Eu and ^{153}Eu) each having large neutron capture cross sections in the thermal spectrum, but additionally the daughter nuclei from each of these (n, γ) reactions has even more potent absorbing properties than their respective parents. As can be seen in Table 6.2, ^{152}Eu and ^{154}Eu have larger capture cross sections than their parent nuclides with substantial half-lives (13.54 and 8.59 years, respectively) which means they will very likely capture again before decaying in an operating reactor setting. Even in the unlikely event of a decay, the decay daughters of samarium (for electron capture) and gadolinium (for β^-) also have significant neutron absorbing properties. In either case, this creates a scenario where the parasitic effects of europium lasts beyond the first capture event, which can help explain the large residual reactivity penalty observed in Figure 6.6. Alternatives to using a europium-based BP may include elements which have a less absorbing capture chain or burn out completely after the first capture event (such as boron).

Table 6.2. Isotopic data for europium [43].

Isotope	Natural Abundance [%]	Half Life [y]	Thermal σ_γ (at 0.0253 eV) [barns]
^{151}Eu	47.8	5×10^{18}	9169
^{152}Eu	0	13.54	12750
^{153}Eu	52.2	Stable	312.7
^{154}Eu	0	8.59	1353
^{155}Eu	0	4.76	3761

Another consideration for a depletion simulation is the number of burnable zones used radially for the tracked materials. Using a single zone averages the neutron absorption events over the entire material volume, which can neglect important geometric effects including self-shielding and spatially dependent burnup. The two principal materials requiring isotopic tracking in the AHTR design are the fuel and BP. A series of cases were run to investigate the impact of using multiple depletion zones, as can be seen in Figure 6.7. An automatic division feature in Serpent was used to segment each spherical material into equal volume radial partitions with the number of

specified zones. Results were compared to the baseline case of using a single depletion zone for both fuel and BP with the eigenvalue differences reported in the figure.

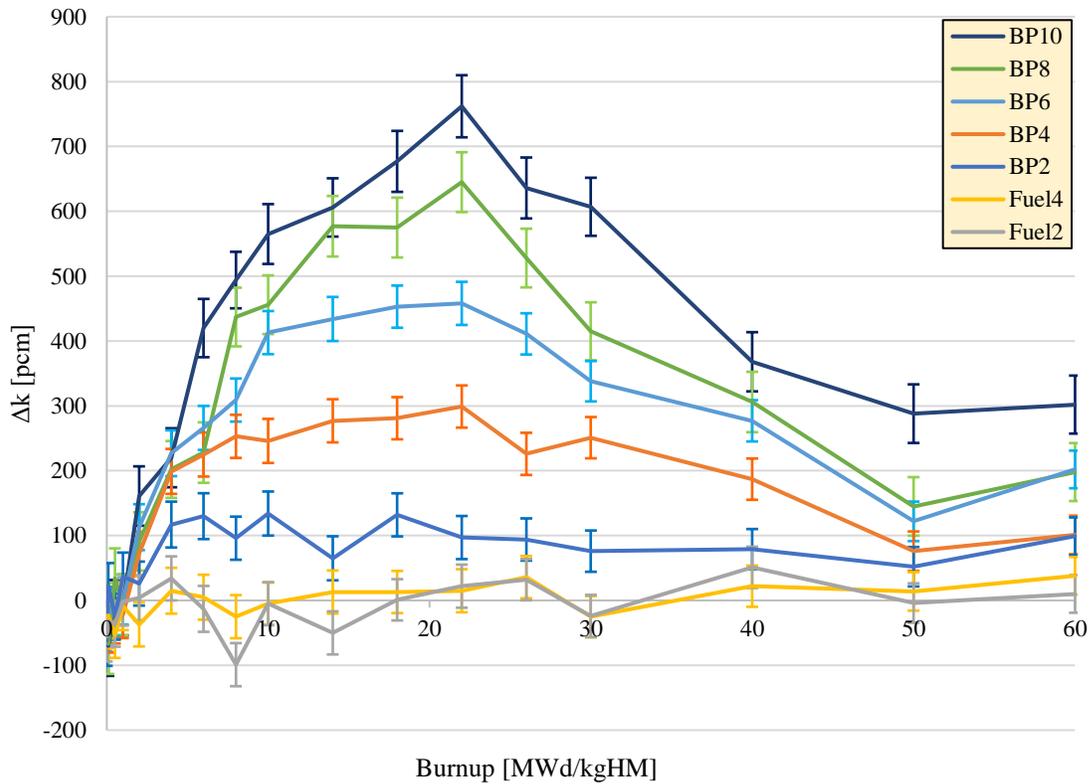


Figure 6.7. Eigenvalue differences due to the number of tracked material zones for both BP and fuel. All cases use full density europia (5.00 g/cm^3).

As can be seen in Figure 6.7, there does not appear to be a significant impact from using multiple depletion zones for the fuel. Two and four depletion zones were considered, and most differences could be justified by statistical uncertainty (about 35 pcm in this instance). This likely means that a single depletion zone for the TRISO fuel particles is sufficient and finer levels of partitioning were not considered after analyzing these sets of results. Further, this implies that due to the small size of the TRISO particles, self-shielding in fuel does not play as important of a role in the AHTR as opposed to other thermal spectrum designs such as LWRs which use larger volumes of lumped fuel.

In the case of using multiple tracked depletion zones for BP particles however, there does appear to be a strong impact on the eigenvalue. Two, four, six, eight, and ten depletion zones were considered for the BP particles. The general behavior is that eigenvalue differences increase up until a fuel burnup of about 22 MWd/kgHM. At that point, the differences diminish but remain non-zero when approaching an expected discharge burnup. The peak observed difference of 762 ± 48 pcm occurs for the case using ten divisions at a burnup of 22 MWd/kgHM. The eigenvalue differences grow with each successively finer level of division. As the number of divisions increases, there is no saturating effect evident in the cases considered. It is possible that further discretization would result in larger differences but is beyond the scope of this work. The important conclusion from this set of case studies is that the europium burnable spheres require division when tracking during depletion and if not done, results may differ by several hundred pcm.

6.2 Control Blade Movement Schedule for Core reactivity Control

Active criticality control is achieved by moving CBs. Selecting a favorable scheme of initially inserted CB groups with withdrawal order can beneficially shape the power profile over a depletion sequence. Keeping peaking factors low over cycle is generally desirable from both safety and fuel utilization standpoints. Due to the one-third radial symmetry of the 3D AHTR Serpent model, all considered insertions also follow one-third radial symmetry. This section addresses how CB insertion schemes are selected, their associated Power Peaking Factors (PPFs), and includes other performance metrics over cycle.

6.2.1 Reactivity Control by Using Full Density Burnable Poison Europa Spheres

The reference AHTR design assumes that the density of the europia BP spheres is 5.00 g/cm³. When no CBs are inserted, the effective eigenvalue of the core is 1.07532 ± 0.00005 . The assembly-wise PPF is 1.88 and the one-third assembly section-wise PPF is 2.05, with both occurring in assembly groups closest to the central reflecting assembly. A visualization of the radial power profile for the case of no CBs inserted can be seen in Figure 6.8. It and other radial power

profiles shown in this work were generated using serpentTools [44], a postprocessing utility developed at the Georgia Institute of Technology for working with a variety of Serpent output files.

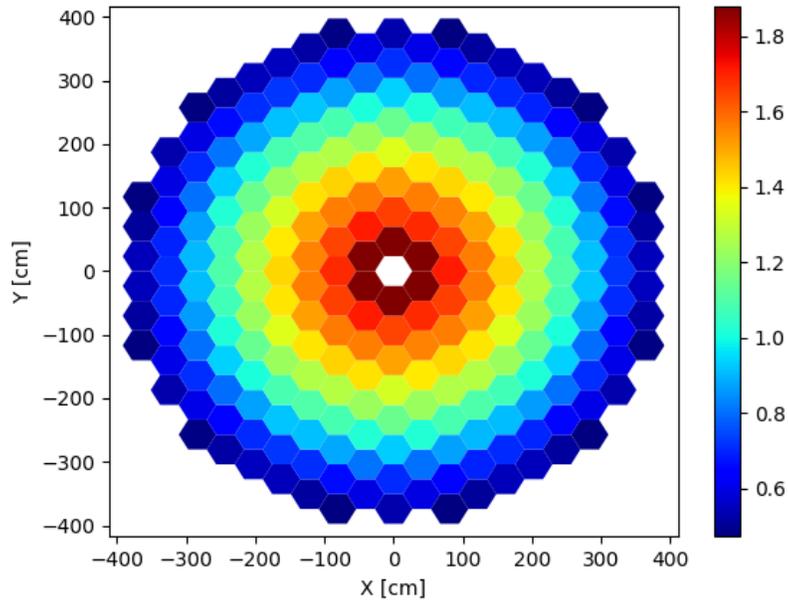


Figure 6.8. Radial power profile for full density europa BP and no CBs inserted.

This section considers four CB insertion schemes and further analyzes one in-depth. Visualizations of the four schemes can be seen in Figure 6.9. These schemes were created manually with the general guiding principle that no CB insertion has a directly neighboring assembly group which is also inserted. All four schemes in Figure 6.9 satisfy this.

As is a consequence of using one-third symmetric assembly group insertions, the four schemes considered in Figure 6.9 also have at least one-third symmetric CB insertion patterns. Additionally, the first two schemes also satisfy one-sixth symmetry in their patterns and further one-twelfth for Scheme 2. While higher orders of radial symmetry are not required, core designers may desire them. However, in the case of AHTR, they might not be feasible without other concessions such as having larger local or global power peaking factors. In general, from the experience of this work (both in this and other sections), one-third core symmetry for active core control is an adequate goal with more cost than benefit coming from trying to extend beyond it.

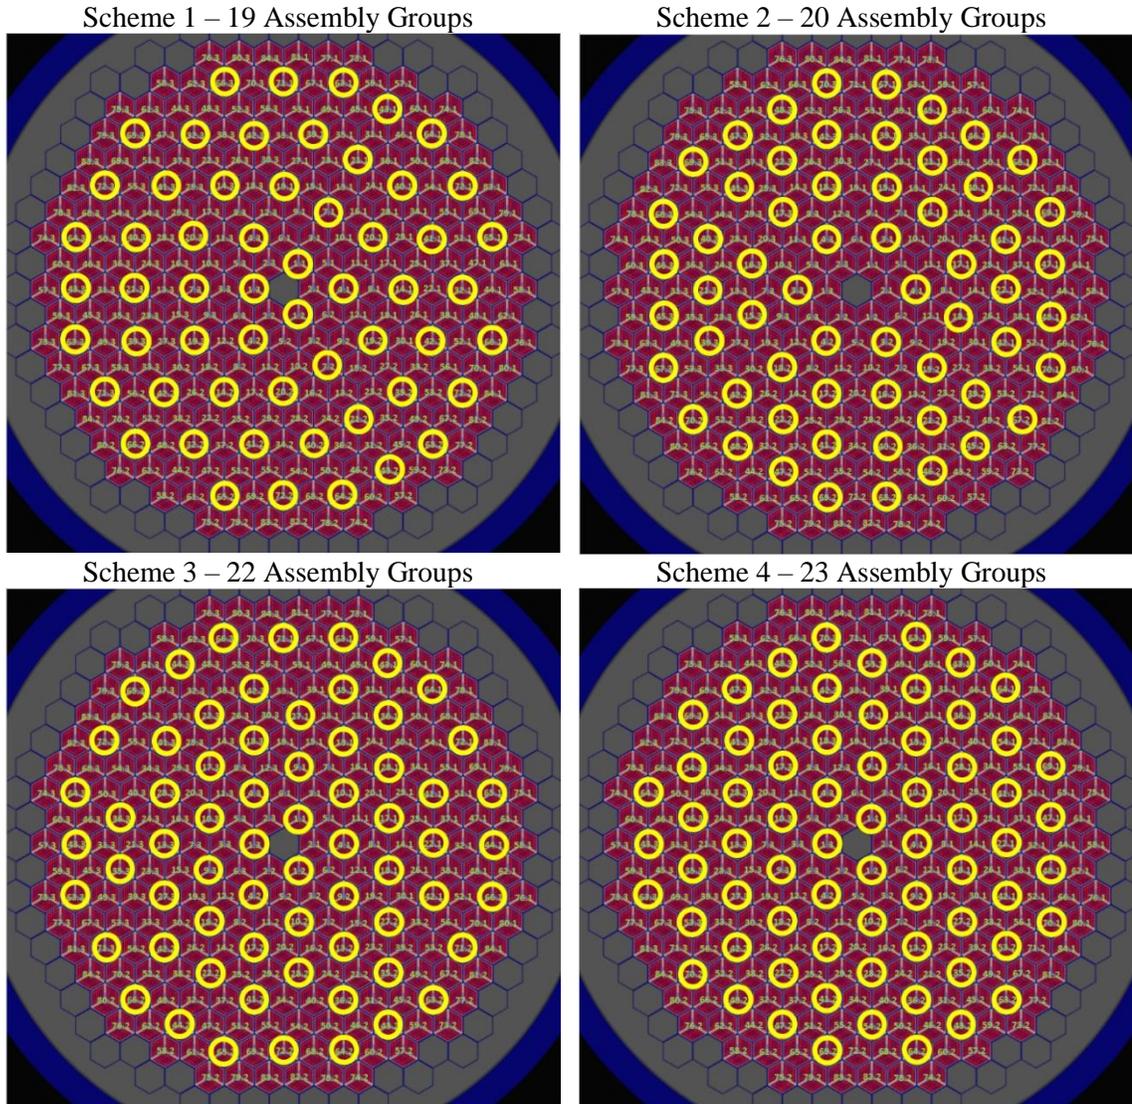


Figure 6.9. Four control blade insertion schemes considered for full density europa.

In Figure 6.9, Scheme 1 was created by using an in-out approach. Some intermediate radial positions used a less-than-tight packing arrangement to allow for both insertion near the central assembly and one-sixth symmetry for the remaining assemblies. This resulted in 19 assembly groups being inserted. Scheme 2 used an out-in approach while trying to maintain close packing where possible. This resulted in the six observable “triangular” zones with 20 assembly groups being inserted. Like Scheme 1, both Scheme 3 and Scheme 4 also used an in-out approach. Both also use non-adjacent close packing in the interior radial regions of the core. The difference between the two is that Scheme 3 favors a slight gap in the CB packing in the third assembly ring from the

periphery and Scheme 4 continues to use non-adjacent close packing to the periphery. Eigenvalue and PPF results for complete insertion of each control scheme can be seen in Table 6.3.

Table 6.3. Eigenvalue and radial power peaking factors for the four considered CB insertion schemes.

Scheme Number	CBs inserted	k_{eff} (marked rods in)	σ [pcm]	Assembly-Wise Maximum PPF*	Section-Wise Maximum PPF
1	57	1.01100	5	1.479	1.581
2	60	1.00968	5	2.660	3.015
3	66	0.99941	5	1.226	1.272
4	69	0.99435	5	1.314	1.413

*Assembly-wise statistical uncertainties: maximum 0.0049, minimum 0.0024, average 0.0035

Of the four schemes considered, only Scheme 3 and Scheme 4 were able to control excess reactivity using their indicated assembly groups. While Scheme 3 is technically subcritical when all its CBs are inserted, there is very little margin (59 ± 5 pcm) available. This is undesirable for a novel reactor design, which will operate using materials, configurations, and a neutron spectrum without extensive testing all together. A wise choice is to select Scheme 4 for further investigation since it has 565 ± 5 pcm of reactivity margin and acceptable power peaking characteristics, which is why the row is highlighted in Table 6.3. Visualization of the power profiles for each of the schemes can be seen in Figure 6.10.

Several broad conclusions can be drawn from the power profiles of Figure 6.10. From Scheme 2, one can see the importance of having CBs inserted in the first radial ring of assemblies. Due to the presence of the central reflecting assembly and only using a single fuel enrichment in fuel assemblies for this work, power production is expected to be quite high near the center of the core. This is confirmed by the power profile with no CBs inserted in Figure 6.8. Without CBs in one of the assembly groups in the first radial ring in the center of the core, the power peaking is very high. From Scheme 1, one can see the importance of having CBs packed close to the center of the core. Compared to Scheme 3 and Scheme 4, Scheme 1 has the highest PPF because not enough control is present close to the center. Scheme 3 has the most favorable power profile, and

it is because of the reduced CB packing in the third assembly ring from the periphery. The close packing from the center to the periphery in Scheme 4 explains its higher central peaking than Scheme 3. While Scheme 4 will be investigated further due to being the only scheme to have adequate reactivity control and meaningful margin, allowing the power to peak near the periphery (but not at the periphery) like in Scheme 3 results in a flatter overall profile.

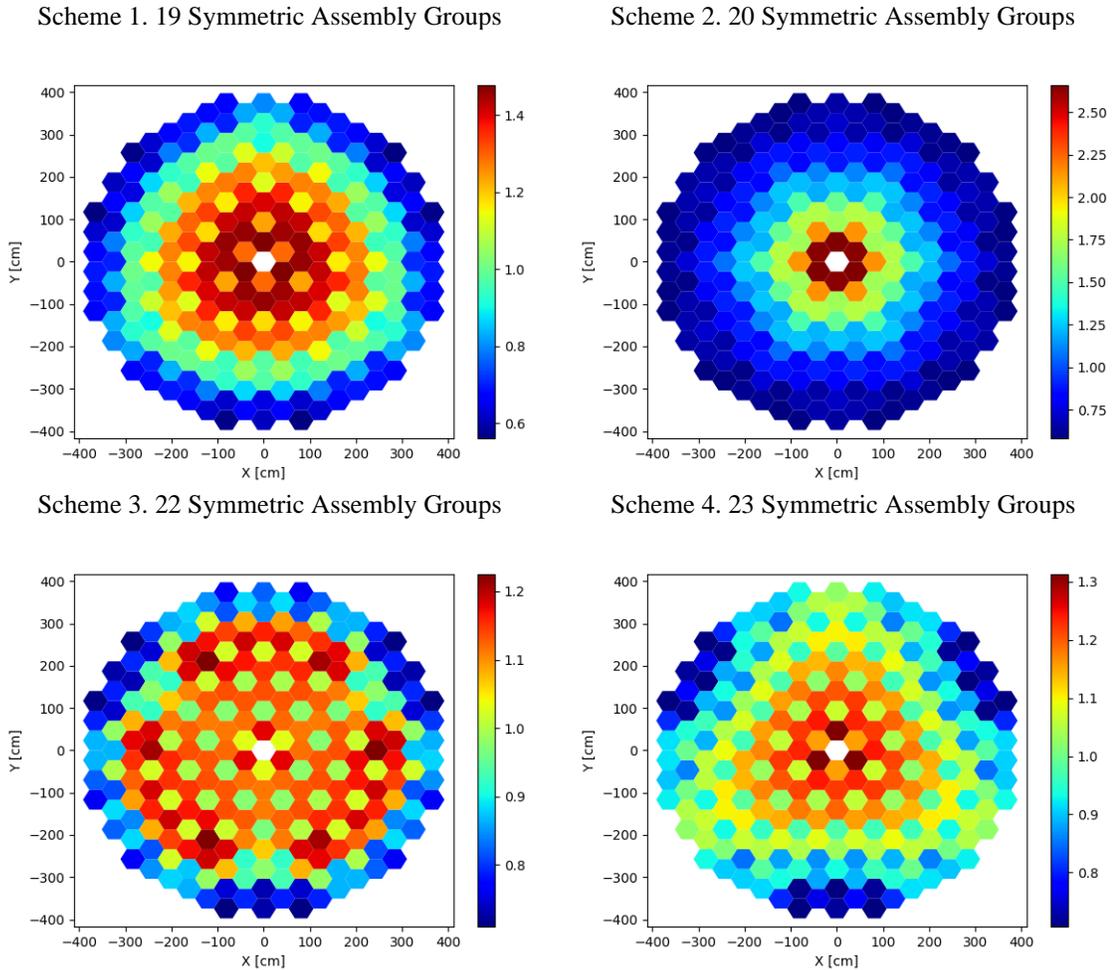


Figure 6.10. Radial power profiles for the four considered CB insertion schemes.

Using Scheme 4, the eigenvalue change from inserting CB assembly groups successively can be seen in Figure 6.11. The average worth of the 23 assembly groups is -352 pcm per group. The insertion order and other schedule details can be found in Supplementary Table D.1 of Appendix D (page 272).

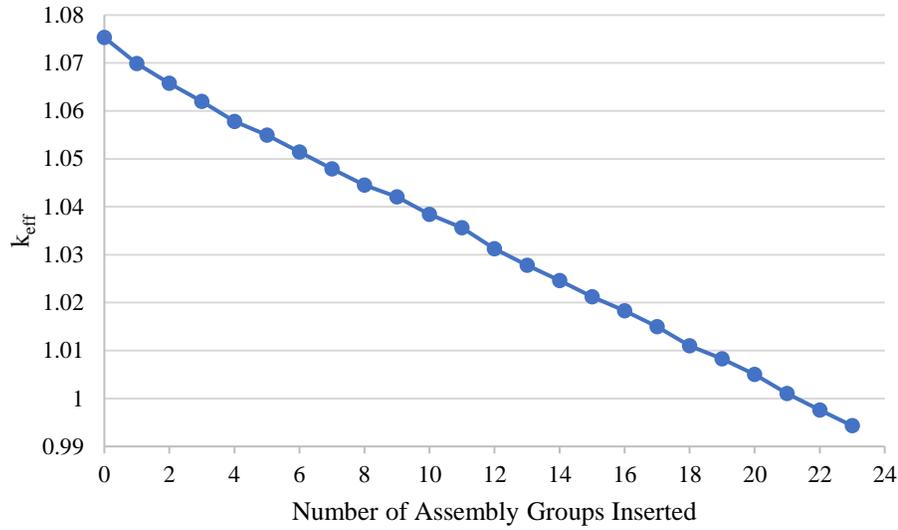


Figure 6.11. Effective eigenvalue impact from successively inserting CB groups.

Radial section- and assembly-wise PPFs for Scheme 4 can be visualized in Figure 6.12. Even just from the insertion of the first assembly group, both PPFs drop significantly. Full insertions keep the two considered PPFs between 1.2 and 1.5. Partial insertions would likely also stay in this range, with individual radial results likely between its respective full insertion/withdrawal results. Visualization of each power profile can be found in Supplementary Table D.2 of Appendix D (page 273).

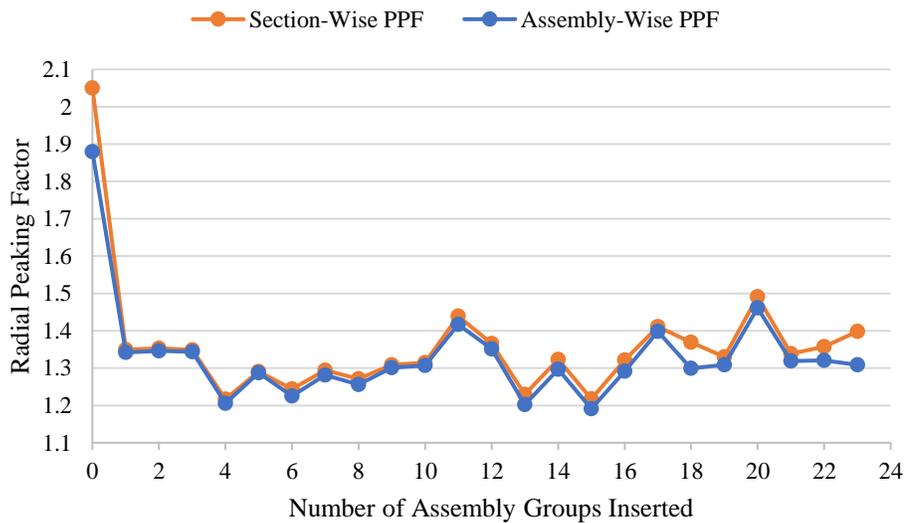


Figure 6.12. Radial section- and assembly-wise PPFs.

6.2.2 Reactivity Control by Using Quarter Density Burnable Poison Europia Spheres

As can be seen from Figure 6.6, there is significant residual reactivity penalty present in the AHTR fuel assemblies when using a BP density of 5.00 g/cm^3 . For this reason, the remainder of the work will assume a BP density one-quarter (1.25 g/cm^3) of that of the reference design. The results in the previous section, which use reference density europia should be informative for other researchers doing similar work, but it is evident that there are unresolved design issues with placing that much BP material in the reactor core.

The automated reactivity control schedule for moving CBs in the core is addressed in the proceeding section. As opposed to the reference density BP study in the previous section which used a manual CB selection technique for finding the next insertion location of a predetermined scheme, the next section uses an automated selection and simulation process. It was deemed distinct enough from the manual selection process to warrant a separate presentation and discussion.

6.3 Automated CB Selection and Insertion Schedule

For a given reactor core design, it can be uncertain where active reactivity control features should be positioned during operation due to the nature of burnup. If power is suppressed in the most reactive locations to achieve a flatter power profile early in the fuel cycle, power peaking might be even larger in those locations later in the power cycle as the active control is withdrawn. This type of analysis is complex as it is design-, geometry- (due to core location), and time-dependent (due to burnup).

For this analysis of the AHTR using single-enrichment fuel assemblies, the core loading pattern is static and the primary consideration for reactivity control over cycle is the CB insertion scheme and schedule. The CB locations and schedule are key to beneficially shaping the power profile over the cycle. For ease on behalf of the user, an automated CB schedule searching capability was added to the C++ script to simplify the process. The automated search runs a Serpent transport simulation to determine the next CB insertion location. The script then modifies the model

geometry to insert the CBs in the three one-third symmetric assemblies and repeat the process. A block diagram of how the process works can be seen in Figure 6.13.

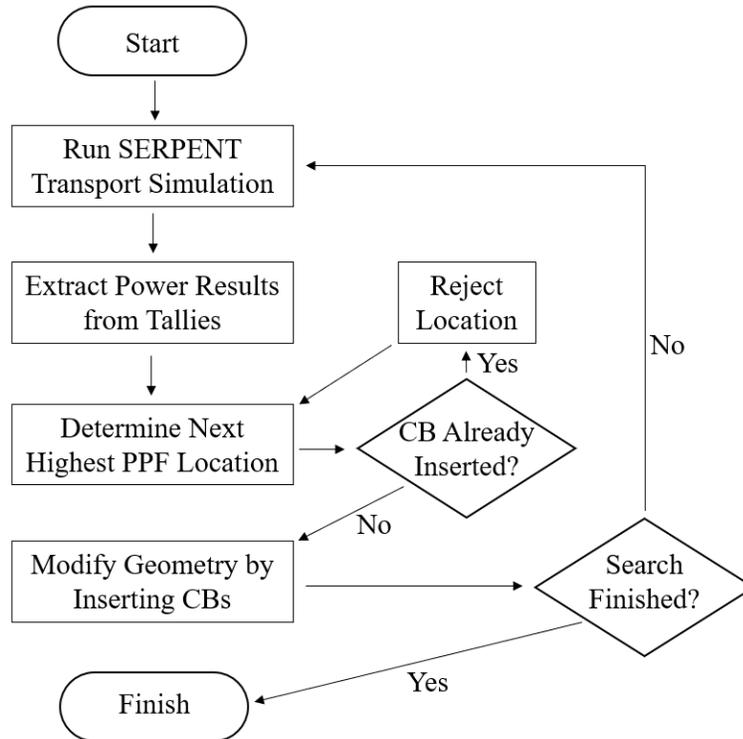


Figure 6.13. Conceptual flowchart of how automated CB insertion search process functions.

The “search finished” criterion in Figure 6.13 is satisfied by reaching one of two user-controlled options: inserting a prescribed number of assembly groups or bringing the eigenvalue within a user-specified threshold. For the work presented in this section, a prescribed number was used since it was desired to determine the insertion schedule for all assembly groups (despite not all being necessary to achieve criticality). As for the eigenvalue check, the search procedure extracts the eigenvalue for the previous transport simulation. If it is below some user-specified tolerance, the search ends. This value can be less than unity since the user might want to know insertion order beyond exact criticality for reasons including temperature feedback (thermal margin) and modeling uncertainty (design margin).

The capability was tested for the one-quarter reference density BP design. The starting basis was no symmetric CB groups inserted and proceeded until all 84 were inserted. In principle, the search would only need to progress until criticality is reached, but to fully showcase the capability, was carried-out for all 84 assembly groups.

The results from the automated CB insertion schedule selection can be seen in Figure 6.14 and Figure 6.15. Figure 6.14 shows how the eigenvalue changes for each successive CB assembly group insertion. The maximum change of -662 ± 7 pcm occurs when the most power-bearing CB group is inserted (next to central moderating assembly), and the minimum change of -146 ± 7 pcm occurs when the least power-bearing CB group is inserted (along the periphery). Overall though, the behavior is fairly linear, with the average worth of an assembly group being -382 pcm to achieve criticality (over 59 assembly groups) and -364 pcm for when all 84 assembly groups are inserted. Both of these values are higher than that of the reference-density europa case considered in the previous section (-352 pcm per assembly group). This is likely due to the softer neutron spectrum in the quarter density europa case due to the presence of less thermal absorbers, which would make the CB insertion more significant. The general linearity of the results is favorable for this work since reactivity changes due to CB withdrawal are well-behaved and predictable, which leads to an easier iterative search process as part of the excess reactivity control scheme of this work. A complete table of results can be found in Supplementary Table D.3 of Appendix D on page 278.

Figure 6.15 shows the radial PPFs at both the section and assembly levels. At the assembly level, results generally stay between 1.15 and 1.50 as the search approaches criticality (59 assembly groups inserted). PPFs are very large for both the cases of all rods full withdrawn (left) and all rods full inserted (right). This is expected from a core using a uniform fuel enrichment in all fuel assemblies, since it would closely follow a truncated-cosine-like profile. Visualizations of individual power profiles can be found in Supplementary Table D.4 of Appendix D on page 280.

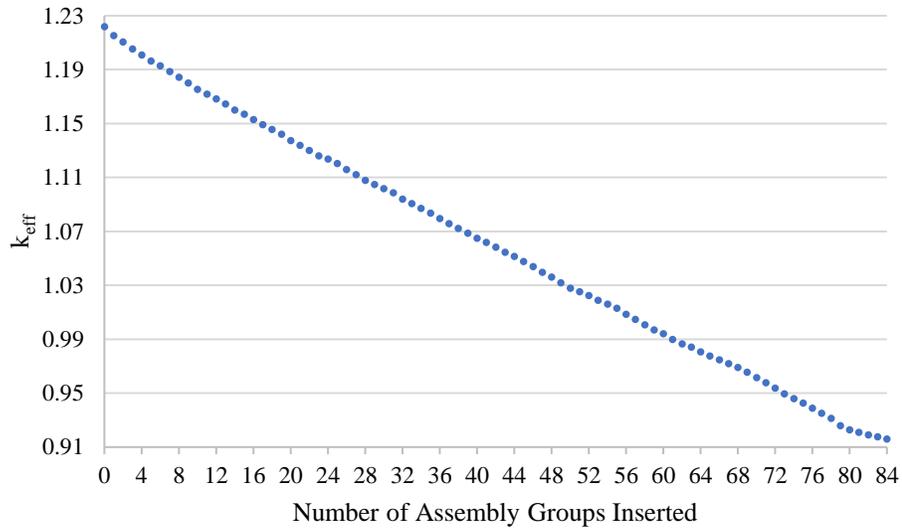


Figure 6.14. Effective eigenvalue impact from automated insertion of CB groups.

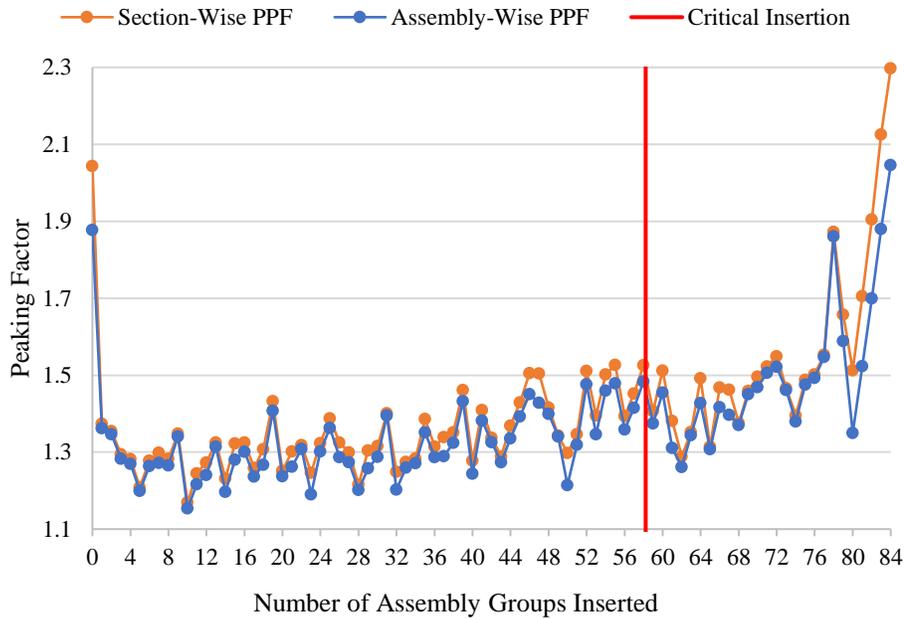


Figure 6.15. Radial section- and assembly-wise PPFs for automated insertion.

6.4 Automated Control Blade Withdrawal Description

In the previous section, the capability to have CB insertions automatically found by searching for the highest power assemblies was implemented. From a design standpoint, this can be beneficial to quickly find a critical core configuration which both controls excess reactivity and

has favorable radial power peaking properties. In practice however, this process is uncommon as CB insertions for operation are normally manually selected to have symmetry; similar to the scheme selection process of section 6.2.1. Under this excess reactivity control method CB insertion locations are preselected, and the remaining work is to determine the withdrawal order. This section introduces the capability to automatically withdraw CBs from a user-specified insertion scheme.

One CB withdrawal procedure was already introduced in the previous section. By finding a favorable CB insertion schedule, the reverse ordering of the schedule (i.e. last in, first out) would likely still have favorable radial peaking properties, but this is not guaranteed by any means. During operation of an actively controlled core, assemblies with CBs will deplete more slowly than for no insertion. The dynamics of power evolution stemming from burnup and CB movement are complex, and assuming the complementary CB withdrawal schedule corresponding to the CB insertion schedule might not be optimal. For this reason, other options should be available to the user for excess reactivity control management.

6.4.1 Details and Testing of Automated Control Blade Withdrawal Capability

The basis for selecting a CB withdrawal schedule in this work is minimizing the radial PPF during the search. It is achieved by withdrawing CBs from the lowest power assembly group among the inserted locations. The idea is that among insertion locations, withdrawing from the assembly group with the least power would be the most reasonable choice to keep radial peaking low. While not true optimization, which would require simulating the withdrawal of each individual CB group and selecting the one with the most favorable performance, this approach would shift power production to the location, which was previously carrying the least load. This process allows for decision making during operation using burnup-dependent neutronic properties.

To test the capability, the automated CB withdrawal search was implemented on the AHTR core using a europia density of 1.25 g/cm^3 in BP spheres and all CB groups inserted. One-third symmetric CB groups were serially withdrawn by selecting the group with the lowest relative

power among inserted groups. This search process is analogous to that conducted in section 6.3, except proceeds in the opposite direction. There is no guarantee that the insertion and withdrawal schedules of the previous section and this section, respectively, would provide comparable results, even without the complexity of burnup effects. Differences should actually be anticipated, stemming from factors including serial withdrawal in a search with many one-sixth core symmetric pairwise groups (arbitrary selection) and the fact that the withdrawal of the lowest power CB group does not always result in this location becoming the highest power assembly.

Eigenvalue change due to successive CB group withdrawals can be seen in Figure 6.16. The overall performance is similar to that of the automated insertions in Figure 6.14: criticality is reached with a partial insertion of the 59th CB group and the change is fairly linear. The average CB group worth is -382 pcm from critical insertion to all withdrawn and -364 pcm for all 84 groups (identical for those of Figure 6.14). The maximum eigenvalue difference observed was 645 ± 23 pcm and the minimum was 103 ± 23 pcm.

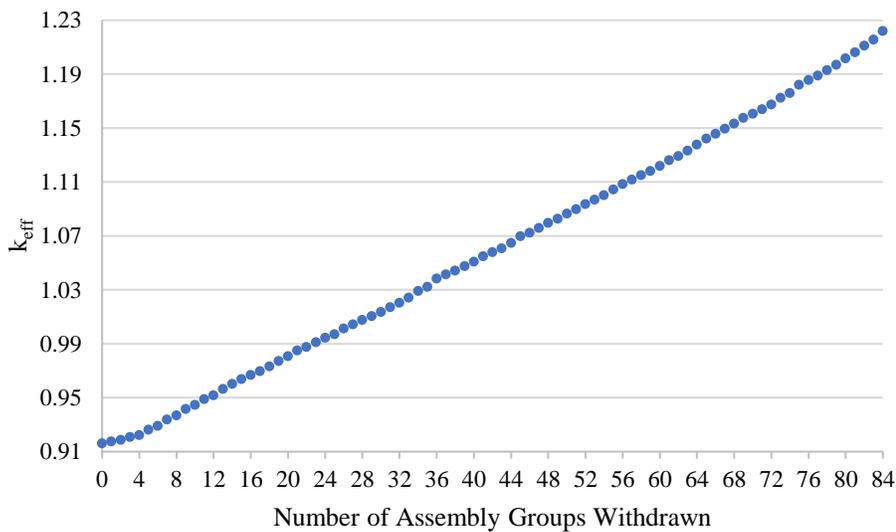


Figure 6.16. Effective eigenvalue impact from automated withdrawal of CB groups.

Radial section- and assembly-wise PPFs for automated CB group withdrawals can be seen in Figure 6.17. When compared to the PPF results of the automated CB insertion schedule search

in Figure 6.15, one can see that this set has generally higher values, which suggests worse performance. Differences were expected between the two, and it is not surprising that Figure 6.15 has lower PPFs. Inserting CBs into the assembly group with the highest power will undoubtedly reduce the power in that location. While the resulting profile could still have a higher peaking factor due to how the power shifts after the insertion, there is still a direct address to the assemblies responsible for the radial PPF. On the other hand, withdrawing CBs like in Figure 6.17 addresses the power peaking from a more indirect approach- withdrawing the inserted CB group with the lowest contribution to the power. This promotes the power to shift to the location, but there is less foresight on how this will impact the resulting power distribution than for insertions. This is not to discredit the CB withdrawal search method implemented here but serves as an advisement that although appealing for its on-the-fly capabilities with respect to power evolution due to depletion, following a user-prescribed withdrawal schedule instead could be more favorable. The question of which choice is superior is obviously case-dependent, but it should be recognized that the initial CB insertion schedule provided by the user will heavily impact the results of either approach.

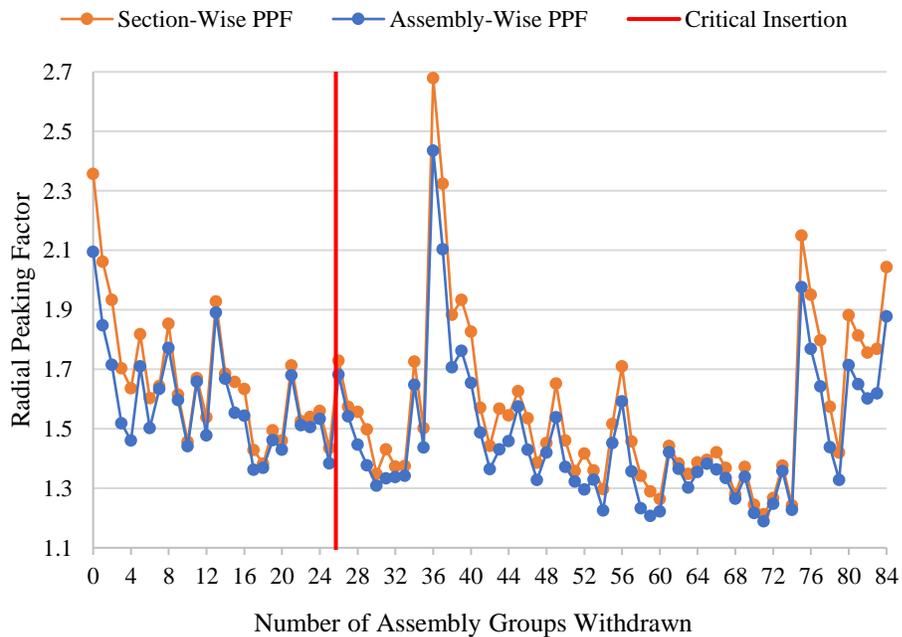


Figure 6.17. Radial section- and assembly-wise PPFs for automated withdrawal.

In Figure 6.17, the average assembly-wise PPF from critical insertion to complete withdrawal is 1.47. This is larger than the corresponding value of 1.32 from the automated CB insertion search of Figure 6.15. Section-wise average PPFs over the same set are also higher: 1.56 for automated withdrawal and 1.35 for automated insertion. Another interesting observation is how the ratio of section-wise to assembly-wise PPF varies between the two studies. For the automated insertions of Figure 6.15, the average section-wise results are only 2.3% larger while 5.7% larger for the automated withdrawals. This would suggest larger relative gradients across assemblies in addition to larger gradients across the core as a whole due to the higher PPFs.

6.4.2 Reaching Critical Core Configuration through Control Blade Movement

In section 6.4.1, the automated CB withdrawal capabilities were explored for the purpose of comparing performance with that of following the automated CB insertion schedule in section 6.3. The ultimate goal of moving CBs, of course, is to control excess reactivity in the core and have a critical configuration. This section discusses how that is accomplished.

When a statepoint or depletion simulation is launched that calls to use criticality iteration via CB movement, an initial CB group insertion guess is required. It is intended that inserted groups produce a subcritical configuration to converge on the critical insertion configuration sooner, but this is not required. When the initial simulation completes, a CB group will be withdrawn and another simulation will be run with the updated geometry. The selected group will be the next in the user-specified schedule. When this simulation completes, the eigenvalue will be extracted. If the eigenvalue is outside the tolerance for target eigenvalue, it is established that another configuration will need to be investigated. The next critical insertion configuration (I_{i+1}) is guessed by Equation 6.2. It depends on the critical insertion guess from the current simulation (I_i) and previous simulation (I_{i-1}) as well as eigenvalues from the current simulation (k_i), previous simulation (k_{i-1}), and the target eigenvalue (k_{target}). Due to the nearly linear reactivity worths of CB groups as seen in Figure 6.14 and Figure 6.16, this method is able to find the critical insertion

configuration using only a few search substeps. When the user-specified eigenvalue tolerance is satisfied, the current insertion scheme is accepted.

$$I_{i+1} = I_i - \frac{k_i - k_{target}}{k_i - k_{i-1}} (I_i - I_{i-1}) \quad (6.2)$$

Here are two comments on how this procedure works. First, the user-specified eigenvalue tolerance needs to be sufficiently large since only discrete insertions of $1/16^{\text{th}}$ of the active core height are considered. If too small, it is possible that the search will not be satisfied and will alternate between two insertion positions. Statistical uncertainty could become a factor here but will not be addressed further. It is intended that convergence within a reasonable tolerance will be acceptable to the user. Previous studies in this chapter showed that CB group worths were usually about 300-400 pcm, giving an average section worth of about 20 pcm. Tolerances could be this low, but the user may be satisfied with larger values which would also require fewer iterations to converge on the critical insertion. Second, the target eigenvalue does not necessarily need to be one. Especially in considerations of burnup, it may be desirable to have the eigenvalue be greater than one to obtain a control scheme more closely matching the burnup step-average eigenvalue rather than the eigenvalue at the beginning of the step. In other cases, it may be desirable to obtain a subcritical configuration from a safety standpoint matching some prescribed threshold. Regardless of the application, the desire to obtain a non-unity eigenvalue may exist for the user and the capability is implemented in the script.

CHAPTER 7

USING THE C++ SCRIPT ATOMICS

Up to this point, much of this work has focused on the methods and their integration into a methodology and implementation within the developed AHTR script. Many of the functionalities in the previous chapters were originally coded as standalone features, which were then combined into a single utility. This chapter will focus on detailing how the C++ script developed for this dissertation functions. The primary features of the script are of course its multiphysics capabilities to account for thermal, neutronic, and material properties feedback in AHTR as well as the CB critical search features to find CB insertion schedules, CB withdrawal schedules, and achieve system criticality by iterating through a user-provided insertion schedule. Based on these features, the C++ script was given the name ATOMICS for Ahtr Thermal behaviOr Modeling and Iterative Criticality Suite.

Documentation beginnings with how users can interface with ATOMICS via an input file and then how it functions as a whole. Individual sections of the user input file are documented and discussed in this chapter with screenshots capturing most of the relevant text. A complete sample user input file can be found in Appendix E (page 297).

7.1 Modifying How the Script is Executed via the User Input File

Many of the parameters of ATOMICS are controlled by a user input file. When the code is executed, the first task that is done after basic initialization is to read the contents of this file and make modifications to the model based upon the user input. In total, the file is 121 lines long and allows for dictating instructions in five broad areas: selecting the script mode, modifying the geometry, supplying material data, providing TH parameters, and choosing depletion and criticality options. Each of these areas will be address in further detail to follow.

There are some possible input values or combinations of values, which will produce errors when Serpent is trying to process the model geometry. To avoid these errors, there are logical checks in the script to help inform the user of possible issues, such as: inputting an invalid number for an option; giving a negative value for a physical quantity such as dimension, density, or thermal conductivity; providing a value other than 0 or 1 for a Boolean quantity, etc. These warnings or cautions to the user can be found in the terminal output which echoes back the user input file values to the user. Comments will be made where applicable in the documentation, but for combinatory geometry definitions like those used by MC codes, some values may define impossible cells which will cause Serpent to fail when creating the model geometry. A common example of this would be having an inner surface larger than an outer surface (such as in the TRISO particles). The script does not prevent or warn of all instances of such inconsistencies, and it is ultimately user's responsibility to be judicious when making modifications.

7.1.1 Selecting the Module Type for the Script

The first block of the user input file deals with high-level options which dictate how the script is executed. An example of this block taken from the input file can be seen in Figure 7.1.

```
1 === User Input File for AHTR Script ===
2
3 -SCRIPT MODE AND BASIC PHYSICS
4 8 // Script Mode. 0 - Statepoint, 1 - CB Insertion, 2 -
5 0 // Use cold dimensions. 0 - False (Thermal Expansion)
6 10000 // Statepoint Particles per Cycle
7 500 // Statepoint Active Cycles
8 100 // Statepoint Inactive Cycles
9
```

Figure 7.1. Input file portion pertaining to selecting modes for running the script.

Five parameters are controlled within this block, the first of which is the script mode (line 4). There are ten possible options:

- Option 0: Statepoint calculation
- Option 1: CB insertion schedule search

- Option 2: CB withdrawal schedule search
- Option 3: Iterative criticality search via CB movement
- Option 4: TH iterative search
- Option 5: Criticality and TH iterative search
- Option 6: Depletion without criticality search and no TH iteration
- Option 7: Depletion with criticality search and no TH iteration
- Option 8: Depletion without criticality search and TH iteration
- Option 9: Depletion with criticality search and TH iteration

The two CB schedule searches (Options 1 and 2) find CB movement orders and are meant to be standalone runs of the script to find the desired schedule. For running simulations of the core beyond CB movement, the modes of choice are either statepoint (Options 0, 3, 4, and 5) or depletion (Options 6, 7, 8, and 9) calculations. It is expected that most simulations would fall in one of these two groups of modes, since CB schedule searches would likely be conducted only once for a given core configuration and applied to all subsequent simulations using that configuration. The second parameter in this block is whether to use thermal expansion (line 5). This impacts the use of temperature-dependent dimensions and densities for solid structures of the AHTR fuel assemblies. Option 0 uses thermal expansion and selecting Option 1 assumes use of cold dimensions and densities (no thermal expansion). Note that this control option does not impact the use of axially-dependent coolant density due to temperature differences. That is controlled further down the user input file in the TH portion. The last three parameters in this block are related to statepoint simulation neutron histories. Line 6 controls the number of particles per generation. Line 7 controls the number of active cycles. Line 8 controls the number of inactive cycles. Later on in the input block related to depletion, there will be a similar set of input parameters related to depletion simulations. The two are distinct because there might be a user desire to run fewer particles for

statepoint simulations within a depletion sequence (criticality iteration and/or TH iteration) than for a depletion step.

7.1.2 Controlling Geometry Features

The second block of the user input file dictates the dimensions of components of the AHTR fuel assemblies and core. An example of this block taken from the input file can be seen in Figure 7.2. The geometry block is divided into seven subgroupings: fuel particle, fuel lattice, BP lattice, planks and assembly, reflector assemblies, axial discretization, and regions beyond the active core. Each of these will be discussed individually.

```

9
10 -GEOMETRY
11 --Fuel Particle
12 0.02135 // Cold Fuel Kernel Radius [cm]
13 0.03135 // Cold Buffer Layer Radius [cm]
14 0.03485 // Cold Inner Pyrolytic Carbon Layer Radius [cm]
15 0.03835 // Cold Silicon Carbide Layer Radius [cm]
16 0.04235 // Cold Outer Pyrolytic Carbon Layer Radius [cm]
17 --Fuel Lattice
18 1 // Particle Lattice type. 0 - cuboidal, 1 - cubic (wrt z, ov
19 4 // Width of Fuel Stripe (in Layers)
20 202 // Length of Fuel Stripe
21 0.09406 // Cold X Lattice Pitch [cm]
22 0.09128 // Cold Y Lattice Pitch [cm]
23 0.09266 // Cold Z Lattice Pitch [cm]
24 --Burnable Poison Lattice
25 1 // Burnable Poison Sphere Usage. 0 - None, 1 - Use
26 0.035 // Cold Poison Kernel Radius [cm]
27 0.09936 // Cold Poison Kernel Axial Pitch [cm]
28 5 // Number of BP columns (must be odd)
29 40 // Integral pitch spacing along the fuel plank
30 --Planks and Assembly
31 2.55 // Cold Plank Width [cm]
32 0.1 // Cold Sleeve Width (distance from coolant to fuel stripe)
33 14 // Cold Distance Between Spacers [cm]
34 22.5 // Cold Assembly Apothem [cm]
35 46.8 // Cold Assembly Pitch [cm]
36 --Reflector Assemblies
37 1 // Reflector Assembly Central Cooling Hole Usage. 0 - False,
38 2 // Reflector Assembly Central Cooling Hole Cold Radius [cm]
39 22.5 // Reflector Assembly Cold Apothem [cm]
40 --Axial Discretization
41 16 // Number of Modeled Axial Partitions in Active Core (up to
42 550.02976 // Cold Active Core Height [cm]
43 25 // Cold Height of Top/Bottom Axial Reflectors [cm]
44 35 // Cold Height of Top/Bottom Axial Support Plates [cm]
45 1 // Axial Flowering / Interassembly Expansion. 0 - False (use
46 --Beyond Active Core
47 478 // Cold Permanent Radial Reflector Outer Radius [cm]
48 479 // Cold Boron Carbide Layer Outer Radius [cm]
49 481 // Cold Core Barrel Outer Radius [cm]
50 519 // Cold Downcomer Outer Radius [cm]
51 520 // Cold Vessel Liner Outer Radius [cm]
52 525 // Cold Pressure Vessel Outer Radius [cm]
53

```

Figure 7.2. Input file portion pertaining to geometric features, options, and dimensions.

7.1.2.a Fuel Kernel

Lines 12-16 of the user input file give the dimensions of the TRISO fuel particles. Each entry provides the cold outer radius (in units of cm) for the particle layers from the inside-out, with the ordering of each successive layer matching that of AGR-2 type TRISO. This means that the layers move from the central fuel kernel, through the buffer, inner pyrolytic carbon, silicon carbide, and finally outer pyrolytic carbon layers. Note that both the number and order of the layers is hard coded into the script. This means that using fewer or additional layers in the particles (such as BISO or QUADRISO particles, respectively) is not currently supported and would need to be manually changed in the particle geometry portion of the Serpent input file generated by ATOMICS. Moreover, if successive layers were sized in an unphysical manner (inner layers sized larger than inner layers), then Serpent would likely either crash or run in an unexpected way due to the combinatorial construction of the unphysical region(s).

7.1.2.b Fuel Lattice

Lines 18-23 of the user input file deal with options and dimensions related to the TRISO fuel lattice. Line 18 allows for use of a cuboidal or cubic TRISO particle lattice. If cuboidal (option 0), the TRISO pitches in the x -, y -, and z -directions are independent. If cubic (option 1), the pitch in the z -direction is assumed for the other two dimensions. The most common reason to use the cubic option is to speed-up the simulation or to use an axial partitioning fraction other than 16. The reason for this restriction is that Serpent constructs cuboidal lattices in an explicit fashion, which makes them much less flexible than cubic lattices which can more easily fill an arbitrary space. Thus, cuboidal TRISO pitches are only supported for models using 16 axial partitions. Line 19 assigns the integer width of the TRISO fuel stripe in the number of layers. Note that at least one layer is necessary (otherwise will have an unfueled core) up to a maximum of 12 layers with the current TRISO dimensions (otherwise the fuel stripes will not fit within the dimension of the reference fuel plank). Line 20 assigns the integer length of the TRISO fuel stripe in the number of

particles. Note that at least one particle length is necessary (otherwise will have an unfueled core). The maximum dimension is 225 particles for a cuboidal lattice (since explicitly constructed and ATOMICS is hard coded to only handle up to this number) and theoretically the width of the fuel plank for a cubic lattice. Unexpected combinatorial behavior will result for fuel stripes extending beyond the length of the fuel plank. Lines 21-23 assign the pitches of the TRISO particles in the x -, y -, and z -directions respectively (in unit of cm).

7.1.2.c Burnable Poison Lattice

Lines 25-29 of the user input file deal with options and dimensions related to burnable poisons at the center of each fuel plank. Line 25 controls the usage of BP spheres (option 0 for none, option 1 to use). Line 26 gives the cold radius of the poison kernel (in units of cm). Line 27 gives the cold axial pitch of the poison lattice (in units of cm). Line 28 gives the number of BP columns to be used (reference number of columns is five). By construction in ATOMICS, the number of BP columns must be odd and the middle column will be positioned at the center of the fuel plank. Line 29 gives the number of integer axial pitches used to separate each BP column. The reason for this admittedly less-than-intuitive construction is for ease and computational speed of modeling the BP columns for each plank from a single BP cubic lattice. Each column is bound by a cylinder to “cut” a stack from the infinite cubic lattice. By using columns an integer number of pitches away, a single lattice can be used for all five poison columns in a fuel plank. This distance is nominally 4 cm in the original ORNL AHTR description, but here is changed to 40 times the axial BP pitch (assumed to be 0.09936 cm cold) or 3.9744 cm cold. This is about the same spacing, which is likely rather arbitrary, and this implementation cuts down on the modeling overhead necessary to create the problem geometry. Since the user can change the axial pitch of these poison particles, the plank spacing can be changed, too. For example, if the BP axial pitch is assumed to be the same as that of the TRISO particles (0.09266 cm cold), then 43 would be a better value for

the plank spacing (resulting dimension of 3.98438 cm cold) to come close to the 4 cm nominal value.

7.1.2.d Planks and Assembly

Lines 31-35 of the user input file deal with parameters used to create the AHTR fuel assemblies. Line 31 gives the cold width of the fuel plank (in units of cm). The reference cold width is 2.55 cm. Note that the plank width must be large enough to accommodate two sleeve widths and two fuel stripes, but must also not exceed 3.25 cm, which is one-sixth of the width of the region where fuel planks and coolant channels are located. If sized too large, there will be no coolant channels and planks will overlap each other, causing Serpent to crash. Line 32 gives the cold sleeve width (distance from the coolant to the fuel stripe, in units of cm). Line 33 gives the cold distance between the plank spacers (in units of cm). Line 34 provides the cold assembly apothem (in units of cm). Note that if this dimension is too small, there will not be a wrapper around the assembly (lower limit of 21.5 cm cold). Additionally, the assembly cannot be larger than the assembly pitch (upper limit of half the cold assembly pitch, which is 23.4 cm). Violating the lower limit should result in geometry issues in Serpent (causing it to fail) while violating the upper limit should result in having assemblies spaced by the assembly pitch with no interassembly gap and truncation of structural carbon-carbon material beyond the assembly pitch. Caution is warranted here, as the assembly size will essentially be overridden by the assembly pitch. Line 35 changes the cold assembly pitch (in units of cm). If made too large, the radial core layout will extend too far into the permanent reflector region and result in some assemblies being partially “cut” by the boron carbide and core barrel features. It would be necessary to also adjust these features beyond the active core accordingly if this issue is encountered.

7.1.2.e Reflector Assemblies

Lines 37-39 of the user input file control options related to the reflector assemblies along the periphery of the active core as well as the central assembly. Line 37 determines whether coolant

holes are modeled at the center of each reflector assembly (option 0 for no hole, option 1 to use coolant holes). Line 38 gives the cold radius of this reflector assembly central cooling hole (in units of cm). Line 39 changes the cold apothem of the reflector assemblies (in units of cm). In the reference design they are the same size as the fuel assemblies, and face similar maximum dimensional constraints as discussed above with the fuel assembly apothem (Line 34).

7.1.2.f Axial Discretization

Lines 41-45 of the user input file control how the 3D model of the AHTR is axially discretized. Line 41 dictates how many axial partitions are used when segmenting the active core (up to 16). As previously discussed with the usage of cuboidal TRISO fuel lattice, only a cubic TRISO lattice is compatible with axial discretization with 1-15 partitions. Discretization with more than 16 partitions is limited by how the data structures within the C++ script are initialized and is currently not supported. Line 42 gives the cold axial core height (in units of cm). Line 43 gives the cold height of the top and bottom axial reflector regions (in units of cm). The axial reflectors are geometrically identical to the active core region, except that there are no TRISO fuel or BP particles present. Line 44 gives the cold height of the top and bottom axial support plates (in units of cm). The axial support plates are homogenized for simplicity: the bottom is FLiBe and graphite homogenized at average inlet conditions, and the top is FLiBe and silicon carbide homogenized at average outlet conditions. Line 45 controls how interassembly thermal expansion is handled for the fuel assemblies (only relevant for axial discretizations with more than one partition). For option 0, axial-average thermal expansion is used in all axial partitions, resulting in a prismatic geometry. For option 1, axial partition-specific thermal expansion is used, which allows for increased spacing moving up through the core due to higher temperatures. This is meant to emulate axial flowering of fuel assemblies, but in implementation is more accurately described as axial staggering of assembly sections (see Figure 3.1).

7.1.2.g Beyond the Active Core Region

Lines 47-52 of the user input file change the cold dimensions of the radial components beyond the active core region (in units of cm). Progressing radially out from the core to the model boundary, the cold outer radius is provided for the following structures: permanent radial reflector, boron carbide layer, core barrel, downcomer, vessel liner, and reactor pressure vessel.

7.1.3 Providing Material Parameters

The third block of the user input file provides modeling options and material properties for components used in the 3D AHTR Serpent model. An example of this block taken from the user input file can be seen in Figure 7.3. The materials block is divided into four subsections: uniform definitions, densities, thermal expansion coefficients, and thermal conductivities. Each of these will be discussed individually.

```
53
54 --MATERIALS
55 --Uniform Definitions (For statepoint or desiring uniformity, no need for unic
56 0 // Fuel. 0 - Unique for each 1/3 assembly section , 1 - Or
57 0 // Other TRISO Particle Layers. 0 - Unique for each 1/3 as
58 0 // Structural Graphite Components. 0 - Unique for each 1/3
59 0 // Burnable Poison Material. 0 - Unique, 1 - Uniform
60 0 // Single Control Blade Material. 0 - Unique, 1 - Uniform
61 0 // Flibe. 0 - Unique, 1 - Uniform
62 --Densities
63 10.9 // Cold Fuel Density [g/cc]
64 1.0 // Cold Carbon Buffer Density [g/cc]
65 1.9 // Cold Inner Pyrolytic Carbon Density [g/cc]
66 3.1 // Cold Silicon Carbide Density [g/cc]
67 1.87 // Cold Outer Pyrolytic Carbon Density [g/cc]
68 1.75 // Cold Graphite Density [g/cc]
69 1.95 // Cold Carbon-Carbon Composite Density [g/cc]
70 1.25 // Cold Europia (Burnable Poison) Density [g/cc]
71 10.28 // Cold MHC (Control Blade) Density [g/cc]
72 2.37 // Cold Boron Carbide Density [g/cc]
73 8.93 // Cold Alloy N Density [g/cc]
74 7.92 // Cold Hastelloy 800 Density [g/cc]
75 --Thermal Expansion Coefficients (x 10^-6)
76 7.6 // Fuel
77 5.5 // Buffer
78 5.5 // Inner Pyrolytic Carbon
79 5 // Silicon Carbide
80 5.5 // Outer Pyrolytic Carbon
81 5 // Graphite
82 5 // Carbon-Carbon Composite
83 7.5 // Europia (BP)
84 4.8 // MHC (CB)
85 5 // Boron Carbide
86 13.6 // Alloy N (RPV Liner)
87 17.3 // Hastelloy 800 (RPV)
88 --Thermal Conductivities
89 3.7 // Fuel [W/(m*K)]
90 0.5 // Buffer Graphite [W/(m*K)]
91 4 // Inner Pyrolytic Carbon [W/(m*K)]
92 16 // Silicon Carbide [W/(m*K)]
93 4 // Outer Pyrolytic Carbon [W/(m*K)]
94 15 // Unirradiated Graphite [W/(m*K)]
95 15 // Fuel Stripe Matrix [W/(m*K)]
96
```

Figure 7.3. Input file portion pertaining to materials options and properties.

7.1.3.a Setting Uniform or Region-Specific Materials

Lines 56-61 of the user input file deal with uniform material options which are only intended to be used for model simplification for testing purposes; not full-resolution simulations. During the testing process, many models started with using a single, assembly-average definition for the various materials used in the active core. This sped-up the initial steps of running Serpent since less data needed to be processed and fewer geometric features initialized; especially when also combined with using a small number of axial partitions. These options have been retained for users, should they want to conduct similar tests where fully resolved details are not required nor desired, or to evaluate the impact of this simplification. Another application could be to zero-power studies, where core properties are mostly uniform. For each of the five lines, option 0 indicates using a one-third assembly section-specific definition for that material and option 1 indicates using a single uniform definition for that material. Line 56 is for using a single fuel material in the entire geometry. It is highly recommended to use unique fuel materials for depletion simulations, otherwise burnup will only be tracked at the core-average level. Line 57 deals with the other layers of the TRISO fuel particles beyond the fuel kernel (buffer, inner pyrolytic carbon, silicon carbide, outer pyrolytic carbon, and matrix graphite). Line 58 is for the structural and other graphite components of the fuel assemblies. Line 59 is for the europia BP spheres. As it was the case for the fuel, it is highly recommended to use unique BP materials for depletion simulations. Line 60 is for using a single material for all CBs inserted in the core. Line 61 is for core FLiBe. To reiterate, these uniformity options override many of the features and levels of detail which are central to this work. They were retained for users should they wish to quickly run for a particular reason but are not intended to generate detailed results.

7.1.3.b Material Densities

Lines 63-74 of the user input file provide the cold densities for the materials used in the AHTR model (in units of g/cm^3). In order of appearance, materials considered are: fuel kernel,

buffer, inner pyrolytic carbon, silicon carbide, outer pyrolytic carbon, graphite, carbon-carbon composite, europia (used as BP), MHC (used in CB), boron carbide, alloy N (INOR-8) vessel liner, and Incoloy Alloy 800H vessel material.

7.1.3.c Linear Thermal Expansion Coefficients

Lines 76-87 of the user input file provide the thermal expansion coefficients for the same materials listed above in the densities subsection. Values have a multiplication modifier of 10^{-6} , which is common practice for thermal expansion coefficients. While thermal expansion coefficients are typically temperature-dependent physical properties, the script only uses a single constant value for all temperatures (hence linear behavior with respect to temperature changes). It is recommended to select a value which closely represents the parameter behavior for each material over its expected temperature range.

7.1.3.d Thermal Conductivities

Lines 89-95 of the user input file provide the thermal conductivities for components of the fuel plank (in units of $W/(m \cdot K)$). In order of appearance, materials specified include: fuel, buffer, inner pyrolytic carbon, silicon carbide, outer pyrolytic carbon, graphite, and fuel stripe matrix (which is likely graphite, but listed separately in case manufacturing techniques impact the thermal performance of this material specifically). FLiBe thermal conductivity will be addressed later with other FLiBe properties. Note that the thermal conductivities for other components of the core are not required. This is due to how the heat transfer is modeled: heat is produced in the TRISO fuel kernels and conducts through the fuel plank to the coolant boundary and is then transferred to the coolant via convection. Thermal conductivity is not needed for carbon-carbon composite because it is not considered as part of the heat transfer process. Other materials with property specifications listed in the previous sections are also not required here if they appear outside the core, since features radially beyond the permanent reflect region are assumed to be at the inlet coolant

temperature. Since at a single temperature, there is no heat transfer modeled and thus a thermal conductivity is not needed for the model.

7.1.4 Providing Thermal Hydraulic Parameters

The fourth block of the user input file deals with TH parameters. An example of this block taken from the user input file can be seen in Figure 7.4. The TH block is divided into two subsections: core power with flow properties and FLiBe physical properties. Each of these will be discussed individually.

```
96
97 -THERMAL HYDRAULIC PARAMETERS
98 --Core Power and Flow Properties
99 1 // Number of Thermal Hydraulic Iterations
100 1.953376E-01 // Core Average Power Density [kW/g]
101 26750 // Mass Flow Rate through Core [kg/s]
102 293 // Cold Component Reference Temperature [K]
103 923 // Core Inlet Temperature [K]
104 1 // Print TH Profile for Highest Power Zone. 0 - No, 1 - Yes
105 1 // Fuel Stripe Temperature Profile. 0 - Homogenized, 1 - Par
106 --Flibe Properties
107 0.0056 // Viscosity [Pa*s]
108 13.525 // Prandtl Number
109 2415 // Heat Capacity [J/(kg*K)]
110 1 // Thermal Conductivity [W/(m*K)]
111
```

Figure 7.4. Input file portion pertaining to thermal hydraulic options and properties.

7.1.4.a Core Power and Flow Properties

Lines 99-105 of the user input file set various core and flow properties. Line 99 gives the number of TH iterations. These are executed after criticality iterations (if any). Line 100 gives the average core power density (in units of kW/gHM). Volume normalization in the model is not straightforward due to the number of features inherently part of ATOMICS as well as the inherent geometric complexity of AHTR, so it is much easier for Serpent to deplete materials with the correct normalization when given power density versus total core power. In the reference AHTR design, the total core thermal power is 3400 MW. For a system with 4 x 202 x 5923 TRISO particles in each fuel stripe (as is the case with the models used in this work), this translates to a power

density of about 0.195 kW/gHM (unit used by Serpent for power density). Line 101 gives the total mass flow rate of coolant through the core (in units of kg/s). Line 102 provides the cold reference temperature for components (in units of K). This is the temperature used for thermal expansions where relevant in the code. Line 103 gives the core inlet temperature for coolant (in units of K). Note that for TH, there needs to be one free parameter among the grouping of core inlet temperature, average core outlet temperature, total core thermal power, and total mass flow rate. In this implementation it is assumed that the core inlet, total mass flow rate, and total core thermal power (via average power density) are known. This leaves the core outlet temperature as being free and dictated by the other three parameters. For the reference input parameters, the average core outlet is about 700 °C (973 K) with an average temperature change across the core of about 50 °C. Line 104 controls whether an output file is printed for the TH profile of the section in the core with the highest power production. The file is named “THChannelProfile.txt” and lists summary data such as where the peak power section is located (axial partition number, assembly group number, and assembly section number), PPF, and the peak temperature. Then, there is an array with a fixed number of 10,000 elements for locations and corresponding temperatures from the center of the coolant channel to the center of the fuel plank. Line 105 controls the usage of the homogeneous (Option 0) versus heterogeneous-reconstructed temperature profile (Option 1) for the TRISO fuel stripe. Both instances use a homogenization technique to reduce the heat transfer problem in the coolant channel to 1D. The particle reconstruction option regains the temperature profile through the centerline of the TRISO particle lattice.

7.1.4.b FLiBe Properties

Lines 97-100 of the user input file change properties of the coolant FLiBe used in the model. Line 97 dictates the viscosity of FLiBe (in units of Pa·s). Line 98 gives the Prandtl number of the coolant flow. Line 99 gives the heat capacity of FLiBe (in units of J/(kg·K)). Line 100 gives the thermal conductivity of FLiBe (in units of W/(m·K)). Note that as with the discussion on thermal

expansion coefficients, these properties of FLiBe have temperature dependence, but only a single value is required by the script. The user should select a value which represents well the parameter over the expected temperature range of the simulation. The default values work well for the operating range of AHTR (650-700 °C), but perhaps should be changed for studies such as cold zero power cases and beyond operating basis scenarios since they change with temperature.

7.1.5 Depletion Simulation Options

The fifth block of the user input file deals with depletion simulation options. An example of this block taken from the user input file can be seen in Figure 7.5.

```
111
112 -DEPLETION SIMULATION OPTIONS
113 0 // Use Fluence-Dependent Thermal Conductivity and Thermal Expansi
114 0 // Use Equilibrium Xenon Treatment for Fuel. 0 - False, 1 - True
115 1 // Number of BP Burnable Zones. (Serpent supports up to 10)
116 1 // Target eigenvalue for Control Blade Movement
117 150 // Eigenvalue tolerance for Control Blade movement. [pcm]
118 0 // Initial Guess for Number of CB Groups to Insert
119 11000 // Depletion Particles per Cycle
120 20 // Depletion Active Cycles
121 20 // Depletion Inactive Cycles
122
```

Figure 7.5. Input file portion pertaining to depletion options and criticality convergence.

Line 113 determines whether fluence-dependent properties are used for graphite in the depletion simulation. As previously discussed in Section 4.5, a fluence and temperature functional response was previously obtained for graphite properties. Option 0 ignores this response (invariant with fluence) while Option 1 uses this response. Line 114 implements an equilibrium xenon treatment for all fuel materials. Within Serpent, the equilibrium xenon feature assumes some steady flux-determined concentration of xenon-135 in all fissile burnable materials. This helps alleviate the impact of numerically-induced spatial xenon oscillations during depletion simulations driven by statistical uncertainties of the transport simulation [45]. This is especially relevant in a large thermal spectrum reactor like AHTR where the physical size of the core is significantly larger than the mean free path, so regions of the core have poor neutronic communication with each other.

Although equilibrium xenon treatment reduces the potential impact of numerical instabilities, it also forces the equilibrium concentration from the outset. This means that there is no period for the xenon to build-in to its equilibrium concentration after startup. Therefore initial (first short step of about one day) depletion results typically have significant eigenvalue differences versus cases not using the equilibrium xenon feature, but these quickly disappear as the equilibrium concentration would normally build-in. Line 115 sets the number of BP burnable zones to use in the simulation. As shown in Section 6.X, this level of tracking has a non-trivial impact on the reactivity results of depletion simulations (see Figure 6.7). Serpent can generate up to ten zones automatically, which is the upper limit for this parameter. Line 116 sets the target eigenvalue for control blade movement. By default, this is set to one but there are instances where the user would want to adjust it. For example, the goal of a statepoint calculation may be to find a CB insertion configuration that satisfies a criticality safety condition with the multiplication factor below some prescribed subcritical value. Another usage is that during a depletion sequence, only one geometry configuration can be used for each depletion step. Instead of using a critical initial step geometry and depleting to subcritical isotopic concentrations (with average subcritical configuration), one could use a slightly supercritical initial eigenvalue for each step which would result in using a configuration averaging closer to critical over the burnup step. This requires a fairly informed knowledge basis of how a given core loading will burn for the depletion steps considered and a level of linearity in the results to justify the chosen value. Line 117 gives the eigenvalue tolerance for the control blade movement (in units of pcm). Note that if this tolerance is set too tight with respect to statistical uncertainty, unphysical oscillations with no convergence are possible. As previously discussed, CB movement is resolved to the one-sixteenth axial level. This discrete insertion methodology is not continuous, so only corresponding discrete eigenvalues are achievable via CB movement. If the average worth of a CB group is assumed to be about 320 pcm, the eigenvalue tolerance should not be given below 20 pcm without significant caution. More realistically, a reasonable lower bound on eigenvalue tolerance would be about 50 pcm due to

statistical uncertainty inherent with using MC and the fact that each axial segment worth will differ (sometimes significantly) from the average assembly worth. For example, axial segments near the midplane in high power bearing assemblies will have reactivity worths well above the average and a critical insertion might not be possible given an average section worth tolerance. A tolerance bound of 50 pcm for a one-sixteenth axial section also agrees well with initial test seen in Figure 6.3, where one-eighth axial sections had an observed reactivity worth of about 100 pcm. Larger tolerances should also converge faster (since fewer CB movements are required) and may allow for more computational resources to be committed to the burnup transport calculation instead of the critical geometry search process. Line 118 sets the initial guess for the number of CB groups inserted in the core. This follows the order of CB insertions in the “CRSchedule.txt” file provided by the user. This initial guess accelerates the iteration process of searching for a critical insertion by obtaining geometries near criticality where extrapolating to the desired eigenvalue should require fewer iterations than if starting from a point where all CB groups are either inserted or withdrawn. Lines 119-121 deal with setting the neutron histories used during depletion simulations. Line 119 gives the number of particles per generation. Line 120 gives the number of active cycles. Line 121 gives the number of inactive cycles. As previously discussed, these values are only used as part of transport simulations for depletion steps. They are distinct from the neutron history parameters used in statepoint transport simulations.

7.2 Implementation of the Methodology in ATOMICS

This section gives a brief overview of how ATOMICS integrates various components of the methodology. The general flow is the reading in of input files, executing code operations, which is where model parameters are updated based upon user specifications and results from previous Serpent simulations, and writing results to output files. File input/output streams are critical to the implementation of the code because it is not running or idle when Serpent transport cases are being simulated. After ATOMICS is run, it will need to be called again after the next Serpent transport

calculation is completed. Relevant parameters between simulations are recorded in the output files, read in again when executed on the next instance, and used during the code execution. A visualization of the general code flow pattern can be seen in Figure 7.6.

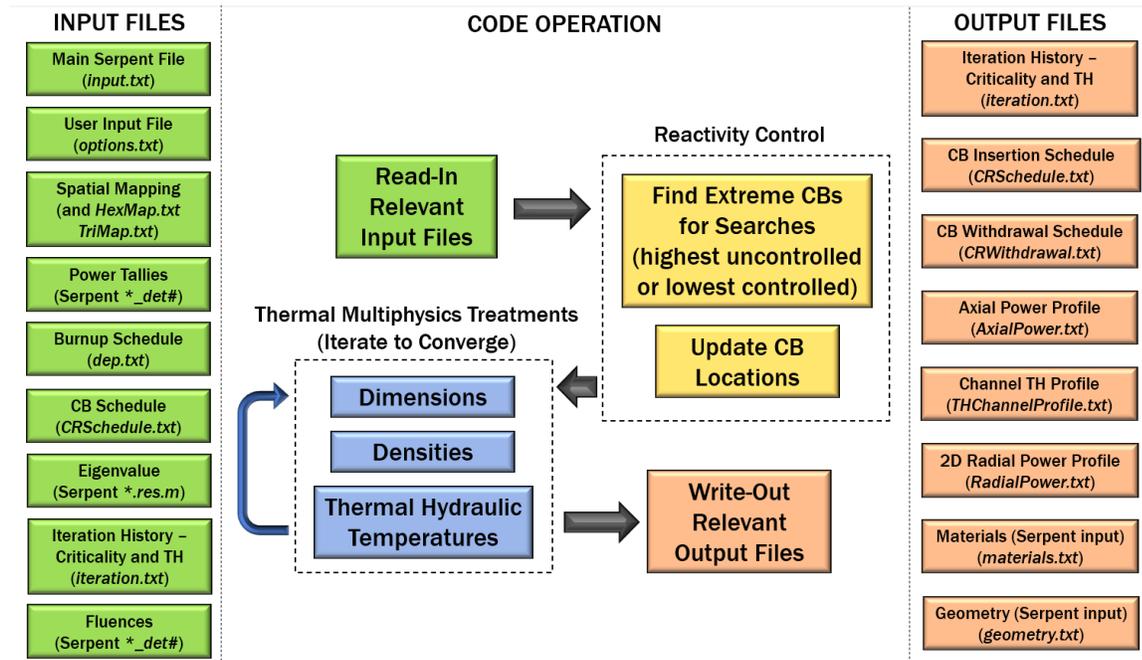


Figure 7.6. A visual depiction of how ATOMICS works and a summary of important I/O files.

7.2.1 Reading of Input Files

Input files provide key parameters that guide the script flow and selection of alternative functional flowpaths. Since the code does not pause or idle during Serpent transport simulations, they are vital to keep the workflow progressing through iteration and depletion steps. A short summary for each of the input files listed in Figure 7.6 will be addressed below. Not all files in this list are used for every run instance. Depending upon the options selected in the user input file, certain files are used while others are not.

7.2.1.a Main Serpent File

The main input file which Serpent should be run with is *input.txt*. This file is created by the user and includes a limited number of parameters, as most modeling features are included in

either the *materials.txt* or *geometry.txt* files. An example *input.txt* file is shown in Figure 7.7. The parameters essential for running are to tell Serpent to include the geometry and materials files as part of the input (lines 3 and 5), to set the vacuum boundary condition at the geometry boundary (line 8), to include the paths to data libraries to be used in the simulation (lines 12, 14, and 15), and to set additional options for the simulation (lines 17 and 20-23). The additional options shown are related to how cross section data is initialized and stored for use during transport and various depletion-specific settings.

```
1 set title "AHTR 3D Depletion"
2
3 include "geometry.txt"
4
5 include "materials.txt"
6
7 % Boundary Conditions
8 set bc 1 % 1 - Vacuum 2 - Mirror 3 - Periodic
9
10 %phoenix - ENDF 7.0
11 % --- Cross section data library file path:
12 set acelib "/absolute/path/to/data/file.xsdata"
13 % --- Decay and fission yield libraries:
14 set declib "/absolute/path/to/data/file.dec"
15 set nfylib "/absolute/path/to/data/file.nfy"
16
17 set opti 2
18
19 % --- Options for burnup calculation:
20 set bumode 2 % CRAM method
21 set pcc 0 % Predictor-corrector calculation on
22 set xscal 2 % Cross sections from spectrum
23 set ures 0 % (1-Use 0-Omit) unresolved resonance table
24
```

Figure 7.7. Example *input.txt* file.

7.2.1.b User Input File

As already discussed in detail at the beginning portion of this chapter, the user input file (*options.txt*) is the primary means for the user to direct the code on how to function. A complete example user input file can be found in Appendix E. The user selects modeling options as well as prescribes physical dimensions and properties to use when constructing the model geometry. The code does not modify contents of this file under any option combination, but it does echo the read values back to the user. This is done for two reasons: first, for the user to immediately identify if

some desired parameters were potentially entered incorrectly with numerous (but not completely exhaustive) error checks; and second, so that parameters for that particular execution can be saved for future reference if the user is also piping the ATOMICS terminal output to an output file. This second point is important due to the nature of having independent options files. Without this terminal echo, there would not be an easy way to trace back the input parameters used for a particular code execution. This ensures traceability of analyses performed and facilitates case comparisons.

7.2.1.c Spatial Mapping Files

Two supporting files which are used to facilitate correct spatial indexing of results are the files *TriMap.txt* and *HexMap.txt*. *TriMap.txt* is used to map the results from the triangular mesh tally (one-dimensional array) to the naming convention used within ATOMICS (84 assemblies with 3 assembly sections per axial partition). The mapping used is identical to the last three columns of Supplementary Table C.1 (page 260). *HexMap.txt* is used to map from the naming convention space used by ATOMICS to a 19x19 array used to show the 2D radial power profile in the output file *RadialPower.txt*. The mapping used in this work can be found in Appendix section F.1 (page 301).

7.2.1.d Power Tallies

This is the Serpent detector output file (**.det*). Notably, this includes the results from using the triangular superimposed mesh tally. The file provides fission power data to be used by the TH portion of the code to update temperatures in the core based on local power production.

7.2.1.e Burnup Schedule

The burnup schedule user input file *dep.txt* contains the depletion steps which will be used by ATOMICS to run a depletion sequence. The depletion step types supported correspond with the same ones offered within Serpent, as can be seen in Table 7.1. An example input file for *dep.txt* can be found in Figure 7.8.

Table 7.1. Burnup step types used within Serpent [46].

<stype>	Step values
bustep	depletion step, burnup intervals given in MWd/kgU
butot	depletion step, cumulative burnup given in MWd/kgU
daystep	depletion step, time intervals given in days
daytot	depletion step, cumulative time given days
decstep	decay step, time intervals given in days
dectot	decay step, cumulative time given in days

The first line of *dep.txt* corresponds with burnup step type used for the depletion steps. The options align with the listings as presented in Table 7.1: Option 1 corresponds with “bustep” basis; Option 2 corresponds with “butot” basis, Option 3 corresponds with “daystep” basis, Option 4 corresponds with “daytot” basis, Option 5 corresponds with “decstep” basis, and Option 6 corresponds with “dectot” basis. In Figure 7.8, the first line has value 2, so the depletion step basis is “butot”. Line 2 is the total number of depletion steps contained in the file. In Figure 7.8, the value 17 corresponds to the 17 depletion steps which follow it. The ensuing lines list the depletion steps in the basis established by the first line. In Figure 7.8, since the basis is “butot”, the cumulative burnup values of the depletion sequence are given in units of MWd/kgU.

```

1 2
2 17
3 0.1
4 0.5
5 1
6 2
7 4
8 6
9 8
10 10
11 14
12 18
13 22
14 26
15 30
16 40
17 50
18 60
19 70
20

```

Figure 7.8. Example *dep.txt* input file.

7.2.1.f CB Schedule

The control blade schedule file prescribes to the code in which order to withdraw CB groups from the core during a depletion sequence. This file is provided by the user and can either be manually created by the user or generated in an automated search fashion with the capabilities detailed in Chapter 6. The general file construction is that for N ordered assembly groups, $N+1$ values are required. The first line gives the number of assembly groups contained in the file (for a full-core ordering, this value will be 84). The subsequent lines (Lines 2 to $N+1$) provide the order in which CB groups should be inserted. Line 2 is the first group to be inserted and the last to be withdrawn. Line $N+1$ is the last group to be inserted and the first to be withdrawn.

7.2.1.g Eigenvalue

The eigenvalue from the previous Serpent transport simulation is obtained from the results file automatically generated with every Serpent run (the **.res.m* file). There are many other useful parameters contained in this file, but no others are leveraged at this time. As part of future work which could extend beyond the efforts detailed in this dissertation, using additional results from the **.res.m* file would be a good starting point for adding new features.

7.2.1.h Iteration History

When CB movement or TH iteration is being utilized during a statepoint or depletion simulation, the file *iteration.txt* logs the iteration and criticality results from each step of the iterative process. Critical insertion iteration steps are recorded with eigenvalues and critical insertion positions to help understand the behavior of the system and extrapolate what the next test case should be. TH iteration steps are then executed as dictated by the user input file *options.txt*. This file will be discussed more in section 7.2.3.a.

7.2.1.i Neutron Fluence

The thermal properties of graphite change with neutron fluence, which can be significant in a system, which contains as much graphite as AHTR. There are a few potential ways (both directly and indirectly) one could obtain the fluence in a region of the core. The method used in this work is directly with a flux tally and calculating cumulative fluence results by summing over all steps for the duration of each burn step to account for fluence. This functionality is not highly tested and should be considered as experimental. Future work could look into tuning this feature and adequately testing it.

7.2.2 Code Execution

The C++ script can be broadly divided into four main functional routine types: reading files, reactivity control, TH, and writing files. Reading files is rather self-explanatory- based on the input options selected by the user, relevant files are read by the script to extract further data to be used by the code. Reactivity control depends on the script mode selected by the user. When searching for a CB schedule (either withdrawal or insertion), the script uses the power tally data to identify the extreme power contributing assemblies. Then, based on the result, CBs are either withdrawn from the lowest power position or inserted into the highest power position (depending on the script mode). When criticality search is selected, excess reactivity is controlled by following the CB schedule provided by the user. CBs are withdrawn to achieve a critical geometry, within the eigenvalue tolerance also supplied by the user. TH uses the power tally results to iterate through resulting temperatures of components, updating dimensions due to thermal expansion with the new temperatures, and density changes also from the thermal expansion. Once these values are self-consistent, the script is finished running and can move on to writing output files. These files are either called when the script is run again (essentially storing/caching data) or are used by Serpent as input files for the next transport simulation. New Serpent simulations are launched for each instance of ATOMICS (both substeps and burnup steps).

7.2.3 Writing of Output Files

Important information from running ATOMICS is written to output files. These can broadly be classified as ATOMICS-specific, Serpent-specific, or summary-specific. ATOMICS-specific outputs help it know what was previously simulated to make data-driven decisions such as the iteration number of the previous run, estimating the worth of CBs to guess what the next CB insertion configurations should be, etc. These files do not follow any special format styling since they are used exclusively for ATOMICS. Serpent-specific outputs are used as part of the input for the next iteration of transport simulations. These mostly deal with updates to the model due to multiphysics effects. They are written for Serpent-specific syntax. Summary-specific files are outputs which are neither used by ATOMICS in subsequent iterations nor used by Serpent for the next transport sequence. They are only useful to the user as a means of summarizing results from the previous Serpent simulation. Each of the output files shown in Figure 7.6 will be briefly discussed below.

7.2.3.a Iteration History

During a CB critical insertion search process or TH iteration, the results from each iterative step are written to *iteration.txt*. The results from the most previous step are normally used to help guess what CB insertion configuration should be tested next. From the standpoint of the user, the file can also be used to see how the search process performed over cycle. Since it records the CB insertion configuration before a depletion step is executed, the user can use this file to reconstruct the simulation if they want to conduct any branching studies from the CB configuration used. An example of an *iteration.txt* file can be seen in Figure 7.9.

```

1 00 00 00 00 00 0
2 00 01 00 57 00
3
4 00 01 00 57 00 1.01096
5 00 02 00 56 00
6
7 00 02 00 56 00 1.00367
8 00 03 00 55 08
9
10 00 03 00 55 08 1.00599
11 00 04 00 56 13
12
13 00 04 00 56 13 1.00016
14 00 04 01 56 13
15
16 00 04 01 56 13 1.00015
17 00 04 02 56 13
18
19 00 04 02 56 13 0.999796
20 01 01 00 56 13
21
22 01 01 00 56 13 0.978094
23 01 02 00 55 13

```

Figure 7.9. Example *iteration.txt* file.

The results written to the *iteration.txt* file are appended to the existing file. This means that prior results are unaltered and only the most recent iteration values are added to the end of the file. An output for an iteration summary consists of eleven values over two lines. The first line is related to the previous transport simulation. It contains (from left to right) the depletion step number, the criticality iteration number, the TH iteration number, the number of CB groups fully inserted, the fraction of partial insertion for a single CB group, and the eigenvalue. The second line is related to the next transport simulation to be run by Serpent. It contains the same value as the line above except for the eigenvalue, which is unknown since the transport simulation has not been run yet. The values (from left to right) are: the depletion step number, the criticality iteration number, the TH iteration number, the number CB groups fully inserted, and the fraction of partial insertion for a single CB group.

The fraction of partial insertion depends on the granularity of axial partitioning. For the example shown in Figure 7.9, sixteen axial partitions are used in the active core region. In Line 19, before the first depletion step is simulated, after four criticality iterations, and then following two

TH iterations; the system has an eigenvalue of 0.999796 for 56 CB groups fully inserted and 13/16th of a 57th CB group inserted. For the sample results shown in Figure 7.9, the initial guess for the number of assembly groups needed to maintain criticality is 57 (as evident from Line 2). The eigenvalue tolerance used is 200 pcm within an eigenvalue of 1, which is why criticality iteration continues four times (Line 13). After the CB position is accepted within the eigenvalue tolerance, TH iterations occur. In this case, two TH iterations were prescribed by the user, after which depletion starts (Line 20). The process continues for each depletion step listed in the depletion schedule user input file *dep.txt*.

7.2.3.b CB Insertion Schedule

When using the CB insertion schedule search mode (Option 1 for Line 4 in the user input file *options.txt*), this is the only time that the file *CRSchedule.txt* is altered by ATOMICS. Normally the file is untouched since it is read as an input for all other modes. When conducting a CB insertions schedule search, ATOMICS reads the power tally results from the previous Serpent simulation to determine the assembly group with the highest power which does not already have CBs inserted. CBs are inserted into this group, the file *CRSchedule.txt* is updated to reflect this by adding that assembly group to the end of the file, and the next Serpent transport simulation is run to start the search process over again.

7.2.3.c CB Withdrawal Schedule

The file *CRWithdrawal.txt* is only used when using the CB withdrawal schedule search option (Option 2 for Line 4 in the user input file *options.txt*). The order of the CB withdrawals is recorded here. Once the schedule search process is complete, the ordered file could be used as a CB schedule file for subsequent simulations. The intended functionality is that the user provides an arbitrarily ordered *CRSchedule.txt* file where all specified CB groups are initially inserted. After each Serpent simulation, the power results are used to find the inserted CB group with the lowest

power contribution. This assembly group is withdrawn as part of the search procedure and the process repeats until all CB groups have been withdrawn.

7.2.3.d Axial Power Profile

Following each instance of ATOMICS, an axial power profile is printed to an output file *AxialPower.txt*. Results are obtained from a superimposed power tally in the detector results file and normalized within a fixed number of 112 axial bins covering the active core region. 112 was selected since it is divisible by 16 ($7 \cdot 16 = 112$), which corresponds to the maximum number of supported active core axial partitions in the model. Integral axial parameters are also included at the end of the file summarizing the axial PPF, AO, and AO statistical uncertainty. Results are binned starting from the bottom/inlet of the core (index 1) to the top/outlet of the core (index 112). A slight caution is that if the number of active core partitions does not evenly divide 112, some axial bins will score across two neighboring axially partitioned sections. Since the tally is a superimposed mesh this does not pose an issue from a simulation standpoint, but some the tally bin results will not be attributable to a single axial partition. This is not considered a concern but the user should be aware of the procedure.

7.2.3.e Hot Channel Thermal Hydraulic Profile

If the user so requests (Option 1 of Line 104 in the user input file *options.txt*), a temperature profile (*THChannelProfile.txt*) is printed for the section in the core which has the highest local power production based on the Serpent power tally results of the previous simulation. The profile is for the temperature distribution within the cooling channel and fuel plate used for the TH model. The printed profile can be for the homogenized fuel stripe (Option 0) or heterogenous TRISO particle reconstruction (Option 1), based on the entry in the user input file *options.txt* (Line 105). The top of the file contains summary information related to the location of the highest power production (axial core partition, assembly group, assembly section), PPF, and peak temperature. Following this summary is an array with 10,000 elements with the location (uniformly spaced) and

corresponding temperature from the center of the cooling channel to the center of the fuel plank. Granularity is fine enough to capture the rapidly changing gradients in the fuel stripe region when resolving the heterogenous TRISO particle profile.

7.2.3.f 2D Radial Power Profile

A 2D radial power profile (*RadialPower.txt*) is output following each instance of ATOMICS. It uses the Serpent power tally results from the prior simulation to produce an axially integrated normalized radial power profile in a 19 x 19 array representing hexagonal core arrangement. The array is made possible by changing from ATOMICS' phase space of assembly indexing to *x*- and *y*-indexing using the mapping in the input file *HexMap.txt*. Following the 19 x 19 array is summary information, including: peak assembly group and corresponding radial PPF, peak assembly section group and corresponding radial PPF, and assembly group statistical uncertainty values (maximum, minimum, and average).

7.2.3.g Materials

The file *materials.txt* is written in Serpent-specific input format for all the material definitions used within the model. It is automatically added to the next Serpent run with an append-like feature called "include" in the main Serpent input file *input.txt*. In addition to the materials definitions, it also includes physics parameters including: power density, neutron history parameters (particles per generation, active cycles, and inactive cycles), and depletion commands (restart command from the previous burnup step, write command for the next burnup step, and transport command for the next burnup step). The length of the materials file varies depending on how many materials are uniquely defined based on parameters in the user input file. For a single axial partition and uniform materials in all sections (Option 1 for Lines 56-61 in the user input file *options.txt*), the materials file is only about 170 lines long. For sixteen axial partitions and all materials resolved uniquely at the axial assembly section level, the materials file is about 29,000 lines long.

7.2.3.h Geometry

The file *geometry.txt* is written in Serpent-specific input format for the model geometry. It is automatically added to the next Serpent run with an append-like feature called “include” in Serpent. In addition to the geometry definitions, the file also includes tallying parameters. The tallies are superimposed meshes, which are geometry based, so require parameters obtained from this section of the code (active core lower and upper boundaries, assembly pitch, etc.). This is the largest file of all used by ATOMICS due to how complex the AHTR geometry is and the redundant nature inherent in defining unique regions. For a single axial region, the geometry file is about 100,000 lines long. For a full feature run with sixteen axial partitions for the active core, the geometry file is about 1,100,000 lines long. A visualization of the order in which ATOMICS is executed can be found in Figure 7.10.

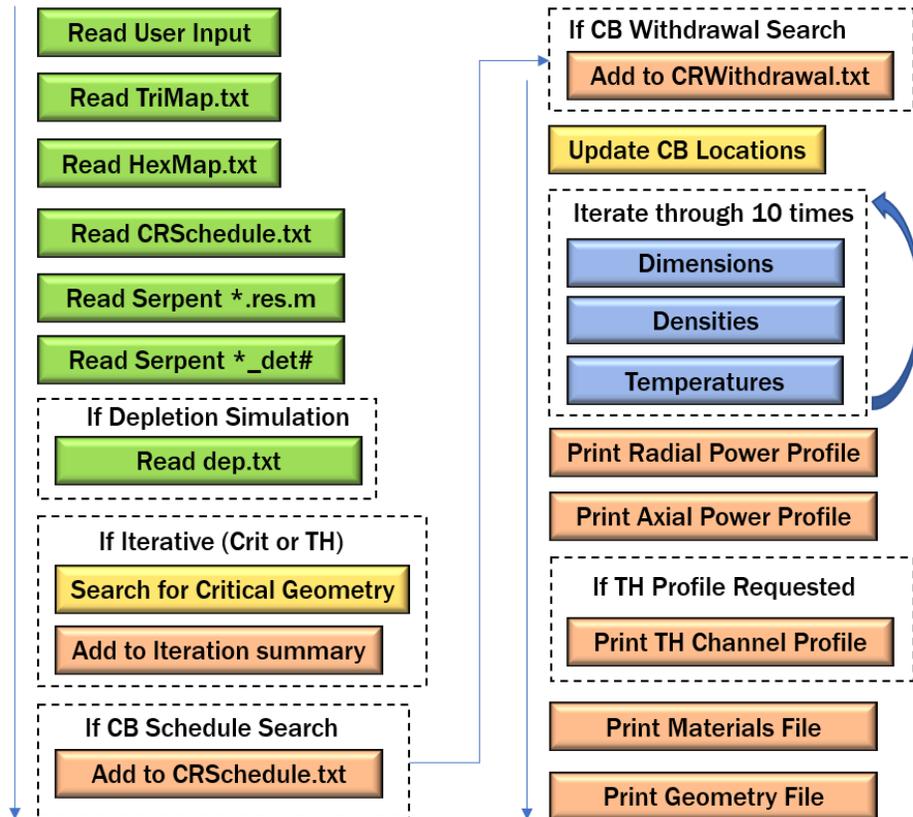


Figure 7.10. Order in which ATOMICS subroutines are executed.

7.3 Using ATOMICS to Run Simulations

This subsection provides a brief overview of how to get started with running ATOMICS in a Linux environment. The steps include updating Serpent 2.31 to include the triangular superimposed mesh tally capabilities, compiling the C++ source code for ATOMICS, and finally running ATOMICS in a Linux environment.

7.3.1 Updating Installation of Serpent

To have the power tallying capabilities necessary for ATOMICS to run correctly, Serpent 2.31 must be recompiled with the updated input files listed in Table 5.2. This process is outlined below:

- In a directory where Serpent 2.31 is already installed, replace the eighteen files with the new versions
- Execute the command “make clean” to remove the previous installation of Serpent 2.31
- Compile again with the command “make”

After a few minutes, an executable file should be produced named “sss2”. This is the resulting executable version of Serpent 2.31 updated with triangular superimposed meshing tallying capabilities. The executable can be renamed so that it is distinct from other Serpent installations and not confused with other existing versions. For simplicity, this executable will be called “sss231t” for the remainder of the section. The executable file sss231t can be moved to another directory, or commands to run Serpent can use an absolute path to this directory or any other directory containing the executable.

7.3.2 Compiling ATOMICS

Compiling ATOMICS in a Linux environment is equally straight-forward. It should be compiled locally to where simulations will be run, since input and output files are essential to the code’s functionality. In Figure 7.11, one can see that the source code file *ATOMICS.cpp* is contained within the directory called “example”. An outline of how to compile ATOMICS is given below:

- Upload the ATOMICS C++ source code to the desired directory where ATOMICS will be installed.
- ATOMICS can be compiled using g++ (GNU C++ compiler) or any other C++ compiler, as the source code is written in standard C++. In Figure 7.11, g++ is used to first compile the C++ file into object code with the command “g++ -c ATOMICS.cpp”.
- Next, the resulting object code file *ATOMICS.o* is used to create the executable file which is simply named *ATOMICS* using the command “g++ -o ATOMICS ATOMICS.o”.

Listing the contents of the directory again shows that indeed the object file *ATOMICS.o* and the executable file *ATOMICS* were successfully created.

```
[kramey7@login-phoenix-3 example]$ ls
ATOMICS.cpp
[kramey7@login-phoenix-3 example]$ g++ -c ATOMICS.cpp
[kramey7@login-phoenix-3 example]$ g++ -o ATOMICS ATOMICS.o
[kramey7@login-phoenix-3 example]$ ls
ATOMICS ATOMICS.cpp ATOMICS.o
```

Figure 7.11. Example of how to compile ATOMICS in a Linux environment.

7.3.3 Running ATOMICS in Linux Environment

ATOMICS can be thought of as a “wrapper” code for Serpent; it runs externally from Serpent and they communicate to each other via shared input and output files. ATOMICS also does not continue running or go idle while Serpent transport simulations are occurring. It must be called again for each iteration when used with Serpent (criticality iterations, TH iterations, depletion steps, CB insertion search steps, CB withdrawal search steps, etc.). This is easy to accomplish with bash shell loops, like the example seen in Figure 7.12.

```

for i in {1..99}
do
  ./ATOMICS
  cp RadialPower.txt RP$i.txt
  cp ST_det0.m ST_$i.txt
  cp THChannelProfile.txt THCP$i.txt
  cp AxialPower.txt AP$i.txt
  sss231t input.txt -omp 24 | tee Output.txt
  cp input.txt_det0.m TR$i.txt
  cp Output.txt O$i.txt
done

```

Figure 7.12. Example bash shell for-loop used for running ATOMICS

For each iteration, ATOMICS is called to create *materials.txt* and *geometry.txt* for Serpent to run. In addition, several output files are produced which are given non-unique names and are overridden between iterations. If these files are desired by the user to be kept for reference to analyze later, they should be copied to another file name so that they can be preserved. In Figure 7.12, after each instance of ATOMICS, some ATOMICS-related output files are copied to uniquely named files. After each Serpent transport simulation, the tally results and terminal output are copied to unique files as well.

CHAPTER 8

RESULTS

This chapter details results obtained from running several different studies using ATOMICS. The first section focuses on geometric changes made to individual parameters controlled by the user input file *options.txt*. The second section focuses on thermal impacts of changing inlet temperature and core power. These can both be classified as model sensitivity studies. The third section focuses on depletion simulations run with various combinations of options. As this is where the multiphysics components come together, this section highlights the full integration of the work as a whole.

Many studies in this chapter include estimates for the AO and various PPFs resolved at the assembly, one-third assembly section, axial, and whole core (3D) level. Statistical uncertainties for each of those parameters are in most cases not reported along with their corresponding results in the interest of conserving space. However, they are roughly proportional to the relative statistical uncertainty of the eigenvalue. Table 8.1 presents the reported relative statistical uncertainties for various power shaping parameters and how they compare to the eigenvalue relative statistical uncertainty. Note that actual statistical uncertainties may be higher than the reported statistical uncertainties, sometimes up to an order of magnitude; furthermore, the uncertainties may amplify with depletion. This is especially relevant for differential quantities, i.e. spatial distributions and factors where particles sampling in the core is correlated across neutron population generations.

Table 8.1. Estimates of reported statistical uncertainty for various power shaping parameters for an example simulation with an eigenvalue relative statistical uncertainty of 10 pcm.

Power Shape Parameter	Spatial Resolution (Fraction of Core)	Reported Relative Statistical Uncertainty [pcm]	Factor of Eigenvalue Reported Statistical Uncertainty
AO	1/2	20	2
Axial PPF	1/112	106	10.6
Assembly PPF	1/84	92	9.2
Section PPF	1/252	159	15.9
Whole Core PPF	1/4032	635	63.5

8.1 Geometric Sensitivity Studies

This section summarizes the results from varying aspects of the AHTR geometry via the user input file *options.txt*. Reproducing these results and similar datasets should be simple due to the one-line value modifications generally used to generate each set. Five parameters are considered here: fuel plank width, the distance from the fuel stripe to the coolant boundary (thin layer of graphite separating the two, also referred to as the plank “sleeve”), axial reflector height, number of burnable poison columns, and burnable poison column pitch. These were selected for further study because if changes were to be made to the AHTR geometry in the future, this set would be a likely starting point for design optimization.

8.1.1 Fuel Plank Width

Each of the fuel plates of an AHTR fuel assembly with its cooling channels on either side has a combined cold width of 3.25 cm. For the reference dimensions, the fuel plank is 2.55 cm wide with two 0.35 cm cooling channels (if the large 0.7 cm channel between two fuel planks can be considered simply as divided in half). This study focuses on the eigenvalue impact due to varying the fuel plank width. Note that since the cold AHTR fuel assembly size is held constant, the cooling channel width also varies as part of this study. This includes the size of the spacer structures since they span the distance across the flow channels. Five plank widths are considered as part of this study: 1.05 cm, 1.55 cm, 2.05 cm, 2.55 cm (the reference dimension), and 3.05 cm. A visualization of the geometric changes can be found in Figure 8.1. Due to the significant variance in the channel width, the size of the spacers also changes from being very large to quite small. Additionally, no further considerations beyond the eigenvalue impact were made regarding these geometric changes. Some potential concerns stemming from these changes include: feasibility of manufacturing any of the given configurations; thin member issues for small fuel planks (buckling, induced vibrations due to flow and potential for wearing, etc.); and adequate cooling capabilities for large fuel planks. These are beyond the scope of this study; the main purpose was to demonstrate the ease by which

complex fuel assembly design changes may be analyzed by simply changing a single parameter in the user input file *options.txt* and using ATOMICS. A summary of the simulations of the five cases of interest can be found in Table 8.2. Each case was simulated without the use of BP spheres, all CBs withdrawn, and otherwise reference dimensions for other assembly components.

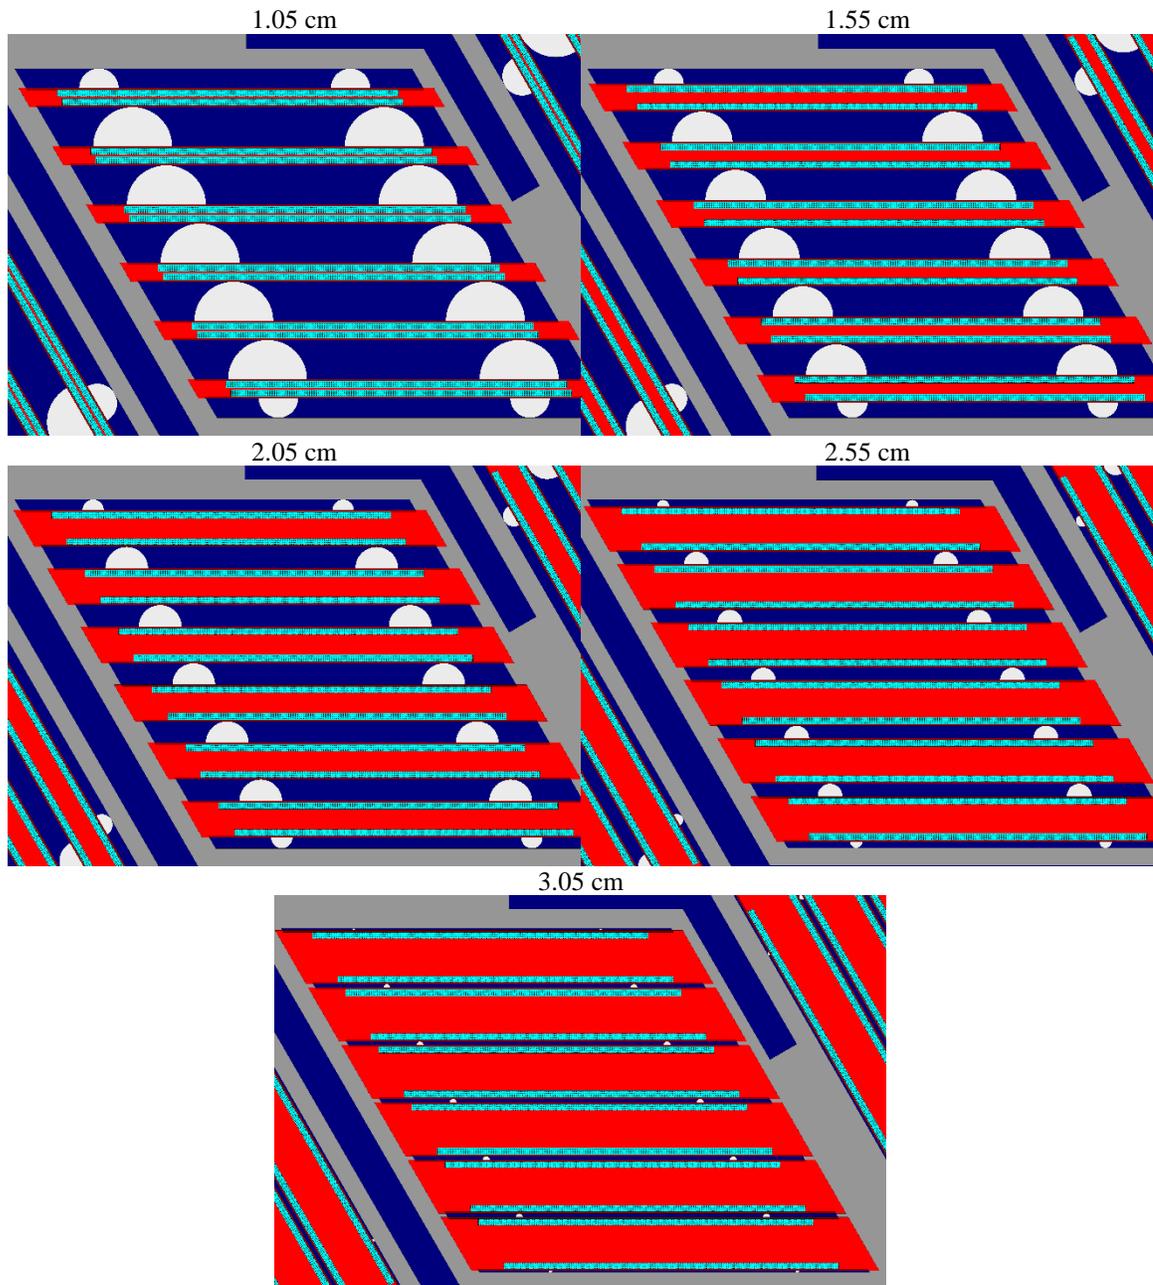


Figure 8.1. Images for assembly sections with five different plank thicknesses.

Table 8.2. Five fuel plank width cases with associated eigenvalues.

Plank Thickness [cm]	Total Cooling Channel Width [cm]	k_{eff}	Δk [pcm]
1.05	2.2	1.29735 [†]	-6631 ± 10
1.55	1.7	1.30523 [†]	-5843 ± 10
2.05	1.2	1.32828 [†]	-3538 ± 10
2.55 (reference)	0.7	1.36366 [†]	-
3.05	0.2	1.41051 [†]	4685 ± 10

[†]Reported statistical uncertainty was 7 pcm for these simulations.

The eigenvalue increases as the plank thickness increases. This is expected, as additional plank thickness adds more graphite to the assembly, which increases the moderation. This result agrees with previous studies of the AHTR fuel assembly which showed that increasing the CHM ratio raised the eigenvalue. A plot of the eigenvalue differences to the reference can be seen in Figure 8.2. The behavior can be well-described as being quadratic (best fit second-order equation has $R^2= 0.9998$). While the increasing behavior is expected, the positive second derivative is less intuitive. As more graphite is added to the fuel assembly, the relative impact of the addition becomes more significant.

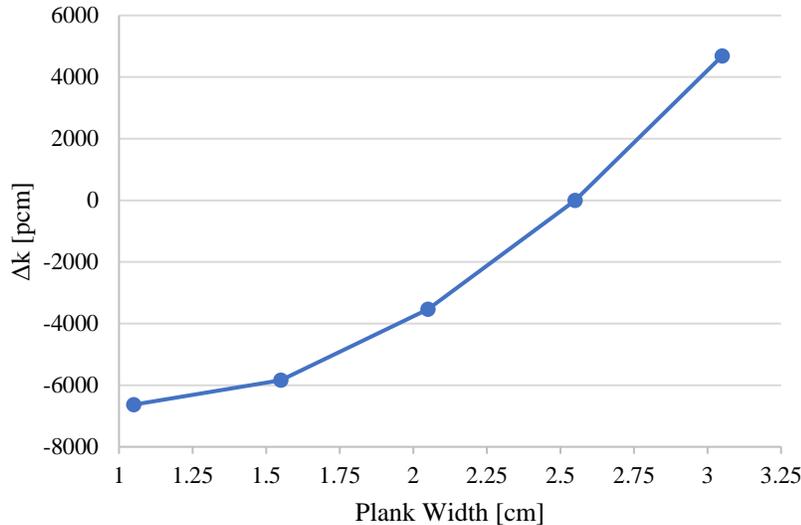


Figure 8.2. Eigenvalue impact due to changing fuel plank thickness.

8.1.2 Distance from Fuel Stripe to Coolant

Similar to the previous set of cases looking at varying the plank thickness, the focus here remains on making modifications to the fuel plank geometry and this study considers the impact of moving the fuel stripes within the fuel planks. The distance from the fuel stripes to the coolant boundaries (also referred to as the “sleeve” of the fuel plank in previous chapters) was changed for six cases, including the reference case of 0.1 cm. Additional cases include: 0.01 cm (very small distance; essentially fuel stripe at the coolant boundary), 0.3 cm, 0.5 cm, 0.7 cm, and 0.85 cm (relatively large distance; comparable to having a single large fuel stripe at the center of the fuel plank). A visualization of the case geometries can be seen in Figure 8.3.

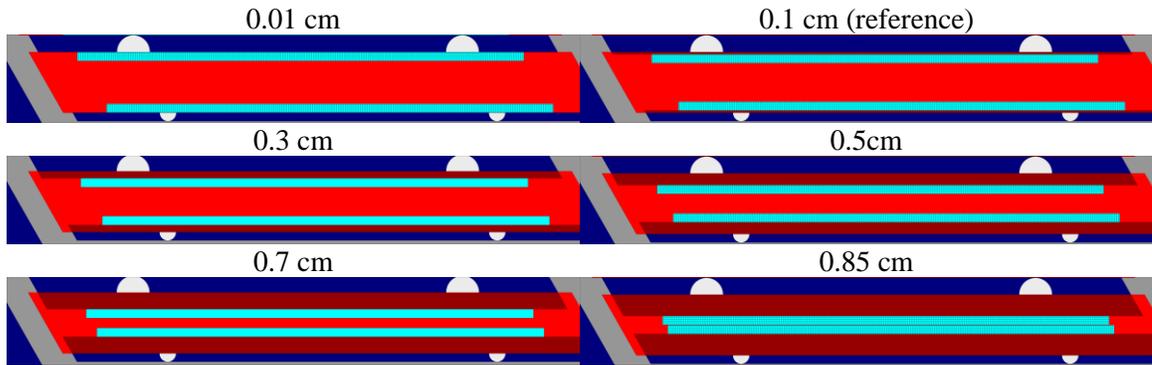


Figure 8.3. Varying separation of the fuel stripe from the coolant.

Simulations were conducted using otherwise reference AHTR assembly dimensions. All cases assumed that all CBs were fully withdrawn, and no BP particles present at the center of the fuel planks. Results for the simulations can be found in Table 8.3.

Table 8.3. Case summary and eigenvalue results for varying the plank sleeve thickness.

Sleeve Thickness [cm]	k_{eff}	Δk [pcm]
0.01	1.36571 [†]	205 ± 10
0.1	1.36366 [†]	-
0.3	1.36153 [†]	-213 ± 10
0.5	1.36455 [†]	89 ± 10
0.7	1.37150 [†]	784 ± 10
0.85	1.38683 [†]	2317 ± 10

[†]Reported statistical uncertainty was 7 pcm for these simulations.

An interesting observation is that the eigenvalue trend of Table 8.3 is nonlinear and even non-monotonic. In order to better visualize the results, the eigenvalue differences are plotted in Figure 8.4. The best-fit second-order equation only has $R^2=0.9775$, so a higher degree fit is needed to capture the trend. Observe that the most reactive case is for a sleeve thickness of 0.85 cm, which is when the two fuel stripes are very close together near the center of the fuel plank similar to if there was only a single fuel stripe. This study did not account for thermal effects or margins. Since the fuel and plank temperature would increase with the sleeve thickness, the negative temperature coefficients for both fuel and graphite would become relevant by lowering differences for the cases with fuel stripes embedded further into the fuel plank.

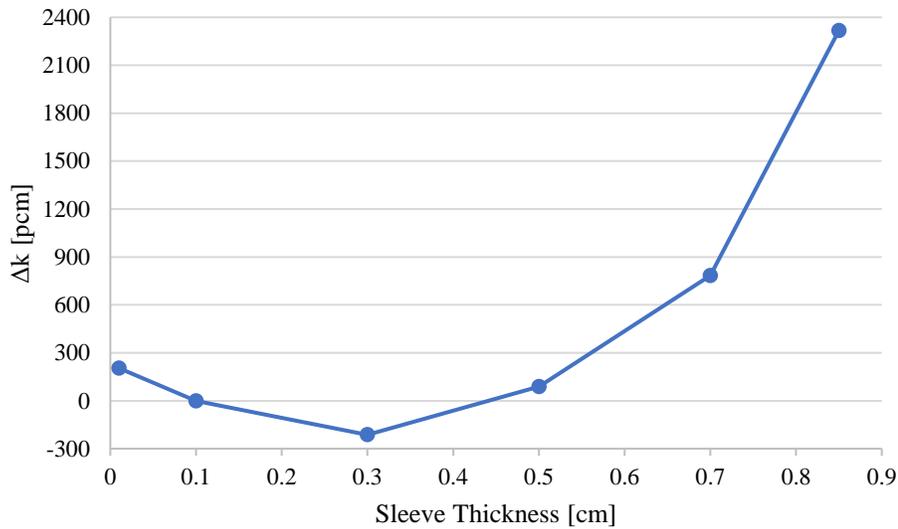


Figure 8.4. Eigenvalue impact of varying the sleeve thickness relative to the reference (0.1 cm).

8.1.3 Height of Axial Regions beyond the Active Core

Immediately above and below the active core region, the assembly geometry extends for another 25 cm with the only difference from the active core section being that the TRISO fuel and BP particles are no longer present. These axial sections are referred to as the axial reflectors. Axially further from the active core beyond the axial reflectors are the axial support plates. In previous studies as well as in this work, these axial support plates are homogenized due to their structural

complexity of numerous small components and overall distance from the active core region. The top axial support plate is a mixture of coolant FLiBe and silicon carbide. The bottom axial support plate is a mixture of coolant FLiBe and graphite.

This study considers the impact of varying the height of both the axial reflector and axial support plate regions. Four thicknesses were considered for both regions and each combination was simulated, producing sixteen datapoints. Again, due to the automated geometry generating capabilities of ATOMICS, this was possible with negligible time needed to prepare all 16 modifications of the reference 3D model. For the axial reflectors, the four values considered are: 1 cm, 25 cm (reference), 50 cm, and 100 cm. For the axial support plates, the four values considered are: 1 cm, 35 cm (reference), 50 cm, and 100 cm. Eigenvalue differences with respect to the reference case are shown in Table 8.4.

Table 8.4. Axial reflector and support plate thickness eigenvalue differences[†] [pcm].

		Axial Reflector Thickness [cm]			
		1	25	50	100
Axial Support Plate Thickness [cm]	1	-443	-151	37	152
	35	-181	-	112	168
	50	-151	12	96	157
	100	-147	23	107	162

[†]Statistical uncertainty is 10 pcm for all case comparisons.

Not surprisingly, the worst performance (lowest multiplication factor and highest axial peaking factor) is observed for the case with very little reflector present (1 cm for both regions). The general trend appears to be that the axial reflectors are more beneficial for the neutron economy than the axial support plates. For each axial support plate thickness row, increasing the axial reflector thickness increases the eigenvalue. No saturation effect is observed, even between the 50 and 100 cm axial reflector thickness results. This suggests that it may be necessary to model axial regions extending well beyond the active core. The same cannot be said about the axial support plates themselves, which appear to reach a saturation point more quickly with diminishing returns. For the cases considered, the differences between 50 and 100 cm axial support plate thicknesses

are within the statistical uncertainty of the simulations used to generate Table 8.4. This suggests that modeling homogenized regions similar in composition to the axial support plates yields no additional benefit beyond the first 50 cm.

In addition to providing mechanical support to keep fuel assemblies in the respective locations, the purpose of these two axial regions is to reflect neutrons back into the active core region to improve neutron economy and to flatten the power profile. The axial power profiles for each of the sixteen cases considered in this study are shown in Figure 8.5.

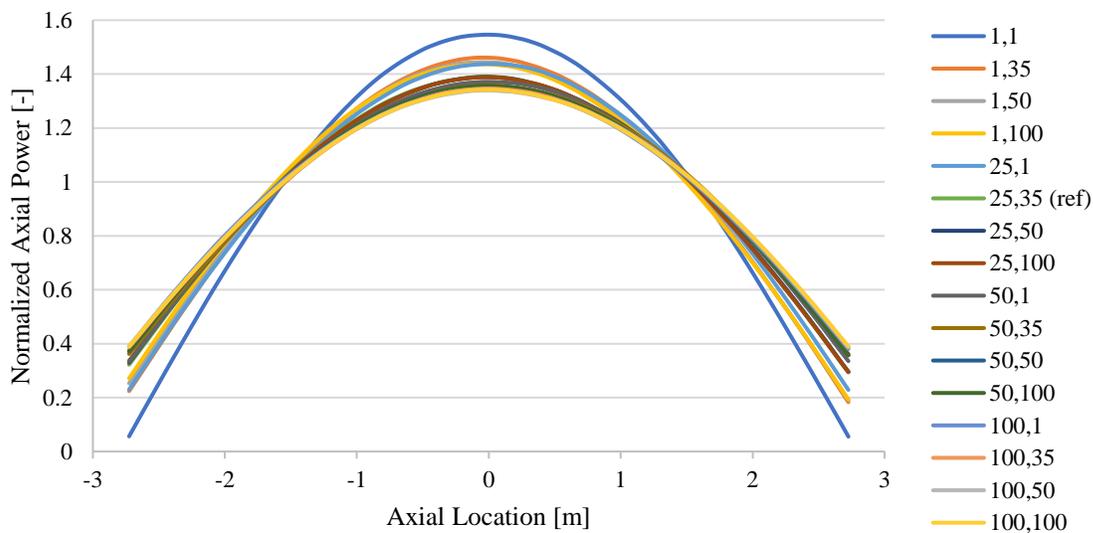


Figure 8.5. Normalized axial power profiles for sixteen axial cases considered. The first index in the legend refers to the axial reflector thickness [cm] while the second index refers to the axial support plate thickness [cm].

A summary of the axial PPFs and relative axial PPF differences between the profiles seen in Figure 8.5 can be found in Table 8.5 and Table 8.6, respectively. The reported values are taken with respect to the reference case (25 cm axial reflector thickness, 35 cm axial support plate thickness).

As expected, the case with the least axial media present (1 cm axial reflector, 1 cm axial support plate thicknesses) has the highest observed PPF. As with the eigenvalue results, there is also a clear trend that increasing the axial reflector thickness is beneficial for the neutron economy,

even up to 100 cm. As well for the axial support plate thickness, there appears to be a saturation effect between 50 and 100 cm. The exception to this is for the series with 1 cm axial reflector thickness, but this is an unrealistic scenario and can be attributed to having at least some reflector media is better than having very little.

Table 8.5. Axial reflector and support plate thickness axial PPF.

		Axial Reflector Thickness [cm]			
		1	25	50	100
Axial Support Plate Thickness [cm]	1	1.546	1.438	1.371	1.345
	35	1.461	1.391	1.359	1.346
	50	1.444	1.389	1.355	1.340
	100	1.436	1.389	1.357	1.342

Table 8.6. Axial reflector and support plate thickness axial PPF relative change [%].

		Axial Reflector Thickness [cm]			
		1	25	50	100
Axial Support Plate Thickness [cm]	1	11.14	3.34	-1.43	-3.34
	35	5.03	-	-2.33	-3.23
	50	3.80	-0.17	-2.62	-3.71
	100	3.25	-0.18	-2.48	-3.53

Due to the composition of the axial reflectors and axial support plates, each also has its own specific impact on the AO. The axial reflectors are identical in composition and size, so they would tend to promote a symmetric power distribution with a small AO. Differences would only be driven by thermal effects. This is seen in the left plot of Figure 8.6, where increasing the axial reflector thickness tends to reduce the axial support plate dependent variability in AO. On the other hand, the axial support plates have different compositions (graphite/FLiBe at the bottom and silicon carbide/FLiBe at the top), which means that they also have different neutron reflective properties. As can be seen again in the right plot of Figure 8.6, thick axial reflector isolates the core from the asymmetric axial support plates; promoting a small AO. Thicker axial support plates tend to push the power toward the bottom of the core. This suggests that the graphite/FLiBe mixture is neutronicly favorable to the silicon carbide/FLiBe mixture.

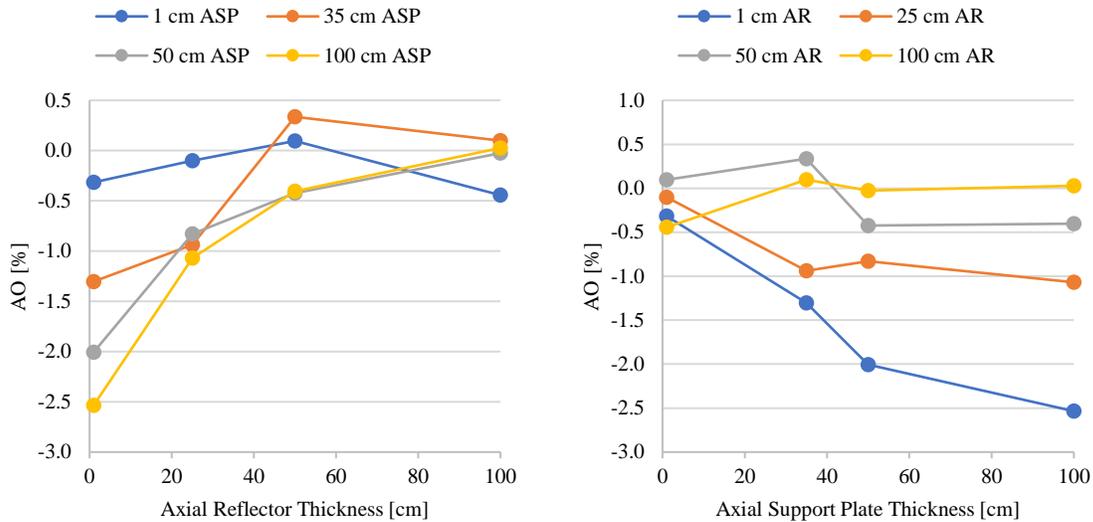


Figure 8.6. Left: Axial offset values for groupings of same axial support plate (ASP) thickness. Right: Axial offset values for groupings of same axial reflector (AR) thickness. Calculated statistical uncertainties for all cases are below 0.007%.

8.1.4 Number of Burnable Poison Columns

This study focuses on the eigenvalue impact due to adding BP columns at the center of the fuel planks. Five BP columns are used in the reference AHTR specifications. Additional cases of one, three, seven, and nine columns are investigated to observe the impact. Visualizations of the cases considered can be found in Figure 8.7. BP columns can be identified by the green particles in the middle of each fuel plank. Particles shown are about thirty-times larger than their physical size to enhance viewability.

Column pitch was adjusted for each case so that each column would have a more equitable impact over the fuel stripe length (18.71732 cm for these cases, from 202 particles with pitch 0.09266 cm). The BP axial pitch was also slightly adjusted to 0.1 cm from the reference cold value of 0.09936 (0.6% increase) in the interest of using round numbers in the test cases. A summary of each case with both its BP spacing and eigenvalue result can be found in Table 8.7. Note that cases were run with all CBs withdrawn and otherwise reference dimensions for components aside from the BP lattice changes listed. BP is europia (Eu_2O_3) with density 1.25 g/cm^3 .

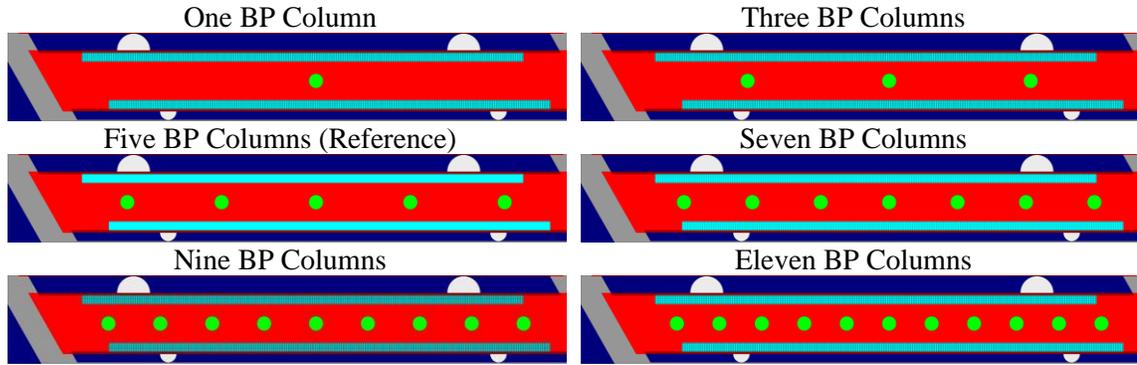


Figure 8.7. Fuel planks with different numbers of BP columns in the center.

Table 8.7. Eigenvalue impact from using a different number of BP columns.

Number of Columns	Column Pitch [cm]	Total Width [cm]	k_{eff}	Δk [pcm]
0	-	-	1.36889 [†]	14797 ± 7
1	-	-	1.33771 [†]	11679 ± 7
3	6	12	1.27696 [†]	5604 ± 7
5	4	16	1.22092 [†]	-
7	2.9	17.4	1.16985 [†]	-5107 ± 7
9	2.2	17.6	1.12306 [†]	-9786 ± 7
11	1.8	18	1.07948 [†]	-14144 ± 7

[†]Reported statistical uncertainty was 5 pcm for these simulations.

As expected, the eigenvalue decreases as more BP columns are added to the fuel planks. A visualization of the eigenvalue impact can be seen in Figure 8.8. The behavior is weakly quadratic with a best-fit of $R^2=0.999991$. The second derivative is slightly positive due to spectral hardening for each additional BP column placed in the model and energy self-shielding.

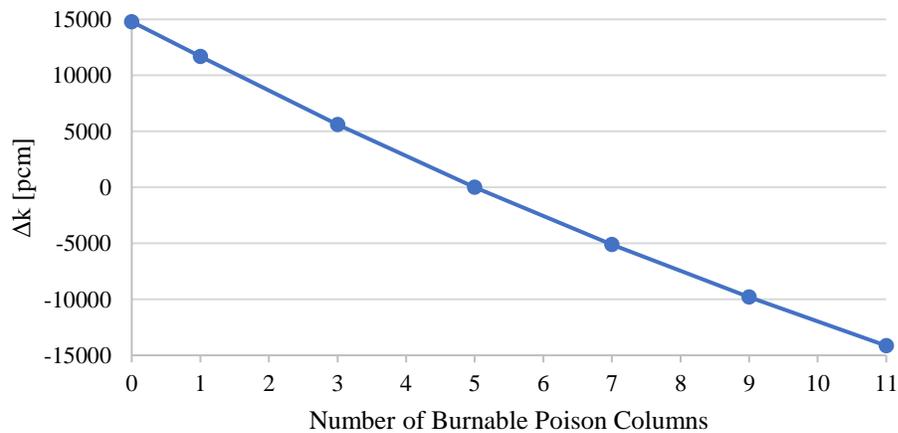


Figure 8.8. Eigenvalue difference from the reference case for adding BP columns.

8.1.5 Burnable Poison Column Pitch

This study focuses on the eigenvalue impact due to varying only the BP column pitch. The previous study prioritized adding BP columns but also adjusted the column pitch in the interest of having a fair comparison with the reference geometry and keeping the total column widths of the cases with higher column numbers (seven and nine) within the fuel stripe width (about 18.7 cm). This study aims to isolate the BP column pitch impact for the reference number of BP columns (five). Once again, in the interest of using round numbers, an axial BP particle pitch of 0.1 cm was used for the five cases considered here. Column pitches of 1 cm, 2 cm, 3 cm, 4 cm (reference), and 5 cm were simulated. A summary of the cases and their associated results can be found in Table 8.8.

Table 8.8. Eigenvalue impact from adjusting BP column pitch.

BP Column Pitch [cm]	Total Width [cm]	k_{eff}	Δk [pcm]
1	4	1.22429 [†]	503 ± 11
2	8	1.22279 [†]	353 ± 11
3	12	1.22098 [†]	172 ± 11
4 (reference)	16	1.21926 [†]	-
5	20	1.21736 [†]	-190 ± 11

[†]Reported statistical uncertainty was 8 pcm for these simulations.

Moving the BP columns further apart has a negative impact on the eigenvalue. This behavior is well described as linear (best fit first-order equation has $R^2=0.9985$) with a gradient of -174 pcm/cm for distance between BP columns. What is somewhat surprising is that the linear trend continues to hold valid for the final case with a total BP column span of 20 cm, which extends beyond the fuel stripe width of about 18.7 cm. Even though the final columns are no longer directly between two fuel stripes, the change is comparable to the result otherwise expected from the first four cases.

8.2 Thermal Sensitivity Studies

This section summarizes the results for a few sensitivity studies related to temperature variation. Studies include varying the coolant inlet temperature for zero power cases, impact of

core power using reference inlet coolant properties, and impact of graphite thermal conductivity on assembly temperature distribution. Results were obtained by using the TH iterative features of ATOMICS. At least four iterations were used for all cases. As will be discussed later in the depletion section though, the temperature profiles converge in fewer than four iterations for the cases considered based on when the iteration differences became comparable to statistical uncertainties of local power for the number of simulated particles (on the order of 10^7 to 10^8 total active particles).

8.2.1 Coolant Inlet Temperature

The coolant FLiBe used in AHTR has a melting temperature of 459 °C (732 K). This sets a practical lower bound for the core inlet temperature since pumping solidified FLiBe through the core is not feasible⁴. Although the FLiBe would not be circulating for temperatures below freezing, extrapolating the results for lower temperatures could provide meaningful insight for the required shutdown reactivity worth. Zero power cases using various coolant temperatures starting with the reference inlet temperature of 650 °C (923 K) were investigated for their impact on eigenvalue and required critical CB insertion. Simulations were run for both all CB withdrawn and critical CB insertion configurations. All cases used a BP europa density of 1.25 g/cm³. The results are summarized in Table 8.9. Eigenvalue (k_{eff}) results shown are for all CB groups withdrawn. CB insertion numbers correspond to a configuration achieving criticality within a tolerance of 100 pcm.

Table 8.9. Summary of results for varying inlet coolant temperature.

Temperature [K]	k_{eff}	σ [pcm]	CB Groups Inserted	Partial Insertion Fraction of Final Assembly Group (1/16 th)
732	1.26430	13	70	2
773	1.25899	13	68	9
823	1.25320	13	66	5
873	1.24737	12	64	3
923	1.24020	12	62	8

⁴ Circulating liquid FLiBe just above the freezing temperature is likely not feasible either, since the viscosity increases significantly at this temperature. There is about a factor of three difference between the viscosity at 950 K and 750 K. Below 750 K, it continues to sharply increase approaching the freezing point [1].

The eigenvalue results for the cases considered are plotted in Figure 8.9. The behavior is well-described as linear for the temperature range considered, with a temperature coefficient of -12.62 pcm/K when no CB are inserted into the core. This result is comparable to that found in similar studies on the AHTR [42]. The dominant driver of this value is the temperature coefficient of graphite.

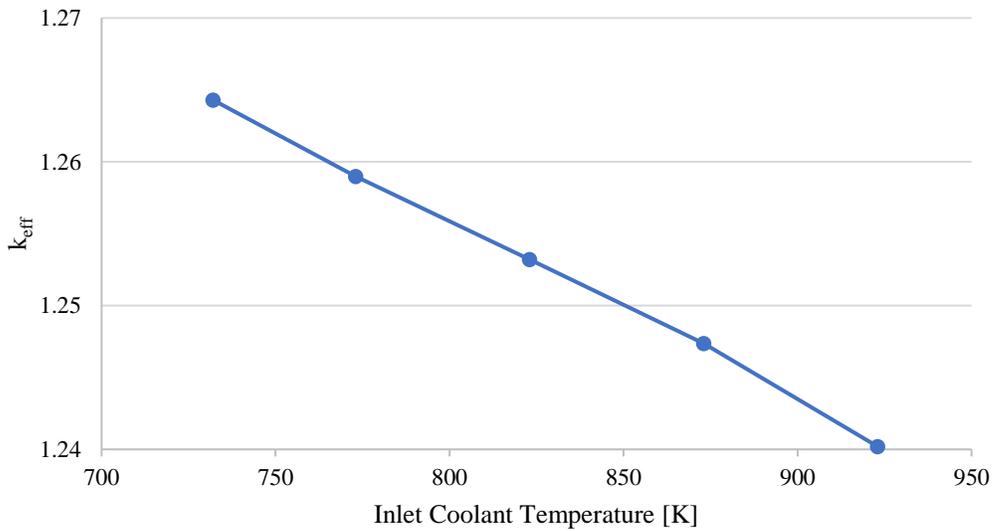


Figure 8.9. Eigenvalue impact of inlet coolant temperature for zero power cases.

The CB critical insertion results from Table 8.9 can be seen in Figure 8.10. Cases assumed an eigenvalue tolerance of 100 pcm for critical insertion. The behavior is well-described as linear for the temperature range considered, with an average value of -0.04 CB groups/K. This equates to one CB group being withdrawn from the core for about every 25 K change in inlet coolant temperature. An observation from this result is that using 1.25 g/cm³ density BP is likely not feasible from a licensing standpoint since if all 84 assembly CB groups were inserted, the core would reach a critical equilibrium point around 110 °C. This design BP loading would likely not be licensed nor operated and shows that a higher BP loading would be necessary (even despite the residual reactivity penalty from using europia BP as discussed in Chapter 6) to allow for a complete cold shutdown (at room temperature) and to account for additional shutdown margin. Although not

completely realistic, this core design and its simulations are still informative about how the AHTR operates and demonstrative of the capabilities of ATOMICS.

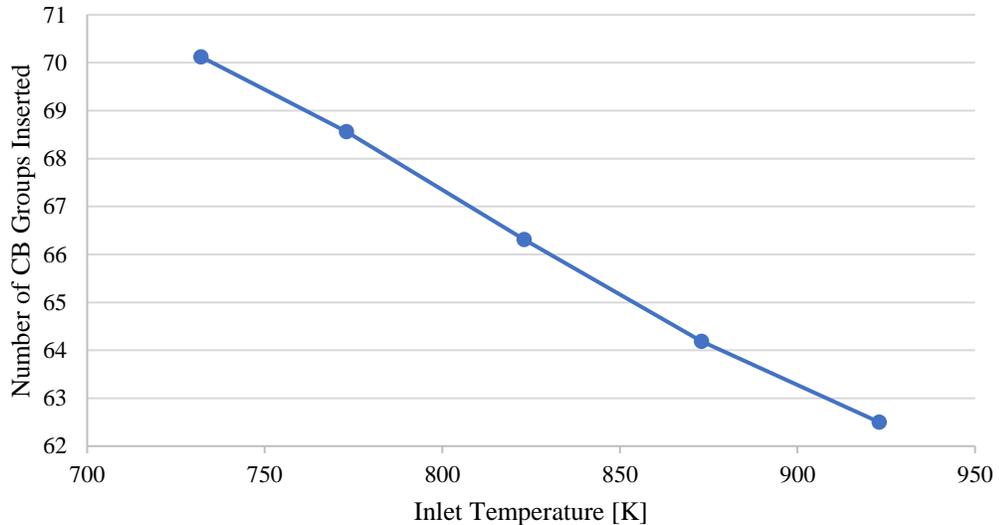


Figure 8.10. Number of CB groups required for criticality when varying inlet coolant temperature for zero power cases.

8.2.2 Core Power

This set of cases focused on changing the core power while keeping constant the core inlet temperature at the reference value of 923 K. Power changed from zero power to twice the reference power in 50% of reference power increments. The reference power is 3400 MW, so each 50% power increase corresponds to an increase of 1700 MW above the previous case. Simulations were run for both all CBs withdrawn and critical CB insertion configurations. A summary of the results can be found in Table 8.10 for all CBs withdrawn and Table 8.11 for critical insertion. All cases used a BP europa density of 1.25 g/cm³. Note that “Peak PPF” corresponds to the PPF of the section in the core producing the most power when resolved to 252 radial zones of one-third assembly sections of 84 one-third core symmetric assembly locations and 16 axial partitions (4032 zones total). These results do not account for the power variability within the section, but previous studies have estimated the local peak to be at most 20% above the section average [47].

Table 8.10. Summary of results for varying core power for all CBs withdrawn.

Power [%]	k_{eff}	σ [pcm]	Assembly PPF	Section PPF	Axial PPF	Peak PPF	Peak Temperature [K]
0	1.2402	12	1.882	2.058	1.383	2.828	923
50	1.22823	12	1.721	1.874	1.357	2.510	1098
100	1.21772	12	1.622	1.759	1.334	2.363	1254
150	1.2076	12	1.588	1.721	1.313	2.241	1396
200	1.19827	12	1.580	1.709	1.292	2.201	1548

Table 8.11. Summary of results for varying core power for critical CB insertion.

Power [%]	CR Insertion Groups	Partial Insertion (1/16 th)	Assembly PPF	Section PPF	Axial PPF	Peak PPF	Peak Temperature [K]
0	62	8	1.318	1.354	1.397	1.900	923
50	59	11	1.402	1.463	1.358	1.987	1037
100	57	4	1.332	1.372	1.351	1.856	1162
150	55	0	1.374	1.433	1.334	1.905	1343
200	52	8	1.369	1.413	1.328	1.920	1414

As the core power increases, the eigenvalue decreases, as expected due to thermal feedback. When power increases, temperatures of materials in the core also increase. Due to the negative temperature coefficient demonstrated in the previous set of studies, reactivity decreases. The reactivity behavior (for critical CB insertion) can be seen in Figure 8.11. The trend is approximately linear with a gradient of about -17.2 pcm per percent full power (-506 pcm/GW).

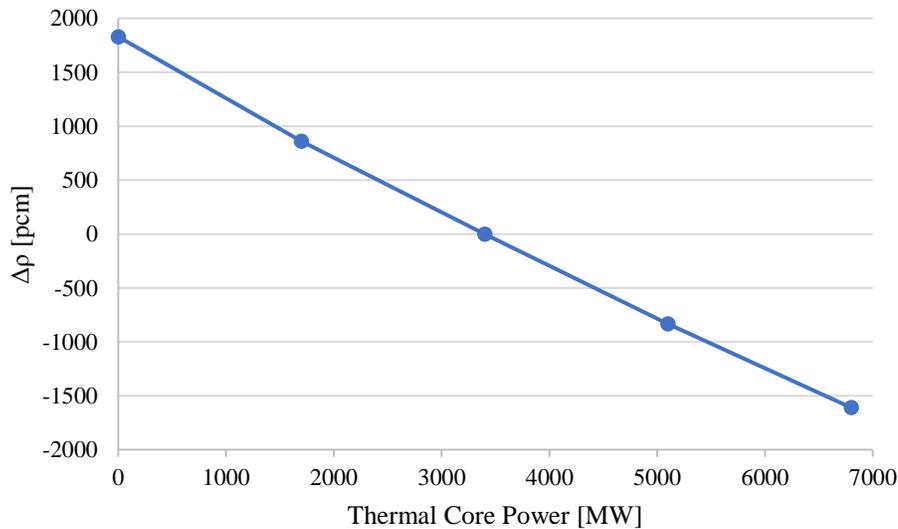


Figure 8.11. Eigenvalue impact due to varying the core power (all CBs withdrawn).

Increasing the core power also increases the peak local section temperature. The peak temperatures for both all CBs withdrawn and the critical CB insertion configuration can be seen in Figure 8.12. For all CBS withdrawn, the trend is approximately linear with a gradient of about 3.13 K per percent full power (92.0 K/GW). The critical insertion results are lower than the results for all CBs withdrawn due to the lower radial PPF for the critical insertion case. Favorably shaping the radial power profile to decrease peaking also decreases the peak fuel temperature. The observed critical insertion results have more variability due to the fact that different CB insertions shape the core power uniquely. Both the radial and axial power profiles are significantly impacted by the critical insertion configuration, which will directly impact the peak section power and its corresponding peak fuel temperature.

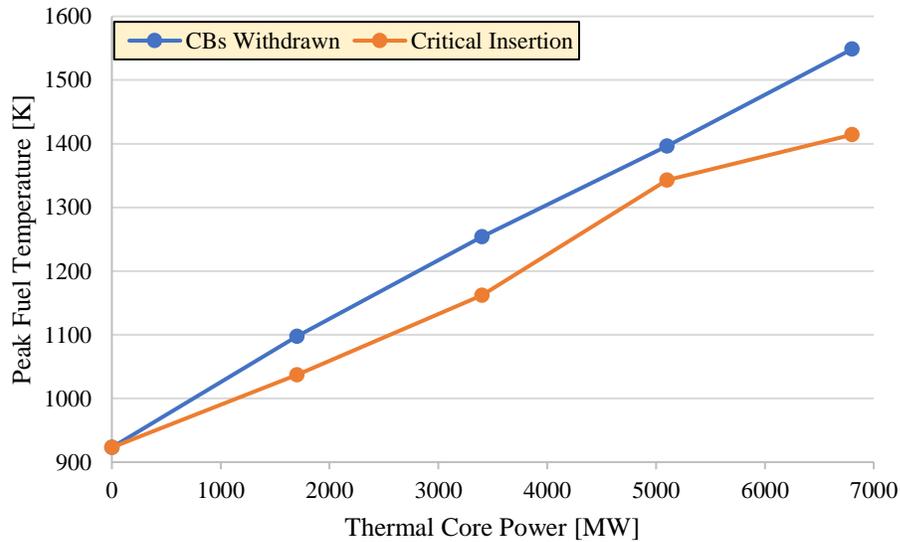


Figure 8.12. Peak temperatures based on core power.

PPFs resolved at the assembly, one-third assembly section, axial, and whole core section peak levels when varying thermal core power can be seen in Figure 8.13. Results are included for both the no CBs inserted (N) and the critical insertion (C) cases. As expected, the case with no CBs inserted has the highest PPFs which can be attributed to the radial performance. Both the assembly- and section-level PPFs decrease as core power increases when no CBs are inserted, which can be

attributed to thermal feedback. Due to the classical chopped-cosine shape for the power, temperatures increase faster at the center of the core than along the periphery. Negative temperature feedback reduced reactivity in this now hotter region which promotes a power profile with less production at the peaked center and more along the periphery, flattening the power profile. In both cases, axial peak varies only slightly with values near 1.4. Critical insertion radial results are comparable to the radial results, also having values near 1.4.

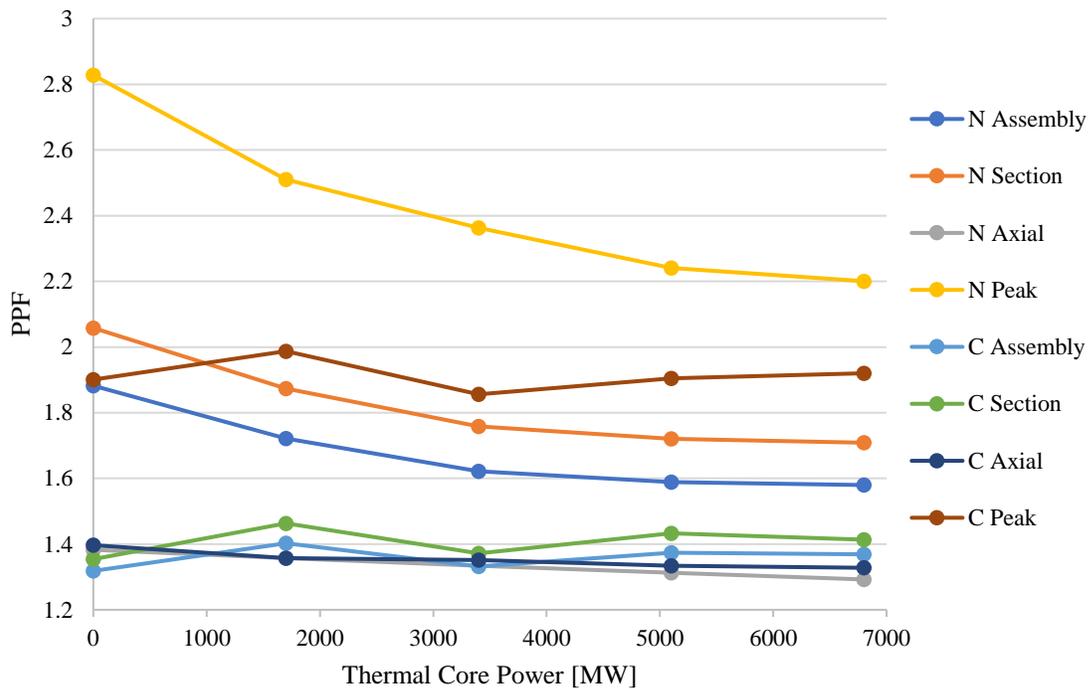


Figure 8.13. PPFs resolved at multiple levels when varying thermal core power for no CBs inserted (N) and critical CB insertion (C).

8.2.3 Graphite Thermal Conductivity

This set of cases focused on the impact of varying the thermal conductivity of graphite in the fuel assemblies. Since a vast majority of the AHTR fuel assembly is graphite, changing its thermal conductivity will have a significant impact on the TH results. Graphite thermal conductivities were varied from 15 (reference) to 120 W/(m·K) in increments of 15 W/(m·K). This range spans cases for fresh graphite with high initial manufactured thermal conductivity to highly

irradiated graphite with low initial manufactured thermal conductivity [48]. The lower value of 15 W/(m·K) is used as the reference for conservatism, with respect to both manufacturing variability and effects due to fluence. Actual thermal conductivities in the core would likely be higher for few-cycle core residency designs and high initial (unirradiated) graphite thermal conductivity. Simulations were run for all CBs withdrawn from the core. A summary of the results can be seen in Table 8.12. All cases used a BP europa density of 1.25 g/cm³.

Table 8.12. Summary of results for varying graphite thermal conductivity.

Graphite Thermal Conductivity [W/(m·K)]	k_{eff}	σ [pcm]	Assembly PPF	Section PPF	Axial PPF	Axial Offset [%]	Local PPF	Peak Fuel Temp [K]	dρ [pcm]
15 (reference)	1.21737	12	1.64	1.78	1.33	-3.7	2.37	1249	0
30	1.22211	12	1.67	1.81	1.35	-4.0	2.41	1180	388
45	1.22365	12	1.67	1.82	1.35	-4.4	2.43	1156	513
60	1.22467	13	1.70	1.85	1.35	-4.3	2.50	1144	596
75	1.22530	12	1.71	1.86	1.35	-4.6	2.52	1133	647
90	1.22565	13	1.72	1.87	1.36	-4.9	2.52	1123	676
105	1.22576	12	1.71	1.86	1.35	-4.1	2.50	1133	684
120	1.22599	13	1.70	1.84	1.36	-5.1	2.49	1118	703

As expected, decreasing the graphite thermal conductivity decreases reactivity. This is due to the negative temperature coefficient of the system and in particular graphite. As graphite thermal conductivity decreases, heat transfer from the fuel kernels to the coolant is reduced. Less efficient heat removal translates to higher component temperatures and a lower eigenvalue for the core. A plot of the reactivity change can be seen in Figure 8.14. The reactivity impact saturates for sufficiently high graphite thermal conductivities relative to the reference case. This is because the thermal conductivity is high enough that further increases have a negligible impact on the temperature profile. These results highlight the importance of carefully characterizing the graphite used in the AHTR design. The thermal conductivity difference alone from 15 to 90 W/(m·K) can account for more than a dollar's worth of reactivity change. Other thermophysical parameters will also have impacts of varying degrees which should be well understood before operation.

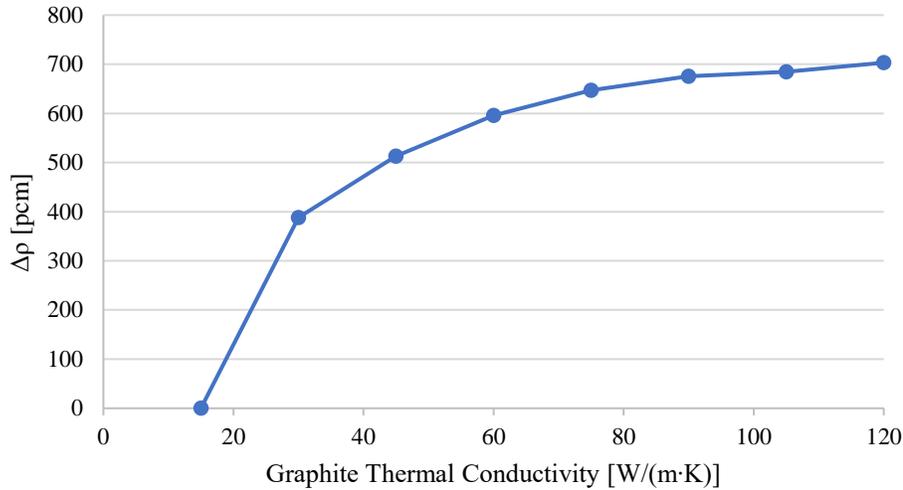


Figure 8.14. Reactivity impact of varying graphite thermal conductivity.

An important safety consideration during operation is the peak temperature observed in the portion of the core producing the most power. Many operating limits are set by this peak heat flux and peak temperature. The temperature distributions in the hottest one-third assembly section for sixteen active core axial partitions can be seen in Figure 8.15. For these cases, reference power is assumed and all CBs are withdrawn. Having no radial reactivity control results in a peaked power profile toward the center of the core, which further increases the hottest temperature observed. A critical system would likely see less extreme temperature variation, but these results illustrate the general impact on the temperature distribution. For the reference case of 15 W/(m·K) thermal conductivity for graphite, the peak fuel temperature is almost 300 K higher than the coolant temperature. Higher thermal conductivities significantly decrease the peak fuel temperature (by more than 100 K for thermal conductivities above 60 W/(m·K)) and greatly reduce the temperature variation between fuel particles⁵.

⁵ The plot for 105 W/(m·K) may appear to have a central plank temperature (right plot value) misplaced between the 75 and 90 W/(m·K) results. This is not a mistake, but instead can be attributed to the fact that its peak temperature location occurred in a section just above the midplane whereas all the other cases occurred in a section below the midplane. This is evident from the graphite “sleeve” location from 3.5 to 4.5 mm. All the other profiles have a lower coolant boundary temperature due to being lower in the core. Thus, the 105 W/(m·K) results are not erroneous but instead sampled from a neighboring location and effectively can be considered as shifted vertically. The reason for this disagreement in location is attributable to the statistical uncertainties in the local power tallies. For a power distribution with an AO near zero, the difference between the regions just above and below the centerline is small. For the simulations used to generate these results, the statistical uncertainty of local power was about 2% near the center of the core. It is reasonable for a few cases to report the local peak power location as being on the opposite side of the AO for small AOs and statistical uncertainties of this magnitude.

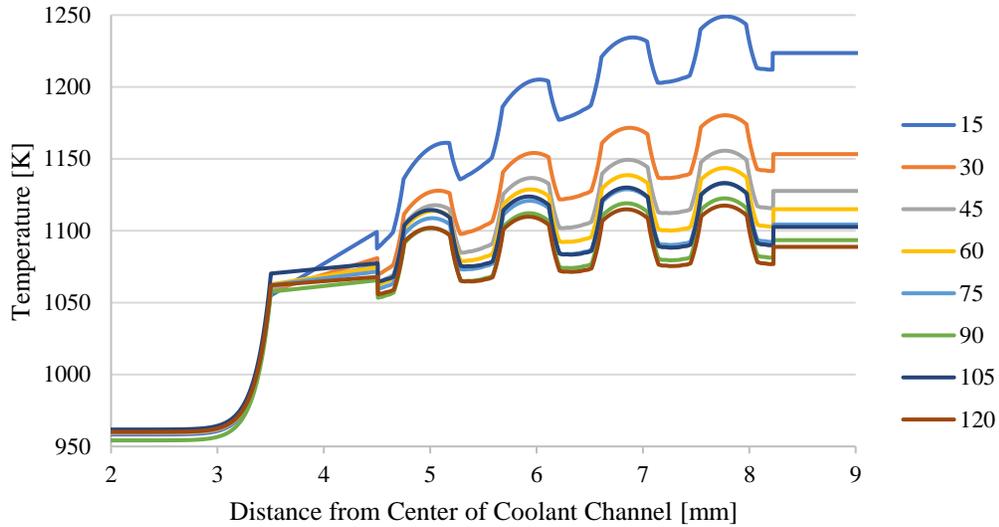


Figure 8.15. Hot channel temperature distribution for different graphite thermal conductivities [W/(m·K)].

This set of cases focusing on graphite thermal conductivity has implications for the neutron fluence-dependent behaviors of graphite. As graphite is irradiated, its thermal conductivity decreases as previously discussed in section 4.5. This set of results may be used to estimate the impact of neutron fluence over a depletion sequence for fuel assemblies residing in the core.

8.3 Depletion Simulations

This section details the results from running several depletion cases of the AHTR core. Discussion begins with a short overview of numerical stability and its challenges in relation to 3D MC depletion simulations. From there, results are presented for five depletion cases. Varying degrees of spatial resolution and physics tracking are considered and compared against each other. The first case uses only a single tracked fuel and single tracked BP material for the burnup sequence. This is a very simplified model, which often serves as a starting point for many fuel cycle studies. The second case uses 252 one-third assembly sections with one-third core symmetry in sixteen axial active core partitions. There are 4032 tracked regions in total, and this level of spatial resolution is denoted as “fine tracking” for ease of discussion. This degree of spatial resolution is comparable to conventional PWR analyses, which resolve the core to a quarter assembly level. The

third case considers TH feedback via temperature variation of components. Thermal feedback substeps are included between burnup steps for local power distribution convergence. The fourth case considers criticality iteration via CB insertion of one-third core symmetric groups. CB movement substeps are included between burnup steps for criticality convergence. The fifth case considered both TH feedback and criticality iteration. The substep search progresses by first finding the critical core configuration via CB movement and then progresses to additional TH iterative substeps between each burnup step.

The five sets of depletion studies all consider seventeen depletion points: 0.1, 0.5, 1, 2, 4, 6, 8, 10, 14, 18, 22, 26, 30, 40, 50, 60, and 70 MWd/kgHM. The total cumulative depletion length is 358 effective full power days (EFPD) for the reference power of 3400 MW_t and the fuel loading considered (9 w% enriched AGR-2 UCO fuel kernel with diameter 0.427 mm and density 10.9 g/cm³, 40% TRISO packing fraction, four TRISO layers, and fuel stripe width of 202 particles). All sets assume an initial HM loading of 17,406 kg, power density of 195.34 W/g, and BP europa (Eu₂O₃) density of 1.25 g/cm³, which is a quarter of the density prescribed in the original AHTR literature [2] [3]. The average burnup system becomes subcritical slightly below 50 MWd/kgHM and before 250 EFPD. The eigenvalue for an average burnup of 70 MWd/kgHM is about 0.92. Each simulation detailed in this section was run using a single compute node of 24 processors. The specific type of CPUs used were Dual Intel Xeon Gold 6226 @ 2.7 GHz. All transport simulations use 100 inactive cycles. This number was deemed as sufficient based on previous 3D simulations of AHTR. Eigenvalue and Shannon entropy by cycle can be seen in Figure 8.16 for different neutron generation sizes.

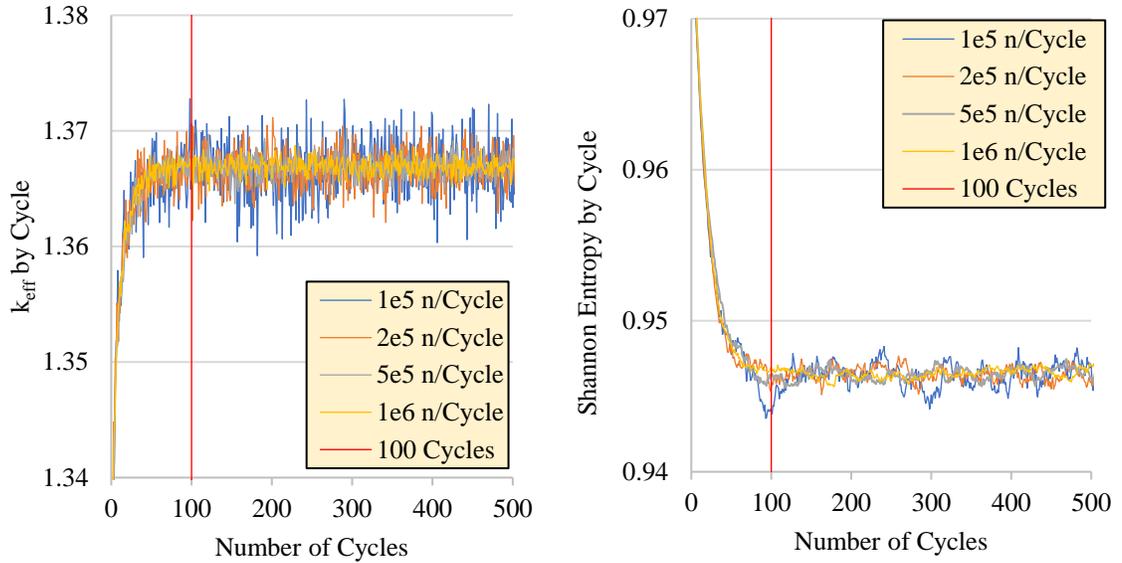


Figure 8.16. Eigenvalue (left) and Shannon entropy (right) by cycle for AHTR transport.

8.3.1 Numerical Stability in 3D Depletion Simulations

This section briefly discusses statistically-driven numerical oscillations which are inherent to 3D MC depletion simulations [49]. This discussion is by no means comprehensive, as there has been extensive efforts put into the topic over the past decade. The physical and numerical drivers will be addressed as well as options and features reviewed available to help mitigate the impact of these oscillations on depletion results.

In depletion simulations, one of the most important fission products to track are ^{135}Xe and its precursor ^{135}I . Due to the extremely high neutron capture cross section of ^{135}Xe , its presence in a reactor significantly impacts the reactivity in the core. For ^{235}U neutron fission reactions, ^{135}Xe is generated directly with a yield of about 0.25% as well as indirectly via precursor decay from ^{135}I (yield 3.1%, half-life 6.6 h) which has its own precursor ^{135}Te (yield 3.3%, half-life 19 s). As a fission product, the density of ^{135}Xe throughout the core is dependent upon the local fission rate which in turn is dependent upon the local flux. Regions with a higher fission rate generate more fission products including ^{135}Xe and ^{135}I , which decreases the local reactivity of the fuel. In 3D MC depletion simulations, knowing local fission rates is important to properly burn fuel materials based upon their power history.

This is challenged by the stochastic nature of MC simulations, which inherently have statistical uncertainty in the results. Physical changes in the local fission rate over cycle as well as statistically driven perturbations due to numerical uncertainty both contribute to changes in the local concentration of ^{135}Xe and other fission products over time. MC (as well as other methodologies) uses discrete time steps to estimate the isotopic changes within the core during operation. These discrete time steps coupled with the changes in local fission product concentration lead to time-dependent fission product reactivity penalties. When a local fission rate is too high, the next time step will both generate more local fission products and have depleted more fission products than should be. Due to the fission yield distribution and half-lives of the ^{135}Xe chain in particular, generating more fission products will mean predominantly ^{135}I and depleting more fission products will mean predominantly ^{135}Xe . This will result in a slightly higher local reactivity (due to less ^{135}Xe) which will later result in a slightly lower local reactivity (due to more ^{135}I which will decay into ^{135}Xe). In subsequent time steps, zones which were previously too reactive will become less reactive and vice versa. Numerical instabilities can occur when this oscillatory behavior amplifies between each step and could potentially become quite significant.

Local fission rate estimates can be further challenged in simulations which account for TH feedback. On top of fission product concentrations, local fission rates now additionally impact the local temperatures and densities which are also relevant for reactivity changes. This can further augment the same challenges detailed in the previous paragraph as physical and numerical differences now have larger reactivity worths; increasing the likelihood that a given perturbation would result in an unstable amplifying behavior. Despite these challenges, TH feedback should not be ignored as performing simulations using average operating properties may reduce accuracy of the results [50]. Features and methods have been introduced to try to alleviate these issues, but many of the ones currently implemented in MC codes do not absolutely eliminate instabilities and the ones that do come with significant computational overhead.

In a depletion simulation, material compositions (and thus cross sections) change between burnup steps. Since only the material composition at the current time step is known and the future composition has yet to be determined, the simplest way to deplete is to assume the cross sections of the current time step are an acceptable approximation to depletion to the next step. This assumption has varying degrees of success, but generally performs worse as the time between steps increases. A standard feature in many MC codes to better reflect the cross sections over the entire step instead of using the beginning of the step value is to use the predictor-corrector (PC) method. PC simulates an additional transport calculation to guess what the material compositions will be at the end of the step by using a constant extrapolation (CE) of the cross sections from the beginning of the step. Once the new compositions are found, their cross sections are averaged with the beginning of step results via linear interpolation (LI) to find average cross sections over the step when the actual depletion calculation is executed (red line in Figure 8.17.a). Although PC can help with numerical stability and potentially increase the time between burnup steps since using better estimators for the cross sections over the depletion sequence, it can still encounter numerical instabilities [51]. As before, coupling neutronics and TH when using PC will only make numerical instabilities more likely [52].

Most of the numerical challenges encountered with using MC depletion can be attributed to its tracking methods of highly absorptive and radioactive (i.e. time-dependent) fission products with relatively small concentrations which are susceptible to instabilities; even without further considerations of TH feedback. One approach to account for the time-dependent nature of important fission products is to implicitly track their concentrations via a physical analog metric instead of explicitly through following direct fission yields and indirect decay chains from fission product precursors. Serpent 2.31 has a feature which allows the code to track the *equilibrium* concentration of the fission products based on the neutron flux instead of following the irradiation history. This allows for instantaneous computation of the fission product concentrations without explicitly having to track losses due to both absorption and decay as well as generation from both

fission yield and precursor decay. These components are instead incorporated into equilibrium approximation equations. Example equations for ^{135}I and ^{135}Xe equilibrium concentrations can be found in Equation 8.1 and Equation 8.2 [45] [53]. n_I and n_X correspond to the concentrations, γ_I and γ_X correspond to the cumulative fission yields, and λ_I and λ_X correspond to the decay constants of ^{135}I and ^{135}Xe , respectively. Σ_f is the macroscopic fission cross section, σ_X is the microscopic capture cross section of ^{135}Xe , and Φ is the total flux.

$$n_I = \frac{\gamma_I \Sigma_f \Phi}{\lambda_I} \quad (8.1)$$

$$n_X = \frac{\gamma_X \Sigma_f \Phi}{\lambda_X + \sigma_X \Phi} \quad (8.2)$$

Equilibrium xenon treatment can help limit the occurrence of numerical oscillations in 3D depletion simulations using Serpent, but it still does not eliminate the possibility of instabilities. Further implicit treatment techniques have been developed and tested which guarantee solution stability [54]. Numerical stability can even be extended toward models using TH feedback [55]. The drawback to using numerically stable implicit methods in general is that there is a significant amount of computational overhead required to conduct the additional iterations necessary to find the correct profile. Methodologies have been developed and tested to only use computationally expensive implicit schemes over their faster explicit counterparts when unstable behaviors are detected in results [56], but they can still be quite taxing.

One stable implicit depletion scheme is the Stochastic Implicit Euler (SIE) method. Implicit Euler uses stable end of step cross sections for the depletion sequence. SIE takes the results from several essentially replica iterations to compute the average end of step cross sections to use for depleting to the next time step (Figure 8.17.c). Simulations can also utilize substeps between depletion time steps. Substep methods begin by estimating the end of step cross sections and assume linear dependence from the beginning to the end of the step (green line in Figure 8.17.b).

Cross sections of substeps can then be estimated as the average of beginning and end of substep values (dark red line in Figure 8.17.b). Having access to additional datapoints allows for higher-order interpolation for substeps. For example, quadratic interpolation can be used when utilizing cross sections from the previous step, current step, and next step [57].

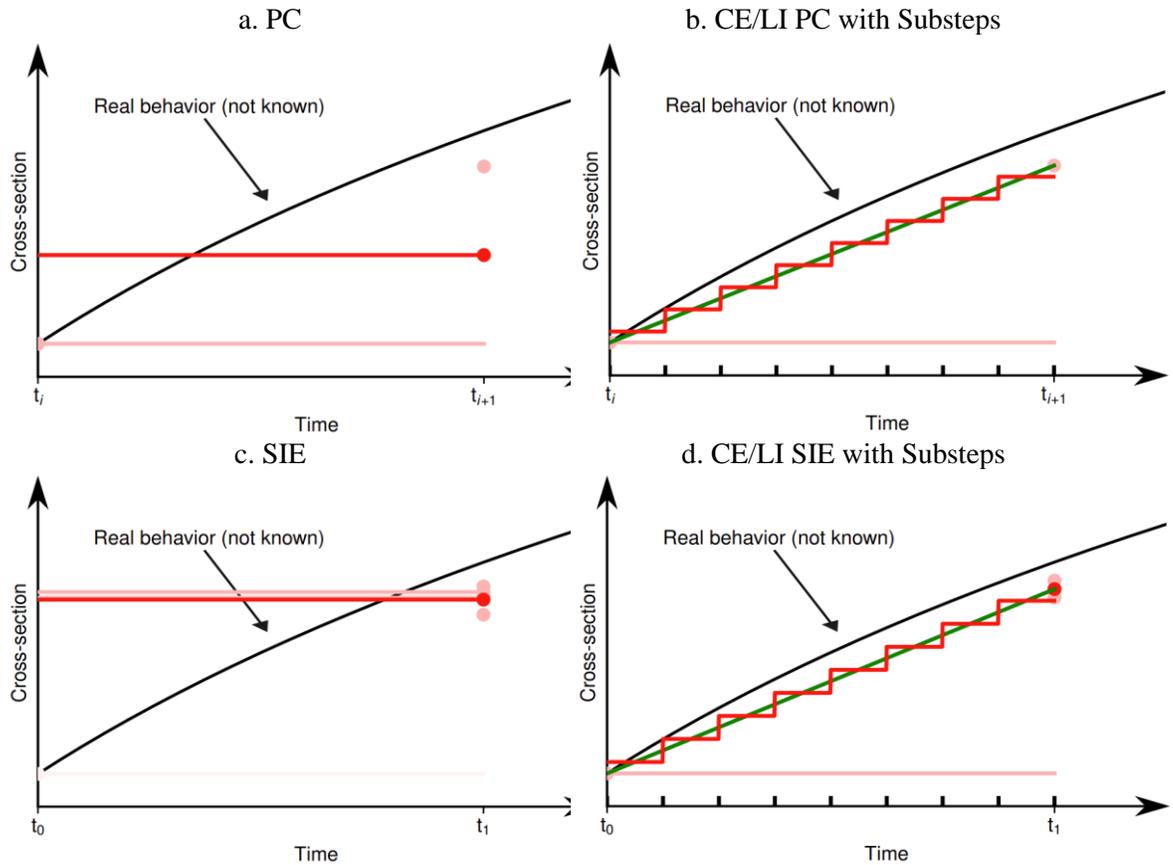


Figure 8.17. Burnup schemes with and without using depletion substeps [57].

The SIE method can be implemented in a few different ways and substep implementation advancements have improved accuracy especially for longer burnup steps [58]. When ATOMICS was developed, Serpent was limited to only using PC with the option to use equilibrium xenon tracking for fissile materials. Higher order methods beyond PC were not yet released with the distribution of Serpent 2.31. This changed with the release of Serpent 2.32 in February 2021, which allows for higher-order predictor and corrector methods. Additionally, substeps can be specified at

both the predictor and corrector level (previously unsupported). A summary of the new options available in Serpent 2.32 can be found in Table 8.13 [59]. In Serpent 2.31, only the first two options are available. The first option (pink line in Figure 8.17.a) simply uses the beginning of step results to deplete to the next step. The second option is PC when using no substeps (red line in Figure 8.17.a). Additional options enhance the numerical stability of simulations, especially when using substeps for longer depletion steps.

Table 8.13. New Serpent time integration depletion options introduced with version 2.32 [59].

Mode	Predictor Method	Corrector Method
CE	Constant Extrapolation	-
CE/LI	Constant Extrapolation	Linear Interpolation
LE	Linear Extrapolation	-
LE/LI	Linear Extrapolation	Linear Interpolation
LE/QI	Linear Extrapolation	Quadratic Interpolation
CE/CE	Constant Extrapolation	Constant Backwards Extrapolation

In this work, both PC and equilibrium xenon methods are used to reduce the occurrence of numerical instabilities. However, these methods do not guarantee solution stability and instabilities are observed for some considered depletion cases. The scope of this work was to create a tool with both TH iteration and criticality search capabilities for the AHTR. Methods to reduce the likelihood of numerical instabilities already a part of Serpent code package (version 2.31 at the time of ATOMICS development and running simulations) were used. Implementation of additional capabilities outside of the Serpent code distribution were outside the scope of this work. Numerical instabilities are obviously undesired but are viewed as an inherent part of some simulated cases considered in this work. The new methods now part of the Serpent code package as of February 2021 could potentially alleviate the concern of the instabilities encountered in the results. As will be addressed in the conclusions of this work (Chapter 9), ATOMICS could be modified to be compatible with Serpent 2.32 and simulations could be run utilizing the new depletion features available in Serpent.

8.3.2 Single Tracked Fuel and Burnable Poison Material

This section details results from simulations run using a single fuel and single BP tracked region. Both materials essentially deplete as the core average since all reactions are tallied against the entire core loading as a lumped region. Simulations were run for two xenon treatments: explicit and equilibrium. Explicit tracking follows the fission product and decay chain just like any other fission product where equilibrium tracking forces the xenon concentration based on the neutron flux, as described in Equation 8.1 and Equation 8.2. Both depletion cases used 250,000 particles per generation, 500 active cycles, and 100 inactive cycles per depletion step. Total wallclock runtimes were 352.4 h for explicit and 367.5 h for equilibrium xenon treatments. Results for each case as well as a comparison of the cases is presented below. Both cases assume all CBs are withdrawn and there is no TH feedback.

8.3.2.a Explicit Xenon Treatment

Results for the case using explicit xenon tracking over the depletion sequence can be found in Table 8.14. The eigenvalue over the sequence is presented in Figure 8.18. Note that the eigenvalue initially drops significantly due to the build-in of equilibrium of ^{135}Xe and other fission products. The eigenvalue plateaus for a period (2-10 MWd/kgHM) as fuel and BP deplete at comparable reactivity rates. After that (14-70 MWd/kgHM), the eigenvalue decreases linearly at a rate of about -411 pcm/(MWd/kgHM).

Since the depletion is tracked as a single material, spatial variation in the profile is expected to be minimal over the depletion sequence. Radial and axial power results are presented in Figure 8.19, confirming this expectation. The radial PPFs increase slightly with burnup, but the total change is less than 10% over the entire sequence. The axial power profiles for each depletion step are nearly identical, suggesting even less variation. Depletion cases presented later in this chapter considering finer spatial discretization will show larger variations in power with burnup, but these

single tracked material results serve as a reference for any initial scoping study depleting the core as a single tracked region.

Table 8.14. Depletion history for a single tracked fuel material with explicit xenon treatment.

Depletion Step	Burnup [MWd/kgHM]	Time [d]	k_{eff}	σ [pcm]
0	0	0	1.22028	7
1	0.1	0.513	1.19155	7
2	0.5	2.56	1.17521	7
3	1	5.12	1.17253	7
4	2	10.2	1.17110	7
5	4	20.5	1.17280	7
6	6	30.7	1.17430	6
7	8	41	1.17418	7
8	10	51.2	1.17297	6
9	14	71.7	1.16457	6
10	18	92.1	1.15219	7
11	22	113	1.13646	6
12	26	133	1.11897	6
13	30	154	1.10085	6
14	40	205	1.05584	6
15	50	256	1.01356	6
16	60	307	0.97368	6
17	70	358	0.93464	6

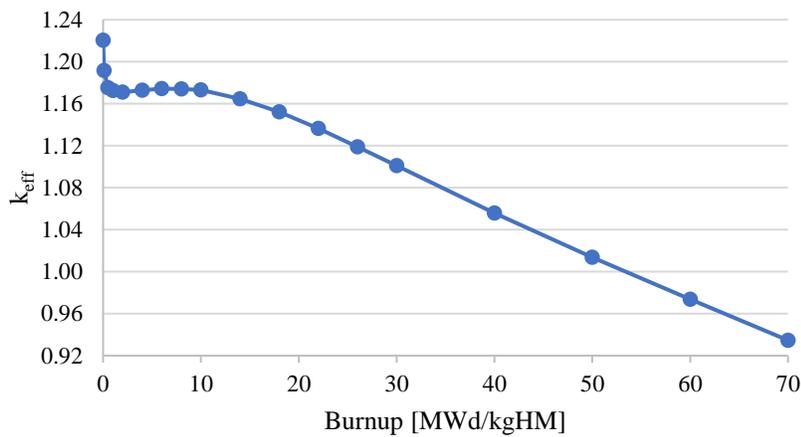


Figure 8.18. Eigenvalue over depletion sequence for single material (core average) tracking with explicit xenon treatment.

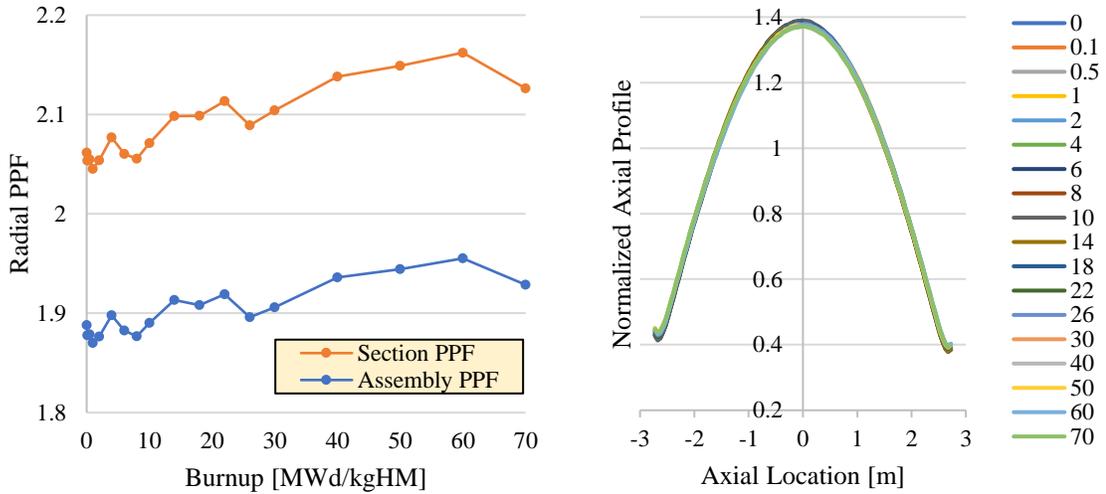


Figure 8.19. Power results for single material (core average tracking) with explicit xenon treatment. Left: radial PPFs for both the assembly and section levels. Right: axial power profiles over cycle (legend given for each burnup step in units of MWd/kgHM).

8.3.2.b Equilibrium Xenon Treatment

Table 8.15 presents the depletion results for the case using equilibrium xenon tracking. Note that there is no initial drop in the eigenvalue for this case since ^{135}Xe is assumed to be at the equilibrium concentration from the beginning of the simulation with no time to allow for build-in. The eigenvalue results over the depletion sequence can be visualized in Figure 8.20.

Table 8.15. Depletion history for single tracked material with equilibrium xenon treatment.

Depletion Step	Burnup [MWd/kg]	Time [d]	k_{eff}	σ [pcm]
0	0	0	1.18769	7
1	0.1	0.513	1.18704	7
2	0.5	2.56	1.18462	7
3	1	5.12	1.18233	7
4	2	10.2	1.18061	6
5	4	20.5	1.18213	7
6	6	30.7	1.18362	6
7	8	41	1.18362	6
8	10	51.2	1.18224	6
9	14	71.7	1.17464	6
10	18	92.1	1.16209	6
11	22	113	1.14634	6
12	26	133	1.12880	6
13	30	154	1.11016	6
14	40	205	1.06414	6
15	50	256	1.02017	6
16	60	307	0.97833	6
17	70	358	0.93750	5

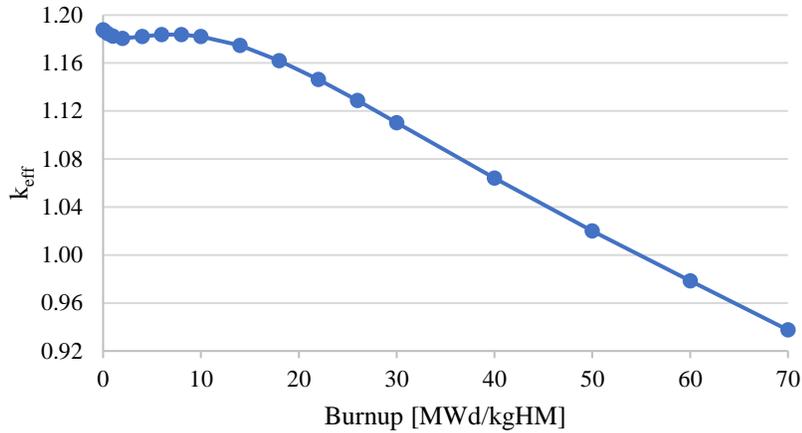


Figure 8.20. Eigenvalue over depletion sequence for single material (core average) tracking with equilibrium xenon treatment.

8.3.2.c Xenon Treatment Comparison

A comparison of the explicit and equilibrium xenon tracking methods within Serpent are presented in Table 8.16. There is initially a very large reactivity difference between the two sets, as expected since the explicit case does not contain any ^{135}Xe at the beginning. As the explicit case allows xenon to build-in, reactivity differences between the two sets of results continue to persist.

Table 8.16. Comparison of xenon treatments for a single tracked material.

Depletion Step	Burnup [MWd/kg]	Time [d]	Explicit Xe*	Equilibrium Xe*	Δk^\dagger [pcm]	$\Delta \rho^\dagger$ [pcm]
0	0	0	1.22028	1.18769	-3259	-2707
1	0.1	0.513	1.19155	1.18704	-451	-379
2	0.5	2.56	1.17521	1.18462	941	798
3	1	5.12	1.17253	1.18233	980	832
4	2	10.2	1.17110	1.18061	951	809
5	4	20.5	1.17280	1.18213	933	792
6	6	30.7	1.17430	1.18362	932	791
7	8	41	1.17418	1.18362	944	801
8	10	51.2	1.17297	1.18224	927	787
9	14	71.7	1.16457	1.17464	1007	861
10	18	92.1	1.15219	1.16209	990	856
11	22	113	1.13646	1.14634	988	866
12	26	133	1.11897	1.12880	983	875
13	30	154	1.10085	1.11016	931	842
14	40	205	1.05584	1.06414	830	783
15	50	256	1.01356	1.02017	661	650
16	60	307	0.97368	0.97833	465	476
17	70	358	0.93464	0.93750	286	306

*Individual eigenvalues have a peak statistical uncertainty of 7 pcm.

†Eigenvalue and reactivity differences have a statistical uncertain of 10 pcm.

A visualization of the reactivity difference between the two xenon treatments can be seen in Figure 8.21. After the initial expected reactivity jump as the explicit treatment allows ^{135}Xe to build-in, a persistent reactivity impact of about 800 pcm continues from about 1 to 30 MWd/kgHM. The impact decreases for higher burnups but does not entirely go away. The cause of this difference is attributable to the normalization process used by the equilibrium xenon treatment:

“Normalization ensures that when the flux in one part of the geometry is underestimated, it must be overestimated equally much somewhere else, but because the dependence of the xenon concentration on the flux is not linear, large overestimations of flux do not increase the xenon concentration as much as equally large underestimations decrease it. Thus the average xenon concentration is always underestimated, resulting in a bias.” [45]

Due to the nonlinearity of xenon concentration with neutron flux, the average xenon concentration is underestimated and thus more reactive than the results produced with explicit xenon treatment. The two xenon treatments produce inconsistent results for this case using core average properties, but as will be seen later in the next depletion study, using finer spatial tracking improves the agreement between explicit and equilibrium xenon tracking in Serpent.

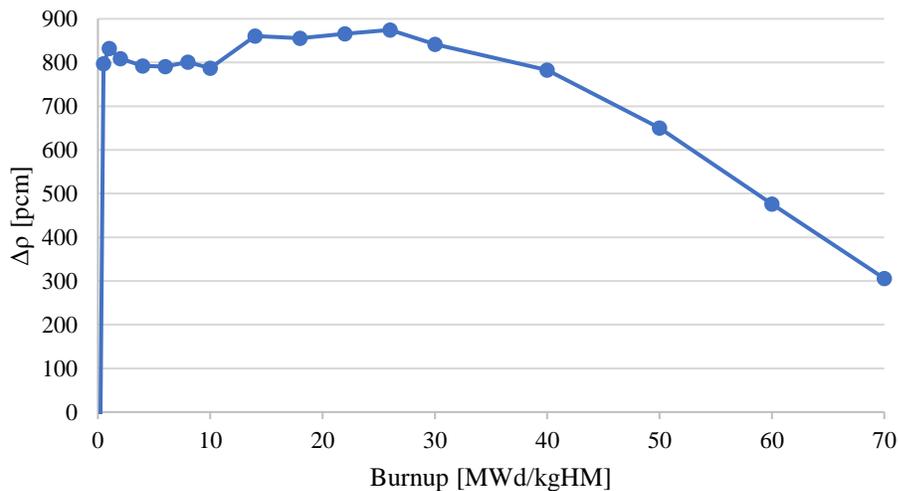


Figure 8.21. Reactivity differences over a depletion sequence using a single tracked material for both explicit and equilibrium xenon treatments.

8.3.3 One-Third Assembly and One-Sixteenth Axial (Fine) Tracking

Similar to the previous simulation set, this case considers depletion without active criticality control via CB movement (all CBs modelled as fully withdrawn) and without TH feedback. However, the spatial discretization is greatly enhanced. Instead of considering only a single core-average tracked material for both fuel and BP, this case subdivides the core into 4032 zones. Resolution is made at the one-third assembly level for 84 one-third core symmetric assembly groups using sixteen axial partitions. As mentioned earlier, this level of spatial resolution is referred to as “fine (spatial) tracking” for brevity in the remainder of results discussion.

Two sets of simulations were run for this case: one using explicitly tracked xenon concentration and another using equilibrium xenon concentration. Both cases use 250,000 particles per generation, 100 inactive cycles, and 500 active cycles for the seventeen burnup points. Predictor-corrector was used. The runtimes required for these simulations were 358.7 h and 367.6 h for explicit and equilibrium xenon treatments, respectively. Results for each simulation are presented in the following subsections with comparisons after.

8.3.3.a *Explicit Xenon Treatment*

Results for using an explicit xenon tracking in each of the 4032 fuel material regions can be found in Table 8.17. Note the large drop in eigenvalue initially as xenon quickly builds into its equilibrium concentration.

A significant concern in any depletion sequence is the numerical stability of the results. This is usually most evident in the axial power distribution as xenon oscillations driven by the statistical uncertainty inherent in explicit MC tracking can lead to unphysical axial power oscillations. The axial power distribution of each depletion step can be seen in Figure 8.22. There do not appear to be significant oscillations in the power distribution over the depletion sequence. The profile begins with the classically predicted truncated cosine shape which flattens due to higher power production in the center of the core burning those regions faster. Since TH is not considered

for this case, core material temperatures as well as coolant temperature and density are both radially and axially invariant. However, due to the differences in axial reflector composition, power is initially peaked toward the bottom of the core. Over the depletion sequence, the bottom of the core initially experiences larger depletion rates which can explain why the power peaks toward the top of the core for the final few depletion steps. The “jumps” seen in later depletion steps are attributable to the axial discretization of fuel materials into sixteen zones. The “jumps” occur because fuel near the center of the core has depleted more (less reactive) than fuel closer to the periphery (more reactive). These reactivity differences result in the power spikes observed for the later depletion steps.

Table 8.17. Depletion sequence summary for fine tracking with explicit xenon treatment.

Depletion Step	Burnup [MWd/kg]	Time [d]	Predictor		Corrector	
			k_{eff}	σ [pcm]	k_{eff}	σ [pcm]
0	0	0	1.22033	11	1.18921	11
1	0.1	0.5125	1.18892	12	1.17249	10
2	0.5	2.56	1.17217	11	1.16970	11
3	1	5.12	1.16968	10	1.16876	11
4	2	10.2	1.16896	12	1.16982	11
5	4	20.5	1.17028	11	1.16862	11
6	6	30.7	1.16929	11	1.16482	11
7	8	41	1.16518	12	1.15944	11
8	10	51.2	1.15978	11	1.14530	12
9	14	71.7	1.14599	12	1.12990	11
10	18	92.1	1.12998	12	1.11301	12
11	22	113	1.11308	11	1.09571	11
12	26	133	1.09542	12	1.07808	12
13	30	154	1.07757	11	1.03612	11
14	40	205	1.03437	12	0.99538	11
15	50	256	0.99299	12	0.95634	11
16	60	307	0.95338	12	0.91769	11
17	70	358	0.91504	11	-	-

The axial PPF over the considered depletion sequence can be found in Figure 8.23. Initially near BOC, the build-in of xenon slightly reduced the PPF. As expected, since the power profiles flattened with burnup in Figure 8.22, axial PPF decreases over cycle. Toward EOC, the axial PPF increases as the power gets more peaked toward the top of the core.

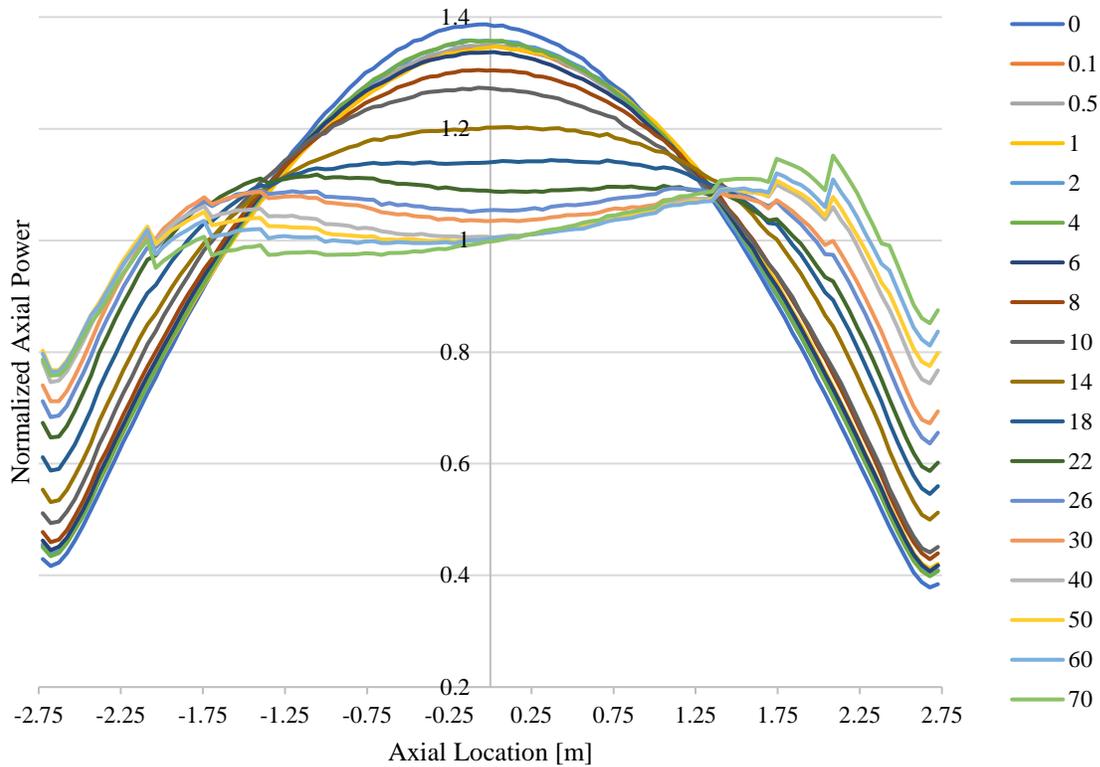


Figure 8.22. Normalized axial power distribution over a depletion sequence using fine spatial tracking. Legend values are for cumulative burnup points in units of MWd/kgHM.

The AO over the depletion sequence can be seen in Figure 8.24. Initially, the power is slightly peaked toward the bottom of the core (AO is negative) due to the different axial support plate compositions (FLiBe/graphite at the bottom and FLiBe/SiC at the top). There is no TH feedback for this case, so all temperatures and densities within the core are uniform. Initial peaking of the power toward the bottom half of the core depletes these regions more, which eventually results in the power shifting toward the top section of the core. This effect is especially pronounced due to lack of both TH feedback and active reactivity control via CB movement. TH feedback would have led to lower material temperatures and denser coolant in the bottom half of the core, which would push the power profile further toward the bottom. Additionally, since CBs enter the core from the top, any partial insertion of a CB group would further push the power toward the bottom of the core.

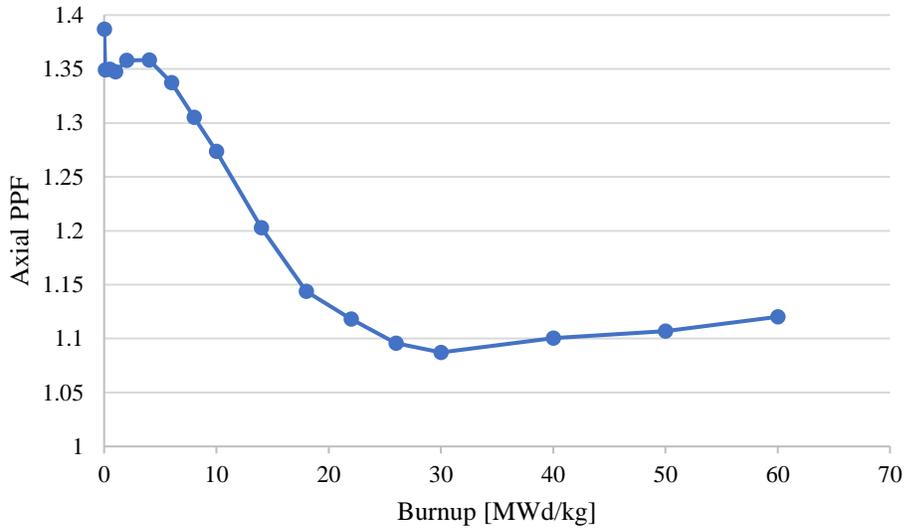


Figure 8.23. Axial PPF over depletion sequence using fine spatial tracking.

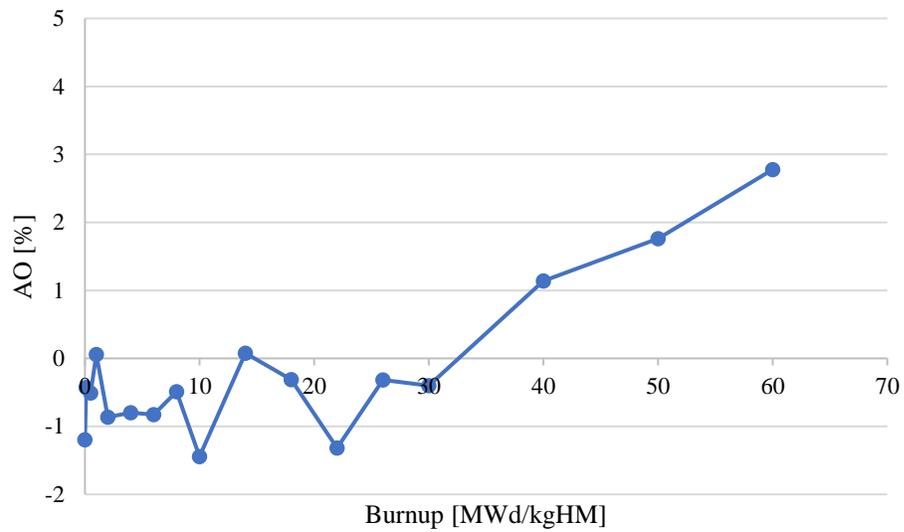


Figure 8.24. AO over depletion sequence using fine spatial tracking.

Analogous to the axial PPF, the radial PPF behavior is presented in Figure 8.25 at both the assembly and one-third assembly section levels. Much like the axial PPF, the radial PPF initially drops with the build-in of xenon and then decreases over the depletion sequence. The initial core loading uses a single uniform enrichment everywhere, which translates into a higher radial PPF than would likely be encountered for optimized operation. Quite quickly though, the radial PPF decreases to a reasonable value of about 1.1 around an average burnup of 18 MWd/kgHM or about

90 EFPD. The large difference in assembly- and section-level results near BOC likely reflect the large power gradient in the core due to the uniform core loading. A flatter profile later in the cycle would result in less significant differences in the assembly- and section-wise peaking, which is what is observed.

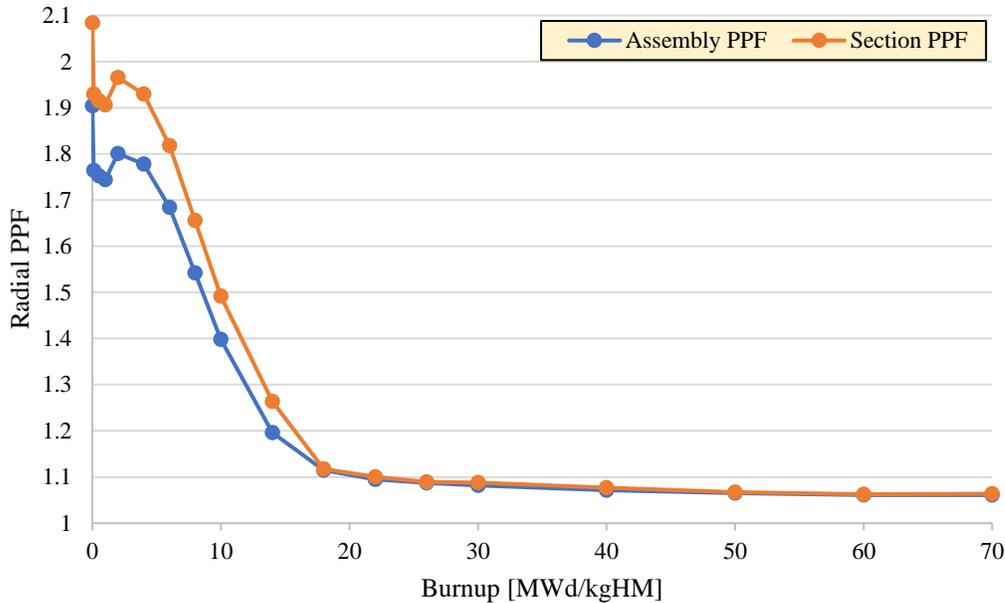


Figure 8.25. Radial PPFs over depletion sequence using fine spatial tracking.

8.3.3.b Impact of Number of Simulated Particles

To investigate the impact of particle statistics on the results, two shorter cases were run with one and two orders of magnitude fewer particles. The reference case used 1.25×10^8 active particles per depletion transport simulation while comparison cases were run for 1.25×10^7 and 1.25×10^6 active particles. The reference case using 1.25×10^8 particles took 358.7 h, the case using 1.25×10^7 particles took 43.9 h, and the case using 1.25×10^6 particles took 11.8 h. Nonlinear total runtimes are due to fixed requirements to initially process the complex geometry as well as transmute isotopics between transport calculations. The fixed requirements are about 8.3 h for processing with transport taking 350.4 h, 35.6 h, and 3.5 h for the individual cases respectively. A

plot of the reactivity impacts from using simulations with reduced number of particle histories can be seen in Figure 8.26.

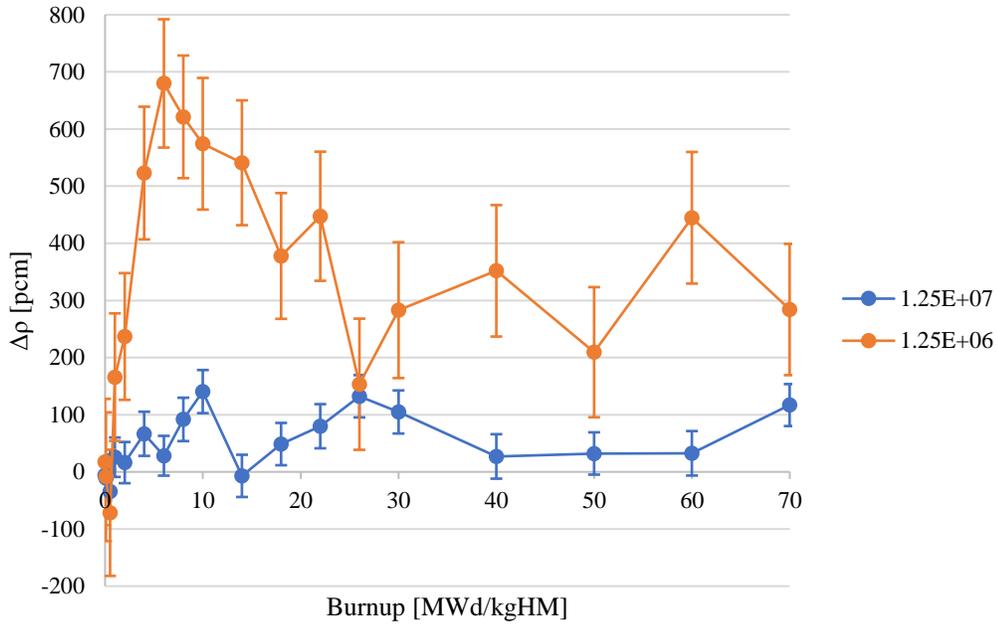


Figure 8.26. Reactivity impact over a depletion sequence due to varying the number of active particles used. Reference used 1.25×10^8 particles whereas the two comparison cases used 10-times fewer (Blue) and 100-times fewer (Orange) particles.

There are large differences observed for the case using the least particles (orange line in Figure 8.26). The peak difference of 680 ± 112 pcm is observed for a burnup of 6 MWd/kgHM (at about 30 EFPD). Differences decrease for this case down to about 350 pcm from 20 to 70 MWd/kgHM. This suggests running fewer particles biases the results by overpredicting the eigenvalue in the system. This is likely due to insufficient particles being run per generation (2,500 for the orange case). The power method used in MC simulations normalizes the fission source based on the results of the previous cycle and having a poor estimation by running too few particles in a single cycle will result in a bias [60]. The particles in each generation cannot adequately sample the entire core and results are biased toward the most reactive areas, artificially increasing the reported reactivity. The middle case (blue line in Figure 8.26) has much better agreement to the reference simulation. Even with considerations of computational overhead, using 10-times fewer

particles still reduces the total required simulation time by a factor of about 8.2. This is enough to be considered a worthwhile trade-off for scoping studies, whereas the case with the least number of particles only gains another factor of about 3.7, which might not be worthwhile. For the middle case, many depletion points are within the statistical uncertainty of the reference result. However, there does seem to be a persistent bias in overpredicting the results; even if only by about 50 pcm for this case. While certainly better than the orange case, reactivity biases are still present albeit small and might be acceptable given potential transport calculation runtime limitations.

8.3.3.c Equilibrium Xenon Treatment

The results from depleting the AHTR using fine spatial tracking and the equilibrium xenon tracking feature of Serpent can be seen in Table 8.18. A comparison between the explicit and equilibrium xenon tracking will be made in the next section. The axial and radial results had comparable values to their explicitly tracked counterparts and are not presented here. Since the explicit xenon tracking results were stable, one would have assumed the equilibrium xenon results would also have been stable which is the case.

Table 8.18. Depletion sequence summary for fine tracking with equilibrium xenon treatment.

Depletion Step	Burnup [MWd/kg]	Time [d]	Predictor		Corrector	
			k_{eff}	σ [pcm]	k_{eff}	σ [pcm]
0	0	0	1.17731	11	1.17670	11
1	0.1	0.5125	1.17663	11	1.17341	11
2	0.5	2.56	1.17341	10	1.17112	11
3	1	5.12	1.17100	11	1.16997	11
4	2	10.2	1.17011	11	1.17124	11
5	4	20.5	1.17148	11	1.17000	11
6	6	30.7	1.17048	11	1.16602	11
7	8	41	1.16642	11	1.16059	12
8	10	51.2	1.16089	11	1.14637	11
9	14	71.7	1.14712	11	1.13081	11
10	18	92.1	1.13112	11	1.11412	12
11	22	113	1.11411	11	1.09681	12
12	26	133	1.09640	12	1.07907	11
13	30	154	1.07854	11	1.03688	11
14	40	205	1.03511	11	0.99583	11
15	50	256	0.99372	11	0.95687	11
16	60	307	0.95411	11	0.91830	12
17	70	358	0.91536	11	-	-

8.3.3.d Xenon Treatment Comparison

Here, the results from the explicit and equilibrium xenon treatments are compared for fine spatial tracking. The predictor eigenvalue is presented in Table 8.19 for each treatment case along with their difference for each burnup step considered. Note that the largest difference occurs at BOC, since the equilibrium xenon treatment assumes equilibrium xenon concentration from the beginning of the depletion sequence with no time for initial build-in toward the equilibrium concentration. Quite quickly after the first few depletion steps (within three EFPD), the two cases agree within about 125 ± 16 pcm. Toward EOC, differences continue to decrease down to 32 ± 16 pcm at 70 MWd/kgHM. This is quite good agreement and demonstrates that if numerical instability issues are ever encountered when using explicit xenon treatment, employing the equilibrium xenon tracking feature of Serpent should produce comparable results over the cycle despite initial differences early in the cycle.

Table 8.19. Comparison of xenon treatments for fine spatial tracking.

Depletion Step	Burnup [MWd/kg]	Time [d]	Explicit Xe*	Equilibrium Xe*	Δk [pcm]
0	0	0	1.22033	1.17731	-4302 \pm 16
1	0.1	0.513	1.18892	1.17663	-1229 \pm 16
2	0.5	2.56	1.17217	1.17341	124 \pm 16
3	1	5.12	1.16968	1.17100	132 \pm 16
4	2	10.2	1.16896	1.17011	115 \pm 16
5	4	20.5	1.17028	1.17148	120 \pm 16
6	6	30.7	1.16929	1.17048	119 \pm 16
7	8	41	1.16518	1.16642	124 \pm 16
8	10	51.2	1.15978	1.16089	111 \pm 16
9	14	71.7	1.14599	1.14712	113 \pm 16
10	18	92.1	1.12998	1.13112	114 \pm 16
11	22	113	1.11308	1.11411	103 \pm 16
12	26	133	1.09542	1.09640	98 \pm 16
13	30	154	1.07757	1.07854	97 \pm 16
14	40	205	1.03437	1.03511	74 \pm 16
15	50	256	0.99299	0.99372	73 \pm 16
16	60	307	0.95338	0.95411	73 \pm 16
17	70	358	0.91504	0.91536	32 \pm 16

*Individual eigenvalues had a peak statistical uncertainty of 12 pcm.

A plot of the eigenvalue over the depletion sequence for both explicit and equilibrium xenon treatments can be seen in Figure 8.27. Note that there is a significant initial difference

between the two as the explicit xenon treatment accounts for the build-in time of xenon. The two quickly converge and have excellent agreement over the remainder of the cycle.

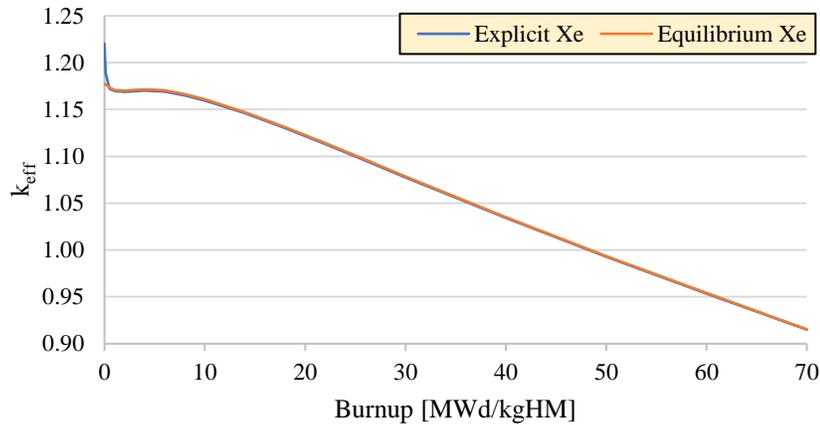


Figure 8.27. Eigenvalue over depletion sequence for fine spatial tracking.

Various PPFs (whole core-, one-third assembly section-, assembly-, and axial-level) over the depletion sequence can be seen in Figure 8.28 for both xenon treatments. There appears to be excellent agreement between the two sets. There are initial differences at BOC of course due to build-in toward equilibrium xenon concentration for the explicit case, but the results closely match beyond the initial few depletion steps.

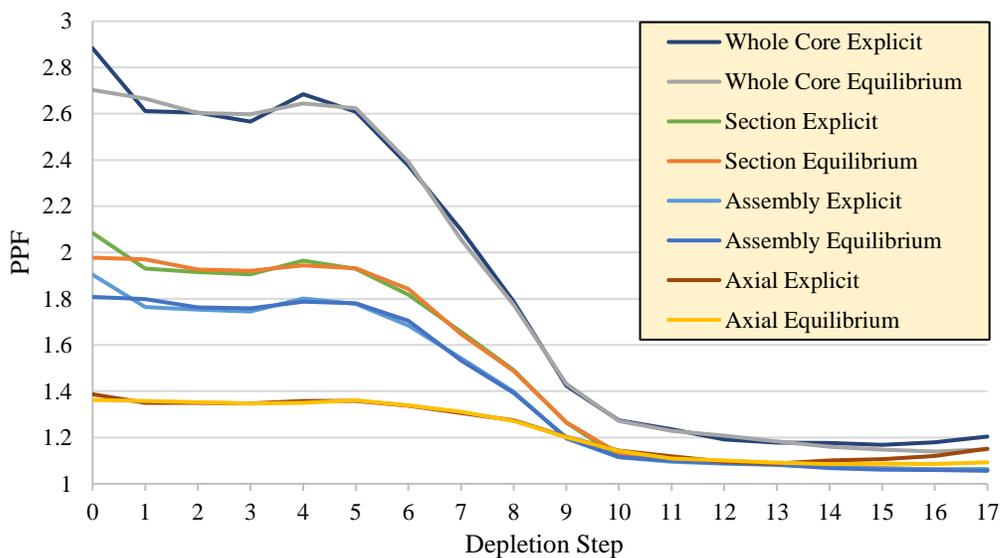


Figure 8.28. PPFs over depletion sequence for both xenon treatments using fine spatial tracking.

8.3.3.e Comparison to Single Material Tracking

In Table 8.20, the results between the single material (core average) tracking and fine spatial tracking are compared for the explicit xenon treatment cases. Differences are initially small as expected since they are primarily driven by statistical uncertainty for no burnup but grow to a few thousand pcm difference over cycle due to the different isotopic evolution in the two cases.

Table 8.20. Comparison between single and fine material tracking cases for explicit xenon usage.

Depletion Step	Burnup [MWd/kg]	Time [d]	Single Tracking*	Fine Tracking*	Δk^{\dagger} [pcm]	$\Delta \rho^{\dagger}$ [pcm]
0	0	0	1.22028	1.22033	-5	-4
1	0.1	0.5125	1.19155	1.18892	263	221
2	0.5	2.56	1.17521	1.17217	304	259
3	1	5.12	1.17253	1.16968	285	243
4	2	10.2	1.17110	1.16896	214	183
5	4	20.5	1.17280	1.17028	252	215
6	6	30.7	1.17430	1.16929	501	428
7	8	41	1.17418	1.16518	900	769
8	10	51.2	1.17297	1.15978	1319	1131
9	14	71.7	1.16457	1.14599	1858	1608
10	18	92.1	1.15219	1.12998	2221	1946
11	22	113	1.13646	1.11308	2338	2079
12	26	133	1.11897	1.09542	2355	2127
13	30	154	1.10085	1.07757	2328	2137
14	40	205	1.05584	1.03437	2147	2054
15	50	256	1.01356	0.99299	2057	2050
16	60	307	0.97368	0.95338	2030	2107
17	70	358	0.93464	0.91504	1960	2119

*Single tracked material cases had a statistical uncertainty of 6 pcm. Fine tracked cases had 11 pcm.

[†]Eigenvalue and reactivity differences have a statistical uncertainty of 13 pcm.

A visualization of the reactivity impact between the two spatial resolutions for depletion is given in Figure 8.29. There is an initial difference of about 250 pcm for 0.5 to 4 MWd/kgHM which can likely be attributed to correctly resolving the spatial distribution of equilibrium ^{135}Xe and other fission products. After this point, the differences steadily grow to around 2000 pcm at 22 MWd/kgHM. This difference can likely be attributed to the spatial depletion of europia BP particles, which depletes more rapidly than fuel. After 22 MWd/kgHM to the final depletion step of 70 MWd/kgHM, the reactivity impact remains constant just above 2000 pcm. This difference is likely attributable to the spatial depletion of TRISO fuel particles.

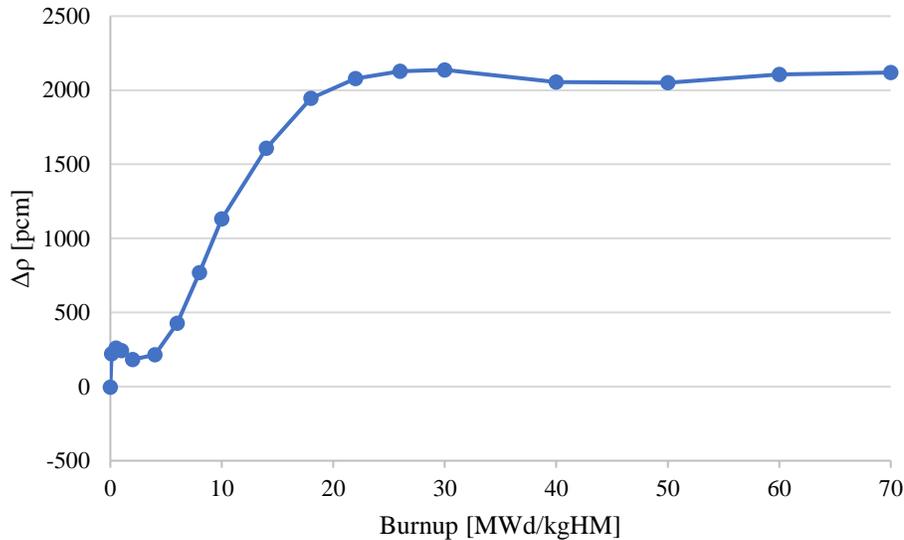


Figure 8.29. Reactivity differences for different material tracking granularity (differences taken with respect to fine material tracking resolution).

8.3.4 Fine Spatial tracking with Multiphysics

The previous section considered cases which used fine spatial tracking in the active core region with 4032 uniquely tracked zones. This section builds on that work by adding TH iterations to the depletion sequence. TH substeps are conducted between burnup steps a user-specified number of times by a parameter included in the user input file *options.txt*. This section considers the impact of applying different numbers of TH iteration substeps and using the same two xenon treatments as done previously.

8.3.4.a Number of Thermal Hydraulic Iterations between Burnup Steps

This subsection investigates how many TH iteration substeps are necessary in an example depletion study to accurately converge on the TH profile between burnup steps. It is important to adequately converge on the TH profile for accuracy of results. Convergence here is loosely interpreted as the temperature and power profile negligibly changing for additional iterations such that the depletion results are not materially impacted. Insufficient convergence would produce results with differences which are greater than those expected from statistical uncertainty. Too many iterations would waste time running transport simulations which do not improve the accuracy

of the results. This means that either the depletion sequence could have completed sooner or run with more particles per step to reduce statistical uncertainty for better converged results in the same amount of runtime. This subsection did not use the PC method within Serpent. This was done purposefully to investigate the numerical stability of simulations without it. All simulations use the equilibrium xenon treatment.

Five different TH substep values were compared in this study from zero to four iterations between depletion steps. For zero substeps, TH feedback is considered for the beginning of the burnup step with depletion immediately commencing based on those calculated values. The non-zero substep cases run iterative transport simulations to better converge on the temperature and power profiles. All cases used 100,000 particles per generation, 100 inactive cycles, and either 200 active cycles for TH substeps or 500 active cycles for depletion steps. The eigenvalue results for the different number of TH iterations can be seen in Figure 8.30.

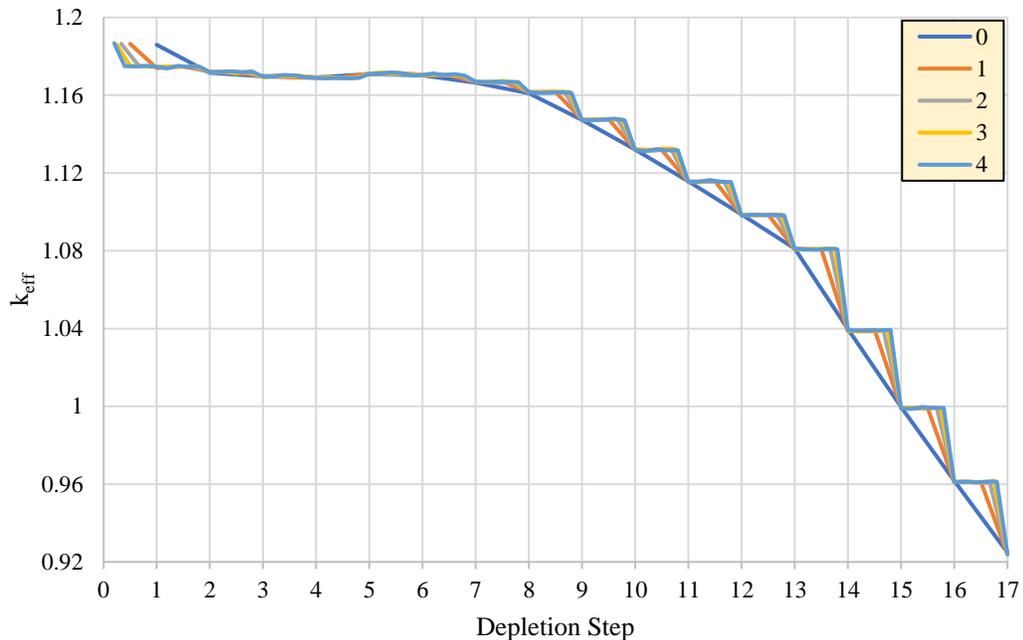


Figure 8.30. Eigenvalue over depletion sequence when using a difference number of thermal hydraulic iterations and equilibrium xenon treatment. Legend values are the number of thermal hydraulic iterations used between burnup steps.

There appears to be good agreement in eigenvalue at each depletion step regardless of the number of TH substeps used. The only noticeable difference is for the first depletion step when using zero iterative substeps since it is conducted using input parameters with no opportunity to first converge on the profile before depletion starts. However, over this large of an eigenvalue range, other differences are difficult to discern. For this reason, reactivity differences relative to the case using four TH iterations per burnup step were calculated and are presented in Figure 8.31.

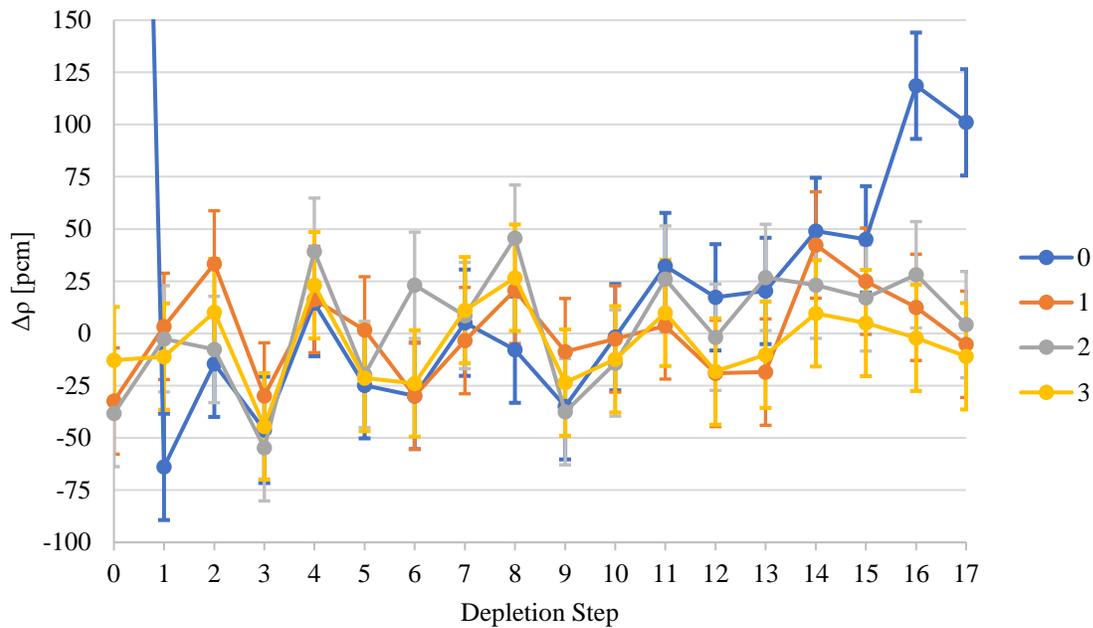


Figure 8.31. Reactivity differences over depletion sequence due to using a different number of thermal hydraulic iterations when using equilibrium xenon treatment. Legend values are the number of thermal hydraulic iterations used between burnup steps, with differences taken relative to the reference case of four iterations.

For one, two, and three TH iterations, there does not appear to be statistically significant differences between the cases considered. These depletion results fall within two standard deviations of the reference case using four TH iterations. The case using zero TH iterations however does appear to show slight differences toward the end of the depletion sequence considered starting at around 40 MWd/kgHM (depletion step 14). The impact is small (maximum observed value of

119 ± 26 pcm at depletion step 16 or 60 MWd/kgHM) but the difference is likely statistically significant.

To better investigate the convergence of the TH profile, the absolute change in local power was considered for each of the 4032 tracked regions between TH substeps. The average of the ratio of change in local power to Serpent-reported local statistical uncertainty was computed and results for each substep are shown in Figure 8.32. For later depletion steps, there are significant differences due to burnup, but TH substeps beyond the first do not meaningfully improve the local power convergence since results comparable to the local statistical uncertainty are quickly reached. The conclusion from these local power and eigenvalue results is that one TH iteration substep is likely sufficient between burnup steps when using a depletion schedule similar to the one considered in this work, since using more substeps did not improve the local power nor the eigenvalue performance relative to their respective reported statistical uncertainties.

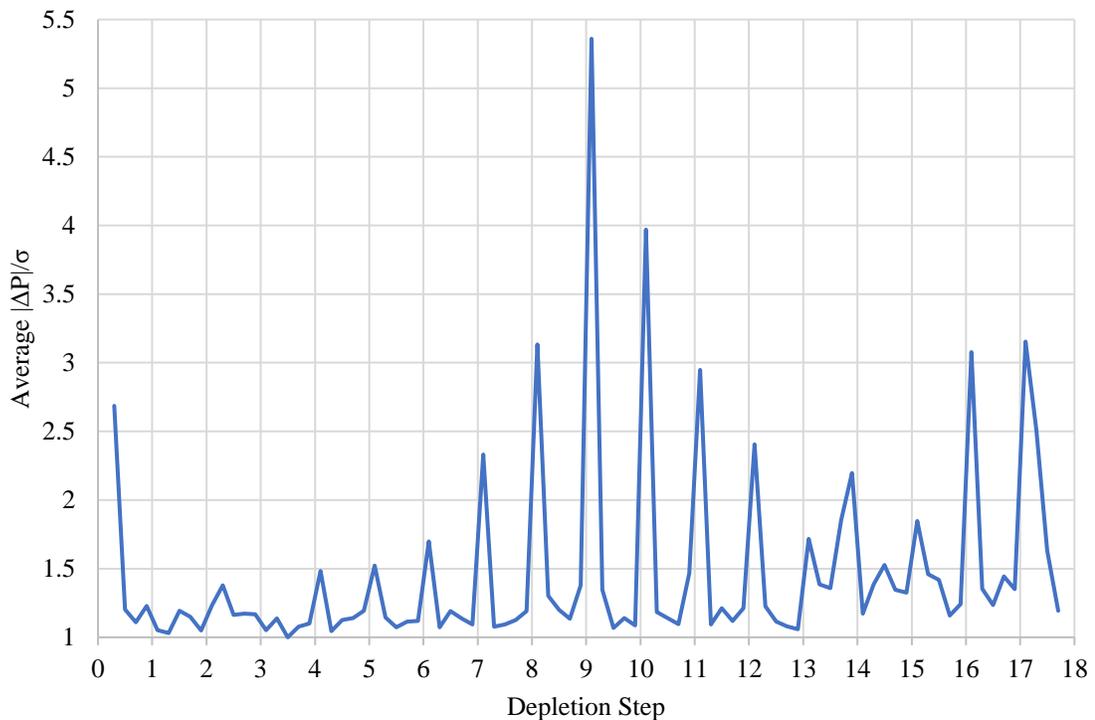


Figure 8.32. Average of 4032 tracked regions' absolute change in local power divided by local power statistical uncertainty.

The simulations run in this section did not use PC method for the depletion steps to test the numerical stability of simulations without it. The axial power profiles over the depletion sequence can be seen in Figure 8.33 for using zero TH iterative substeps. At first glance the profiles appear comparable to those found previously in Figure 8.22, but on closer inspection one can see that there are classical signs of axial numerical instability especially for the final four depletion steps. In Figure 8.34, consider the profiles for 40 (gray line) and 60 (light blue line) MWd/kgHM. Both profiles are peaked toward the bottom. Now consider the profiles for 50 (yellow line) and 70 (lime green line) MWd/kgHM. Both profiles are peaked toward the top. Alternating axial shifts in power over cycle such as these are clear indicators for numerical instability, although not so large in magnitude to discredit the results.

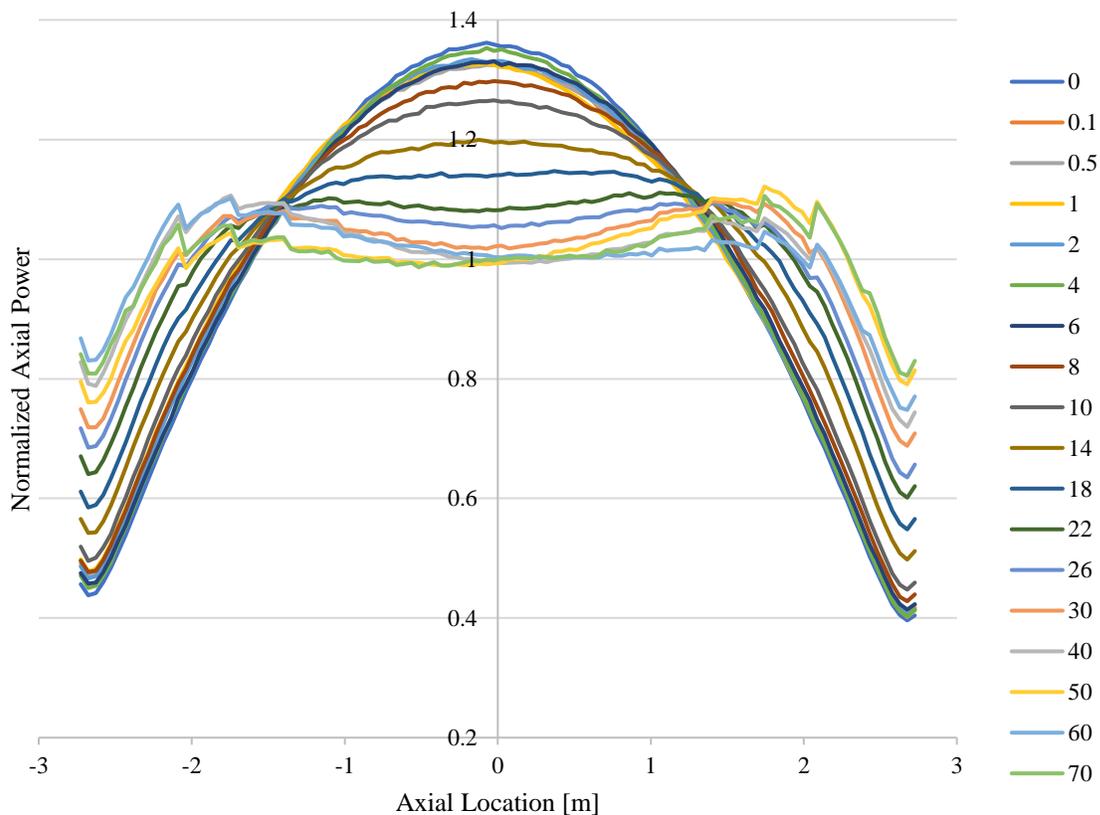


Figure 8.33. Normalized axial power profiles for thermal hydraulic feedback and zero iteration substeps. Legend corresponds to cumulative burnup amounts given in units of MWd/kgHM for the seventeen depletion steps considered.

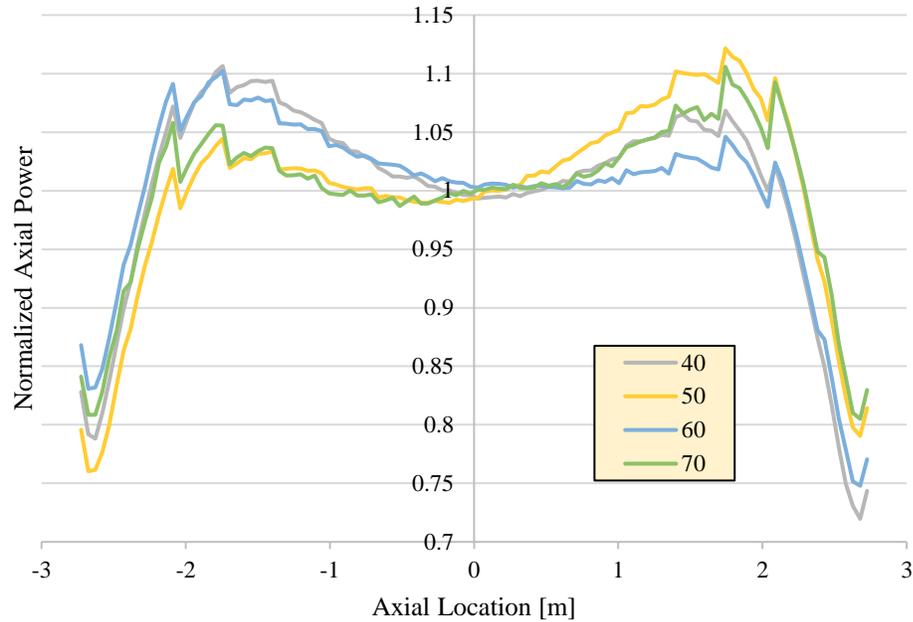


Figure 8.34. Normalized axial power profiles for final four depletion steps. Legend corresponds to cumulative burnup amounts give in units of MWd/kgHM.

Observing the axial numerical instabilities is more straightforward for an integral parameter such as AO. The AO results for each TH iteration case are shown in Figure 8.35. A peculiar observation is that for higher TH iteration substeps, the numerical instabilities appear to begin sooner. This might be due to relatively high statistical uncertainties in local power production. Regardless of the cause, it appears the unstable behavior for four TH substeps begins at the ninth depletion step, at the twelfth depletion step for both two and three TH substeps, and at the thirteenth depletion step for both zero and one TH substeps. More iterations cause the observed axial oscillations to have an earlier onset and is indicative that TH search method is introducing more instability than not using it. This might be due to the fact the TH iterations use fewer total transport particles than depletion steps. TH substeps obtain more poorly converged results and successive iterations compounds the issue. For this reason, it could be beneficial to not use fewer particles during the substep calculations. Instabilities would be expected from depletion of a large reactor without using PC or other instability mitigating methods. As will be shown by means of comparison to studies to follow, these numerical instabilities are attributable to lack of using the

PC method during the depletion steps. Thus, it is suggested to at least use the PC method (or a better one if available, as discussed previously with recent additions to Serpent with the release of version 2.32) when running a full-core depletion simulation of the large AHTR core. The onset of these observed instabilities is likely due to AHTR's size and resulting loose coupling between distant fissile zones within the core.

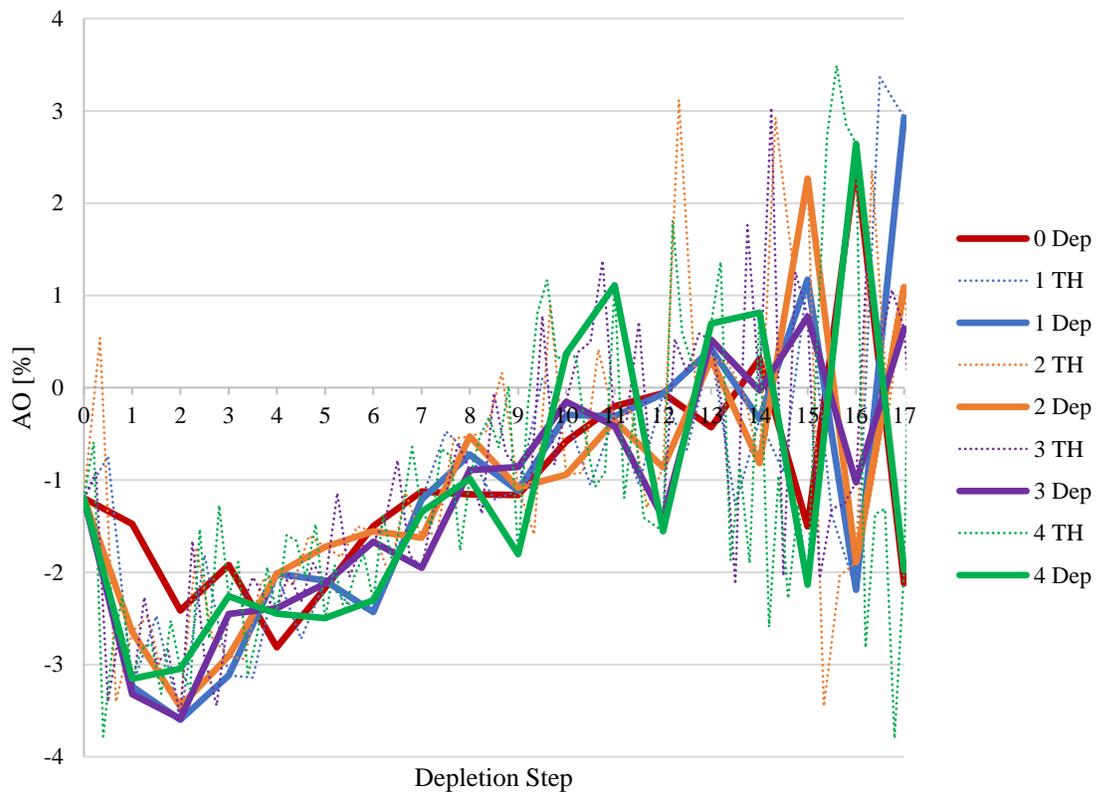


Figure 8.35. AO behavior over depletion sequence when using a different number of thermal hydraulic substeps. Cases used equilibrium xenon treatment but no predictor-corrector method.

8.3.4.b Explicit Xenon Treatment

The results in this section detail simulations run using PC, explicit xenon treatment, and one TH substep between depletion steps. The results are notably more stable than those of the previous study due to the use of PC. The axial power profile for each depletion step can be seen in Figure 8.36. Results are similar to those in the case without TH feedback (Figure 8.22) where the

power initially has the classical cosine shape peaked toward the bottom of core, flattening out over the cycle, and eventually shifting toward peaking toward the top of the core at EOC. Significant axially alternating behaviors are not observed, so the results do not appear to suffer from numerical instabilities.

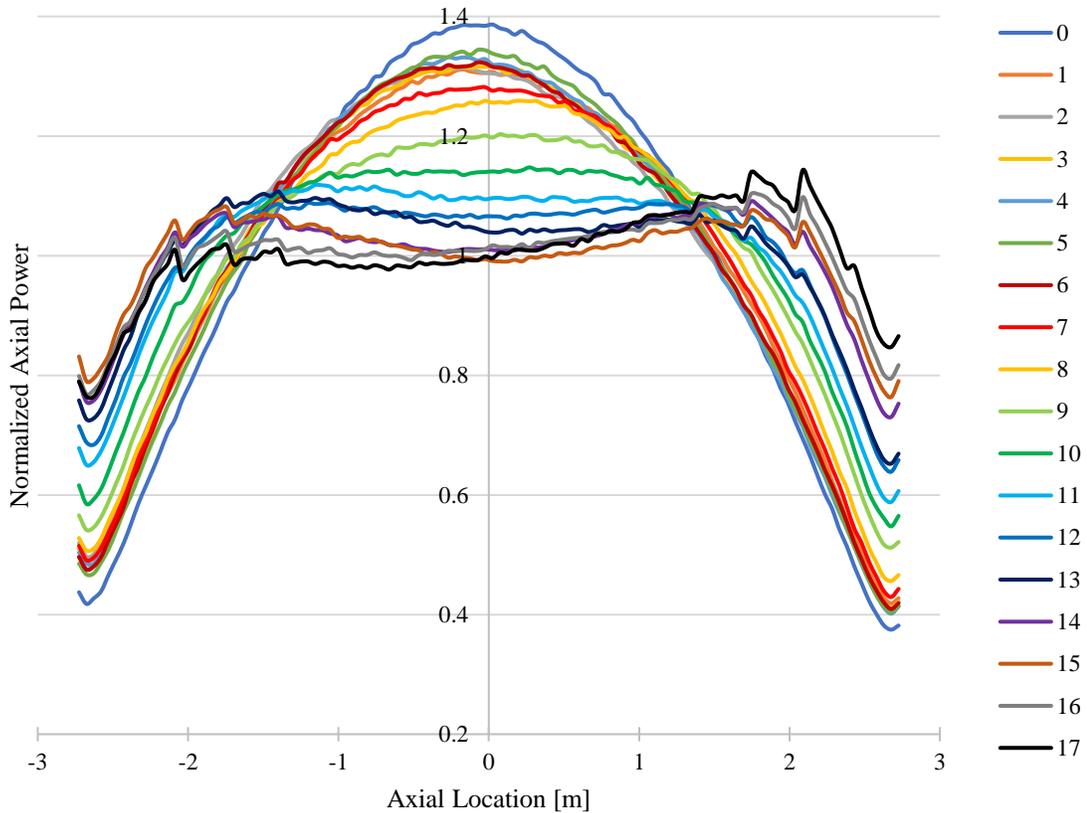


Figure 8.36. Normalized axial power profiles over depletion sequence using one thermal hydraulic substep and explicit xenon treatment. Legend values correspond to the burnup step number.

The PPFs for the whole core, one-third assembly section, assembly, and axial levels of resolution can be found in Figure 8.37. One can see the immediate impact TH feedback has on radial peaking by the significant drop in PPF for the TH substep conducted before the first depletion step. After that point, the PPFs decrease similarly to the performance observed for no TH feedback (Figure 8.25).

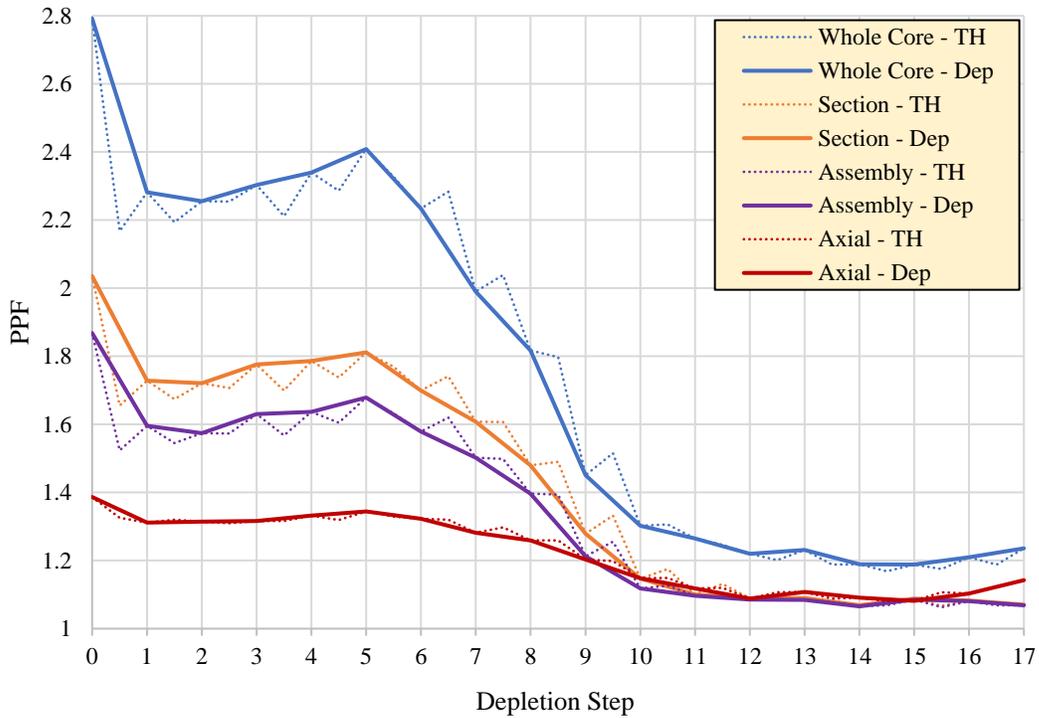


Figure 8.37. PPF values for whole core (4032 regions), section (252), assembly (84), and axial levels (112) over the depletion sequence when using one thermal hydraulic substep, explicit xenon treatment, and predictor-corrector.

A clearer way to discern the axial numerical stability of the simulation is to consider the AO performance over the depletion sequence. This is shown in Figure 8.38 and compared with the results from the case with no TH feedback (Figure 8.24 and repeated as the red line in Figure 8.38). One can see that TH feedback initially shifts the power to the lower core half than for the case without TH feedback, which is attributable to the negative temperature coefficient for the AHTR system. Over the depletion sequence, AO shifts higher in core and becomes positive near EOC.

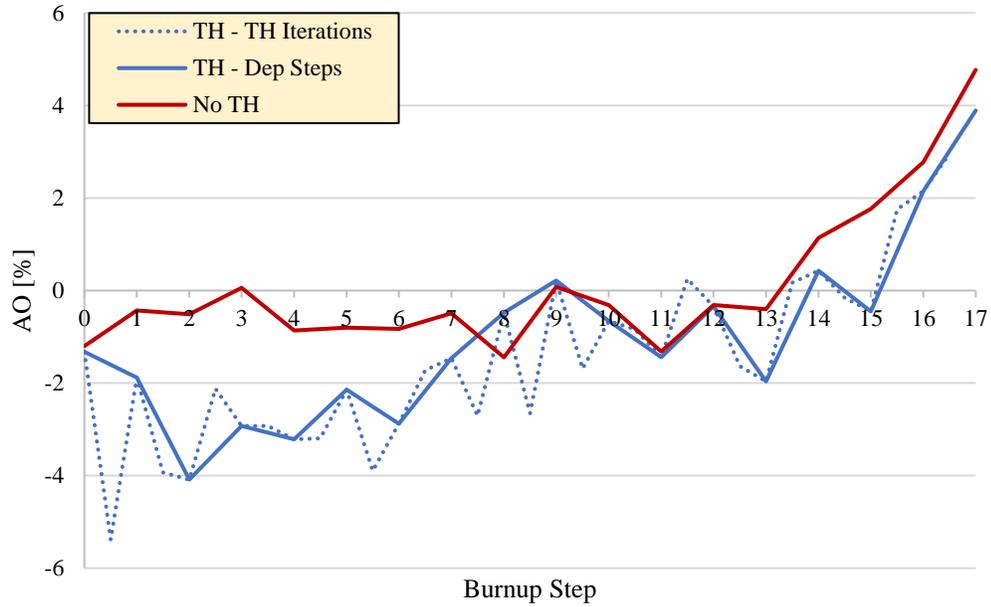


Figure 8.38. AO of the power profile at each depletion step using one thermal hydraulic substep and explicit xenon treatment.

8.3.4.c Equilibrium Xenon Treatment

Analogous to the explicit xenon treatment case in the previous subsection, this subsection also considers a depletion sequence run using PC and one TH substep but with equilibrium xenon treatment instead. The results are comparable between the two xenon treatments and differences are addressed in the proceeding subsection. Axial power profiles over the depletion sequence can be found in Figure 8.39.

Spatial PPF results for equilibrium xenon treatment are presented in Figure 8.40. Like their explicit counterparts, the radial PPF drops at BOC due to TH feedback. It increases slightly until the fifth depletion step (4 MWd/kgHM or about 20.5 EFPD), and then decreases over the remainder of the depletion sequence.

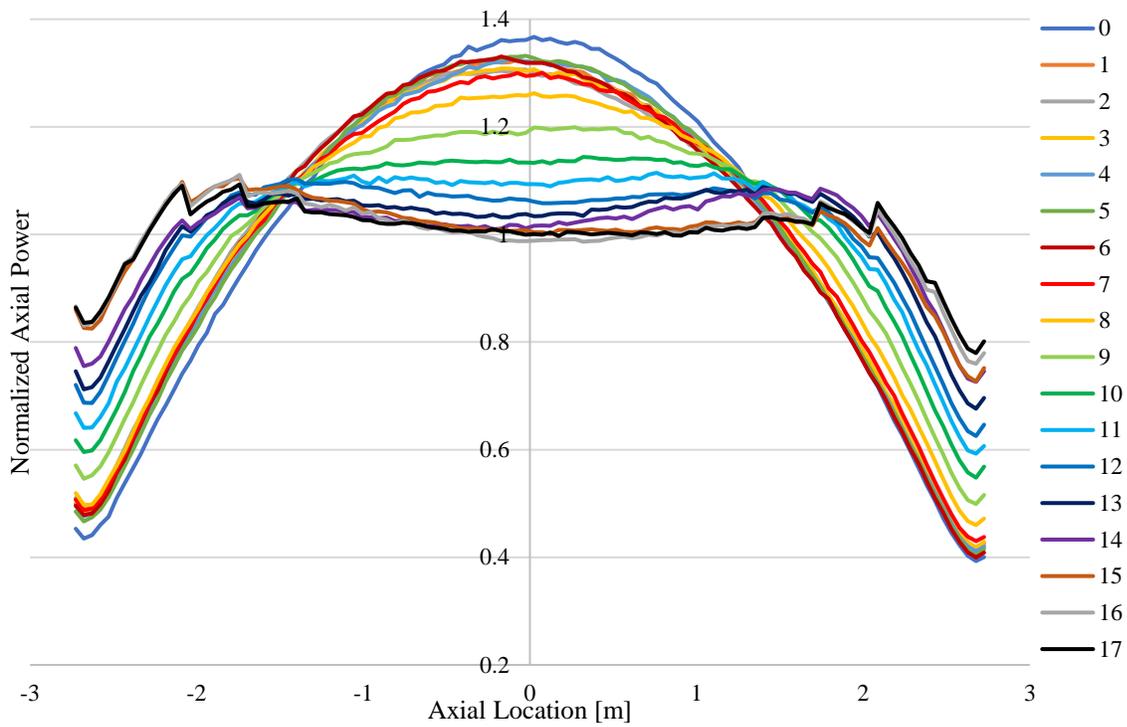


Figure 8.39. Normalized axial power profiles over depletion sequence using one thermal hydraulic substep and equilibrium xenon treatment. Legend values correspond to the burnup step number.

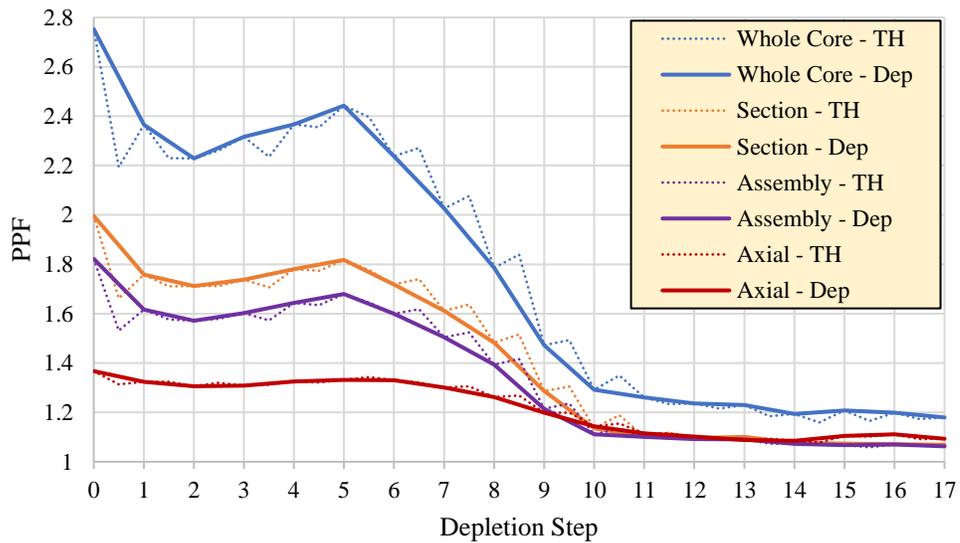


Figure 8.40. PPF values for whole core, section, assembly, and axial levels over the depletion sequence when using one thermal hydraulic substep, equilibrium xenon treatment, and predictor-corrector.

The AO obtained when using equilibrium xenon treatment can be found in Figure 8.41. Unlike their explicit xenon treatment counterpart, these results appear to have slight numerical stability issues near EOC. The final three depletion steps report a negative AO whereas all over considerations predict a positive result (both for results previously presented and those still yet to come). This suggests that accurate axial power convergence was challenged during this depletion study should be noted for similar depletion studies.

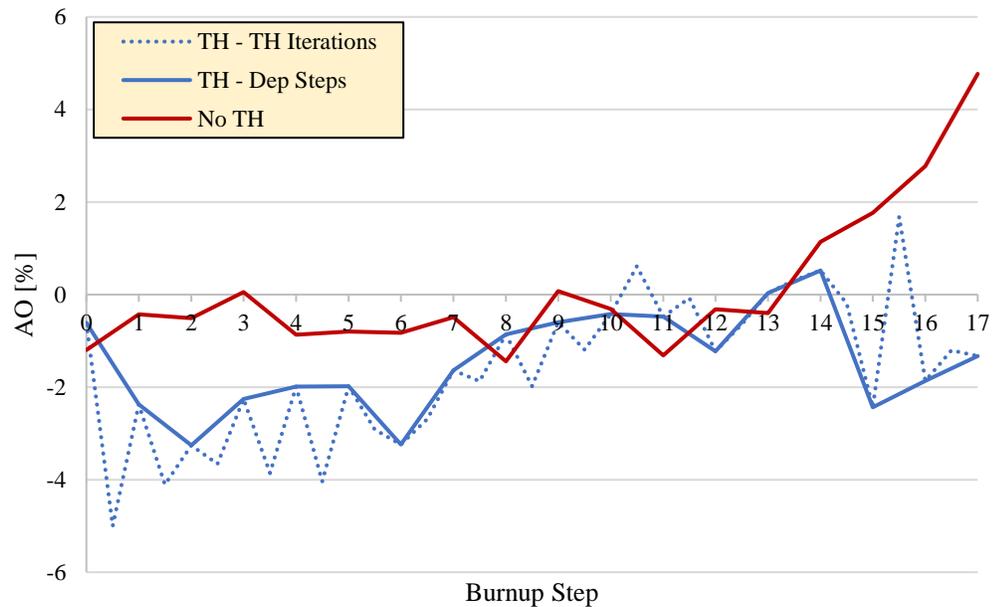


Figure 8.41. AO of the power profile at each depletion step using one thermal hydraulic substep and equilibrium xenon treatment.

8.3.4.d Xenon Treatment Comparison

Results for both the explicit and equilibrium xenon treatments are compared in this section. Overall, agreement is good with the largest difference occurring in the AO near EOC. Eigenvalue differences over the depletion sequence are summarized in Table 8.21. BOC differences are large as expected as the explicit treatment builds-in to the equilibrium concentration, but quickly drop by 0.5 MWd/kgHM (about 2.5 EFPD).

Table 8.21. Eigenvalue results of different xenon treatments when using one thermal hydraulic substep.

Burnup [MWd/kgHM]	Time [d]	Explicit Xenon		Equilibrium Xenon		Δk [pcm]	$\Delta \rho$ [pcm]
		k_{eff}	σ [pcm]	k_{eff}	σ [pcm]		
0	0	1.21646	18	1.18122	17	-3524	-2940
0.1	0.513	1.18737	17	1.18144	17	-593	-501
0.5	2.56	1.17107	17	1.17835	18	728	620
1	5.12	1.16850	17	1.17623	19	773	659
2	10.2	1.16801	18	1.17566	18	765	653
4	20.5	1.17071	17	1.17835	18	764	650
6	30.7	1.17053	17	1.17789	17	736	627
8	41	1.16779	18	1.17498	19	719	614
10	51.2	1.16326	17	1.16968	18	642	550
14	71.7	1.14940	18	1.15546	17	606	526
18	92.1	1.13445	18	1.13927	18	482	424
22	113	1.11725	18	1.12174	19	449	401
26	133	1.09916	18	1.10423	18	507	460
30	154	1.08179	18	1.08600	18	421	388
40	205	1.03720	17	1.04137	18	417	401
50	256	0.99557	18	0.99852	19	295	296
60	307	0.95539	18	0.95821	18	282	295
70	358	0.91619	17	0.91821	18	202	220

A plot of the reactivity differences over the depletion sequence between the two cases can be seen in Figure 8.42. The initial (large negative) difference between the two cases is not shown since it is fully expected and to better show the behavior over the remainder of the depletion sequence. Once the explicit xenon treatment reaches the equilibrium concentration, reactivity differences are quite high (about 650 ± 21 pcm). The reactivity differences decrease over the depletion sequence, down to 220 ± 27 pcm at 70 MWd/kgHM. Differences between the two xenon treatments are larger for TH feedback than previously observed without TH feedback (Table 8.19 and Figure 8.27). The reason for this difference is not known but the user should be aware of the discrepancy between the two sets.

Relative differences in various spatial power parameters between the two xenon treatments are shown in Figure 8.43. After consistent differences for the first three depletion steps, the differences over the remaining depletion steps do not seem to follow any discernable trends. Therefore, it is likely that the observed relative differences are driven by statistical uncertainty rather than some systematic trend beyond the first few steps near BOC.

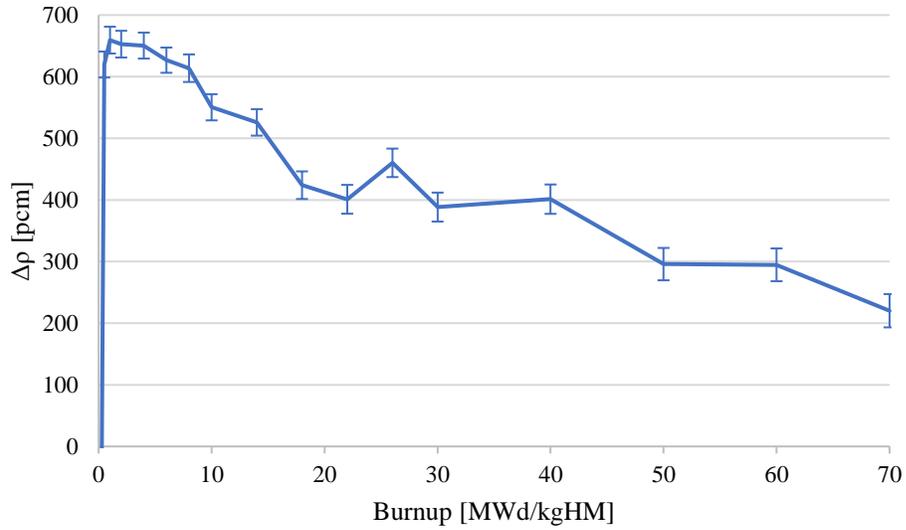


Figure 8.42. Reactivity differences between xenon treatments when using one thermal hydraulic substep.

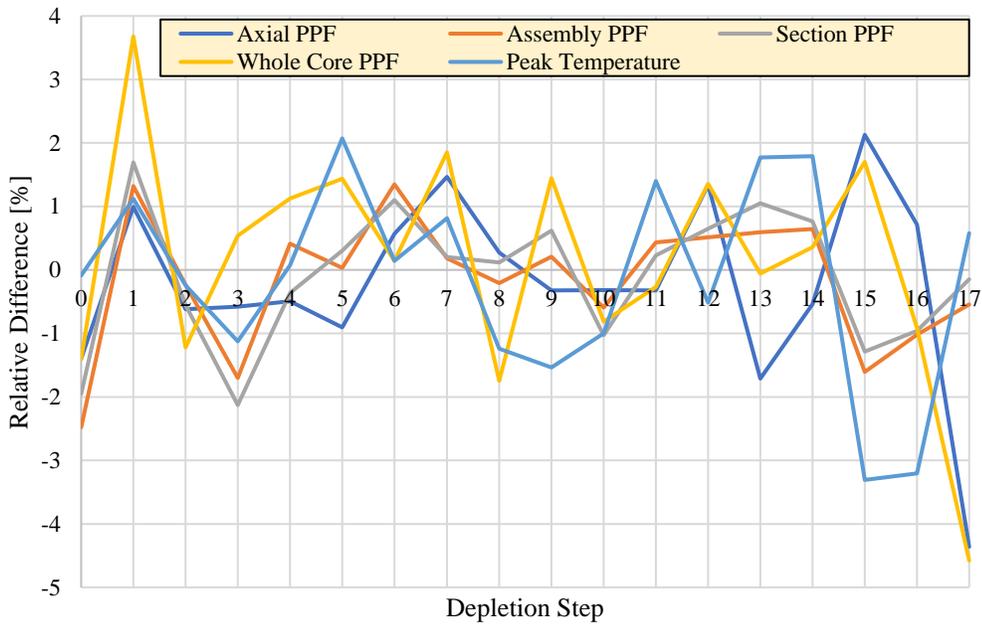


Figure 8.43. Relative (to explicit xenon treatment) differences between xenon treatments for various spatial power parameters when using one thermal hydraulic substep.

A summary of the average relative differences and average absolute relative differences for each parameter shown in Figure 8.43 is given in Table 8.22. Low average relative differences suggest that the differences over the depletion sequence observed in Figure 8.43 can likely be

attributed to statistical uncertainty in the simulations since there are no strong systematic trends. More significant average absolute relative differences suggest that the actual statistical uncertainties of these parameters are likely higher than those expected by the reported statistical uncertainties estimated in Table 8.1.

Table 8.22. Relative (to explicit xenon treatment) average differences summary for various spatial power parameters when using one thermal hydraulic substep.

Parameter	Average Relative Difference [%]	Average Absolute Relative Difference [%]
Axial PPF	-0.23	1.06
Assembly PPF	-0.15	0.78
Section PPF	-0.09	0.84
Whole Core PPF	0.15	1.37
Peak Temperature	-0.14	1.22

A comparison of the AO behavior over the depletion sequence for the two xenon treatments can be seen in Figure 8.44. The results when using no TH feedback are also presented for comparison (red line). Note that initially, the two xenon treatments agree quite well. However, toward EOC, the equilibrium xenon treatment seems to predict the axial power shape incorrectly by shifting the power toward the bottom of the core instead of toward the top. The cause of this behavior seen in the equilibrium xenon case is not known but it is the only case disagreeing with five other cases using fine spatial tracking which predict that AO becomes positive at EOC.

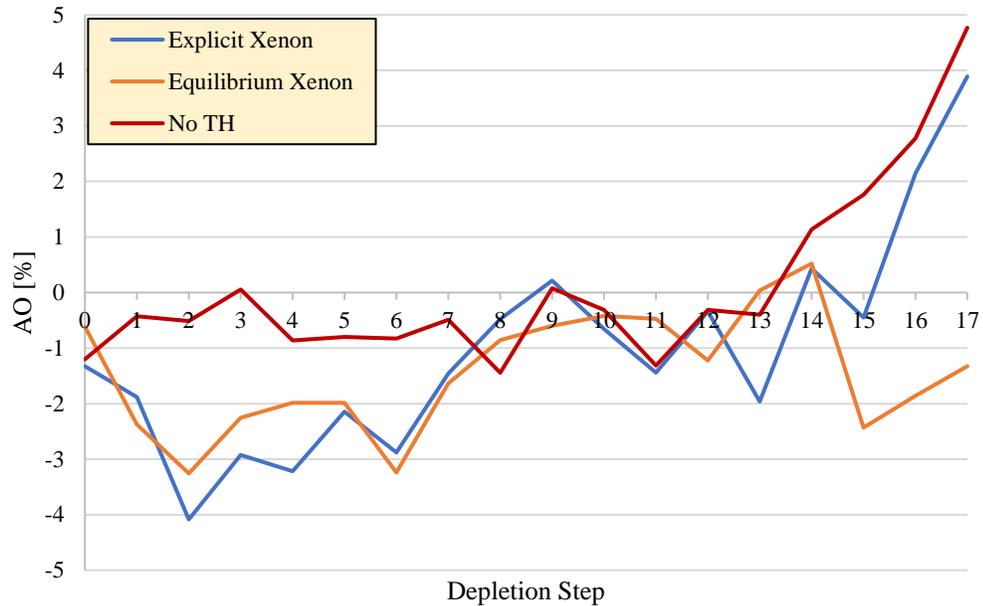


Figure 8.44. AO of the power profile at each depletion step using one thermal hydraulic substep.

8.3.5 Fine Spatial Tracking with Criticality Control via Control Blade Movement

This section details the results from two simulations run using the same fine spatial resolution detailed in the previous study (4032 total regions) with the addition of active criticality control. As detailed in Chapter 6, active criticality control is achieved by moving CBs in the core according to a user-prescribed schedule. The two simulations run here each use the xenon treatments considered previously: explicit and equilibrium tracking. Each simulation uses criticality search substeps between depletion steps to search for the critical insertion at the beginning of the step, within a user-prescribed tolerance. Both cases assumed an eigenvalue tolerance of ± 150 pcm about a target eigenvalue of 1.00000. Each also used 100,000 particles per generation, 100 inactive cycles, and the number of active cycles depending on the case: 200 active cycles for criticality search and 500 active cycles for depletion transport. Fewer cycles were used for the criticality search substeps since those calculations were only interested in a single integral parameter (eigenvalue) whereas the depletion transport calculations needed to resolve local power production at the 1/4032 core level. The explicit xenon tracking case required 45 total transport

calculations (28 critical search, 17 depletion) and finished in 392.3 h. The equilibrium xenon tracking case required 46 total transport calculations (29 criticality search, 17 depletion) and finished in 400.3 h. Results for both simulations are detailed below with comparisons following.

8.3.5.a Explicit Xenon Treatment

The results for the depletion case using explicit xenon treatment can be found in Table 8.23. Note that some depletion steps require multiple CB movement substeps while others require none. The number is dependent upon how many searches were required for the beginning of burnup step critical configuration to be found. For a few of the initial depletion steps where reactivity changes only slightly due to comparable reactivity losses due to fuel depletion and reactivity gains due to BP depletion, no CB movement substeps are required at all as the previous step's critical insertion configuration is still within the eigenvalue tolerance (± 150 pcm for this case).

The eigenvalue over criticality search substeps and depletion steps can be seen in Figure 8.45. Note that the criticality search algorithm is discernable by looking between depletion step points. After a depletion step, if the resulting configuration is not within the eigenvalue tolerance of the target value, a single CB group is withdrawn from the core. If that configuration is not within the tolerance either, the next guess at a critical configuration is extrapolated from the previous two configurations and their respective eigenvalues. This is especially evident in the second half of depletion steps considered (steps 10 through 16). For the final depletion step, all CBs are withdrawn due to being subcritical, which explains why the eigenvalue does not reach the eigenvalue tolerance before conducting the depletion transport calculation.

Table 8.23. Depletion sequence summary using CB movement and explicit xenon treatment.

Depletion Step	CB Movement Step	k_{eff}	Assembly Groups Inserted	Assembly PPF	Section PPF	Axial PPF	AO [%]	Core PPF
0	1	1.00153	57	1.418	1.457	1.398	-2.32	2.032
0	2	1.00494	56	1.358	1.396	1.391	-2.68	1.926
0	3	0.99982	57.4375	1.366	1.418	1.425	-8.17	2.021
1	1	0.97717	57.4375	1.358	1.401	1.417	-8.06	1.916
1	2	0.98087	56.4375	1.351	1.378	1.383	-6.24	1.933
1	3	0.99837	51.25	1.316	1.345	1.376	-4.79	1.920
1	4	0.99869	50.75	1.315	1.390	1.342	-3.26	1.906
2	1	0.98597	50.75	1.289	1.361	1.353	-2.49	1.831
2	2	0.98940	49.75	1.239	1.352	1.361	-5.21	1.848
2	3	1.00068	46.6875	1.537	1.638	1.350	-5.48	2.214
3	1	0.99912	46.6875	1.437	1.524	1.349	-5.47	2.069
4	1	0.99885	46.6875	1.431	1.516	1.347	-5.25	2.080
5	1	1.00141	46.6875	1.416	1.499	1.345	-5.25	2.027
6	1	1.00247	46.6875	1.412	1.486	1.370	-6.26	1.981
6	2	1.00645	45.6875	1.393	1.448	1.358	-4.38	1.969
6	3	1.00037	47.3125	1.316	1.342	1.378	-4.69	1.981
7	1	0.99979	47.3125	1.303	1.323	1.388	-5.30	1.853
8	1	0.99639	47.3125	1.286	1.297	1.352	-4.11	1.743
8	2	1.00049	46.3125	1.367	1.422	1.343	-5.58	1.981
9	1	0.98983	46.3125	1.255	1.297	1.325	-4.42	1.638
9	2	0.99432	45.3125	1.418	1.461	1.256	-5.63	2.099
9	3	0.99852	44.0625	1.514	1.552	1.225	0.55	1.920
10	1	0.98501	44.0625	1.408	1.444	1.232	1.68	1.699
10	2	0.98839	43.0625	1.413	1.422	1.183	1.38	1.756
10	3	1.00000	39.625	1.400	1.435	1.166	-3.46	1.705
11	1	0.98473	39.625	1.294	1.328	1.170	-4.15	1.520
11	2	0.98794	38.625	1.495	1.509	1.127	-2.13	1.824
11	3	1.00044	34.875	1.358	1.374	1.183	4.47	1.632
12	1	0.98451	34.875	1.280	1.293	1.178	4.28	1.456
12	2	0.98808	33.875	1.520	1.549	1.110	-0.09	1.773
12	3	1.00033	30.5625	1.372	1.381	1.115	-1.18	1.570
13	1	0.98300	30.5625	1.296	1.301	1.125	-3.00	1.460
13	2	0.98627	29.5625	1.444	1.461	1.123	1.04	1.770
13	3	1.00033	25.375	1.313	1.325	1.129	-3.67	1.623
14	1	0.95855	25.375	1.188	1.194	1.100	-1.75	1.366
14	2	0.96241	24.375	1.335	1.342	1.123	-2.66	1.768
14	3	0.99394	14.6875	1.381	1.412	1.119	-1.73	1.520
14	4	1.00029	12.8125	1.281	1.297	1.154	3.97	1.522
15	1	0.95946	12.8125	1.203	1.206	1.128	2.22	1.333
15	2	0.96334	11.8125	1.362	1.411	1.095	-0.70	1.535
15	3	0.99769	2.375	1.350	1.371	1.123	-4.10	1.635
15	4	1.00037	1.75	1.421	1.444	1.096	0.81	1.598
16	1	0.95835	1.75	1.262	1.266	1.087	0.63	1.352
16	2	0.96217	0.75	1.394	1.402	1.118	-3.22	1.946
17	1	0.92461	0	1.341	1.353	1.133	4.69	1.487

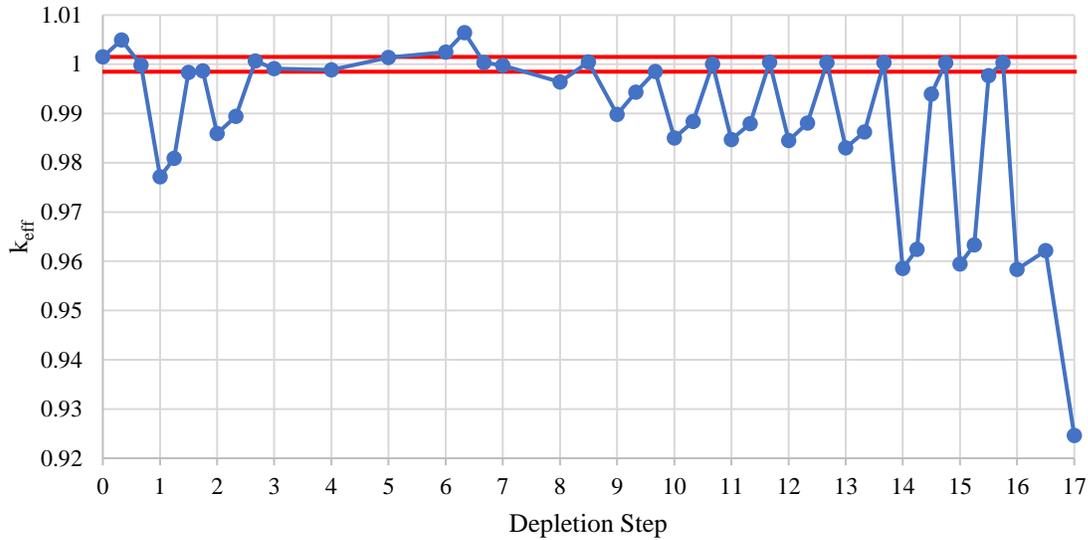


Figure 8.45. Eigenvalue over depletion sequence using criticality iteration via CB movement and explicit xenon treatment. Red lines correspond to the criticality tolerance for the criticality search (± 150 pcm).

A plot of the PPF performance over the depletion sequence can be seen in Figure 8.46. Results for the depletion steps are shown with solid lines and corresponding results for criticality searches are shown with dashed lines. The depletion results appear to be stable for the depletion steps considered. One can see that for the later depletion steps which have larger larger burnup per step, the radial results (orange and blue for section- and assembly-level averaging, respectively) jump considerably during the criticality search procedure (dashed lines) but settle to a lower value once criticality is found for the depletion step (solid lines). The axial results show less variation over the depletion sequence, starting with an initial value around 1.4 and decreasing to about 1.15. Criticality iteration substeps do not appear to significantly impact the axial power distribution as much as the radial distribution during the critical configuration search process.

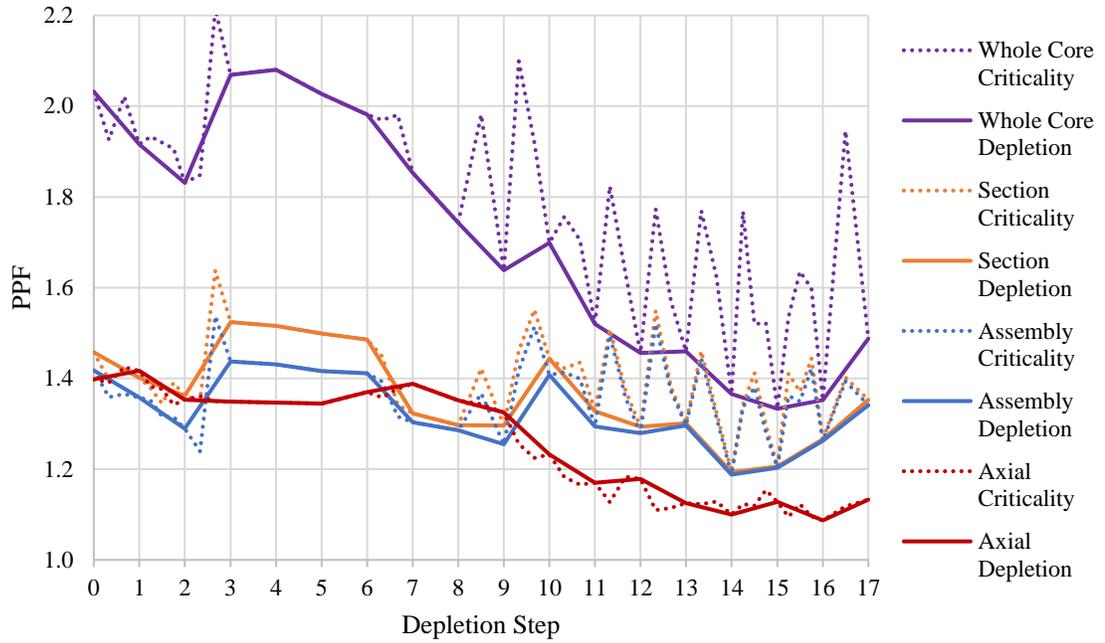


Figure 8.46. PPFs over depletion sequence using criticality iteration via CB movement and explicit xenon treatment.

The axial power profiles for each of the seventeen depletion steps can be seen in Figure 8.47. The results appear to be stable and comparable to those found without criticality consideration for fine spatial discretization (Figure 8.23). Axial power near BOC has a classical cosine shape peaked toward the bottom of the core. Over the depletion sequence, the profile flattens and eventually becomes peaked toward the top of the core for the last few considered depletion steps.

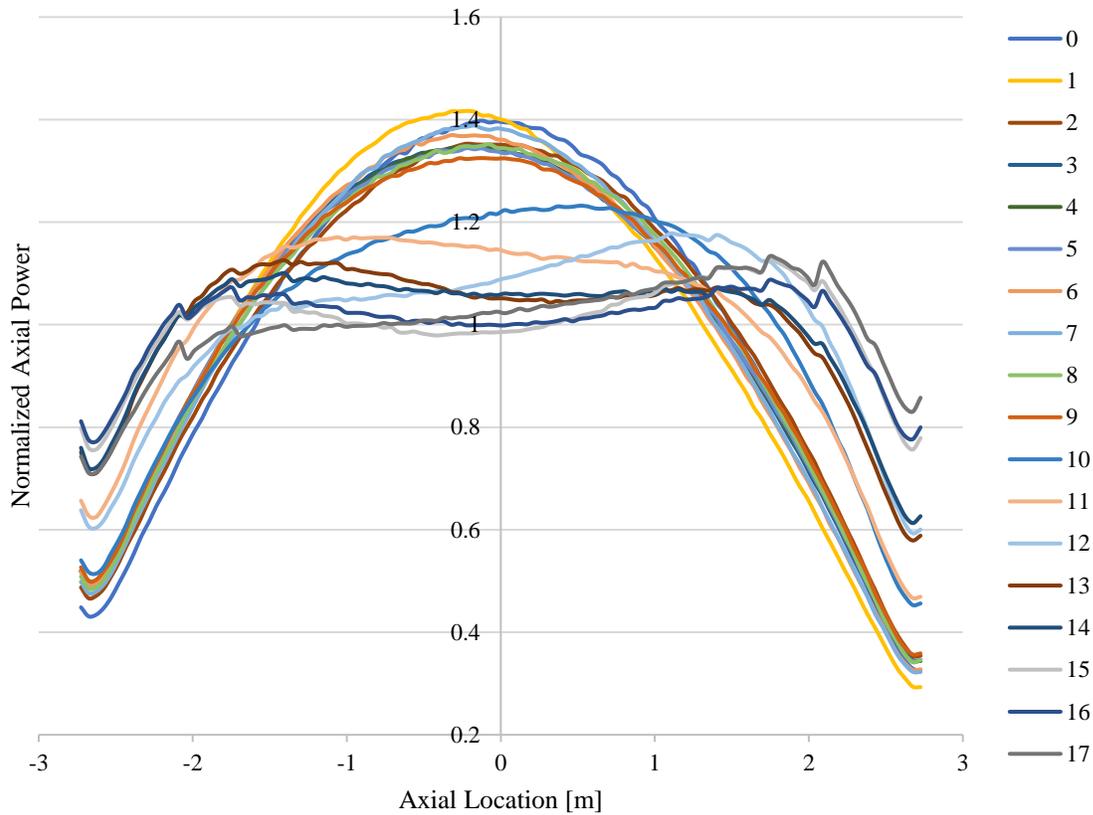


Figure 8.47. Normalized axial power profiles over depletion sequence using criticality iteration via CB movement and explicit xenon treatment. Legend values are for each depletion step.

The AO performance over the depletion sequence can be seen in Figure 8.48. It is initially negative, but then becomes more oscillatory for later depletion steps. One might immediately suspect that numerical instabilities are present here, but this is not necessarily the case. For reference, the partial insertion fraction of the last CB group inserted into the core is shown in orange. For partial insertions near the top or bottom of the core, it is similar to if the group was fully withdrawn or inserted, respectively. For these cases, there is no significant axial peaking from the CBs since the axial CB configuration is more symmetric. For partial insertions near the middle of the core, there is a significant axial shift toward the bottom of the core due to the asymmetric axial loading of the CBs. This justifies the AO behavior observed, especially for the later depletion steps where the power would want to shift toward the top of the core, as was the case in Figure 8.24. Partial insertions near the middle of the core push the power toward the bottom, whereas

partial insertions near the top and bottom of the core have positive AO behaviors which would be expected of a symmetrically-controlled system at these later depletion steps.

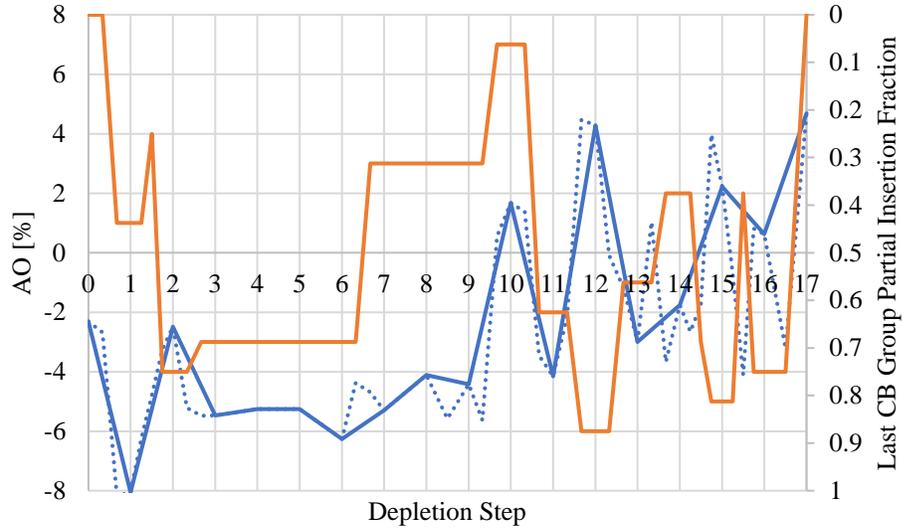


Figure 8.48. AO (left axis) over depletion sequence (solid blue) with criticality iteration substeps (dashed blue) via CB movement and explicit xenon treatment. Partial insertion fraction (right axis) of the last CB group (orange) is shown for reference for impact on the AO.

8.3.5.b Equilibrium Xenon Treatment

The discussion for equilibrium xenon treatment is very similar to that of explicit xenon treatment. Detailed discussion for each result presented here can be found in the previous subsection, so this one is kept shorter. Results from the simulation case using criticality iteration via CB movement and equilibrium xenon treatment can be found in Table 8.24. The eigenvalue over the depletion sequence when using criticality iteration via CB movement and equilibrium xenon treatment can be found in Figure 8.49. The PPFs at the whole core, section, assembly, and axial levels over the depletion sequence when using criticality iteration via CB movement and equilibrium xenon treatment can be found in Figure 8.50. The normalized axial power profiles for each depletion step for the case of using criticality iteration via CB movement and equilibrium xenon treatment can be found in Figure 8.51.

Table 8.24. Depletion sequence summary using CB movement and equilibrium xenon treatment.

Depletion Step	CB Movement Step	k_{eff}	Assembly Groups Inserted	Assembly PPF	Section PPF	Axial PPF	AO [%]	Core PPF
0	1	1.00322	47	1.424	1.510	1.368	-2.41	2.097
0	2	1.00731	46	1.510	1.595	1.361	-3.32	2.148
0	3	0.99985	47.8125	1.379	1.393	1.368	-2.51	1.955
1	1	0.99992	47.8125	1.395	1.413	1.366	-2.32	1.934
2	1	0.99713	47.8125	1.388	1.406	1.349	-4.36	1.909
2	2	1.00048	46.8125	1.502	1.595	1.353	-2.73	2.189
3	1	0.99862	46.8125	1.448	1.537	1.353	-3.45	2.099
4	1	0.99863	46.8125	1.420	1.503	1.345	-3.18	2.049
5	1	1.00124	46.8125	1.432	1.511	1.359	-3.83	2.051
6	1	1.00258	46.8125	1.384	1.458	1.368	-4.55	1.961
6	2	1.00664	45.8125	1.415	1.475	1.364	-3.42	2.013
6	3	1.00069	47.4375	1.326	1.348	1.376	-6.13	1.889
7	1	0.99970	47.4375	1.316	1.331	1.383	-6.56	1.803
8	1	0.99663	47.4375	1.301	1.318	1.351	-5.90	1.704
8	2	1.00023	46.4375	1.380	1.447	1.327	-6.46	1.896
9	1	0.98948	46.4375	1.293	1.338	1.326	-6.85	1.629
9	2	0.99368	45.4375	1.289	1.337	1.251	-4.17	1.949
9	3	0.99910	43.9375	1.455	1.492	1.239	0.94	1.907
10	1	0.98601	43.9375	1.390	1.424	1.228	-0.67	1.660
10	2	0.98856	42.9375	1.375	1.385	1.169	-0.33	1.680
10	3	1.00439	38.4375	1.414	1.442	1.191	-4.55	1.761
10	4	1.00010	39.6875	1.314	1.359	1.152	-2.25	1.634
11	1	0.98468	39.6875	1.262	1.301	1.161	-2.71	1.502
11	2	0.98847	38.6875	1.485	1.507	1.113	-1.02	1.812
11	3	0.99835	35.625	1.302	1.332	1.132	0.08	1.577
11	4	1.00123	35.125	1.296	1.315	1.169	2.45	1.556
12	1	0.98418	35.125	1.252	1.274	1.156	2.36	1.450
12	2	0.98720	34.125	1.419	1.446	1.121	0.31	1.682
12	3	1.00216	29.875	1.446	1.451	1.161	3.88	1.712
12	4	0.99990	30.5	1.343	1.348	1.134	-3.60	1.565
13	1	0.98384	30.5	1.307	1.314	1.108	-1.46	1.490
13	2	0.98666	29.5	1.374	1.386	1.117	1.01	1.662
13	3	1.00220	24.75	1.429	1.462	1.116	0.61	1.686
13	4	0.99985	25.4375	1.272	1.290	1.105	-1.86	1.587
14	1	0.95881	25.4375	1.208	1.212	1.097	-0.92	1.358
14	2	0.96240	24.4375	1.261	1.287	1.122	2.29	1.620
14	3	0.99647	14	1.292	1.335	1.188	6.86	1.543
14	4	1.00013	12.9375	1.326	1.340	1.213	8.53	1.599
15	1	0.95946	12.9375	1.217	1.226	1.185	6.92	1.407
15	2	0.96259	11.9375	1.373	1.414	1.184	7.21	1.672
15	3	1.00978	0	2.094	2.178	1.187	7.89	2.522
15	4	0.99741	2.5	1.329	1.348	1.092	0.43	1.510
15	5	0.99926	2	1.401	1.414	1.216	10.01	1.712
16	1	0.95772	2	1.294	1.300	1.168	6.22	1.478
16	2	0.96117	1	1.377	1.402	1.132	4.55	1.587
17	1	0.92430	0	1.340	1.355	1.147	5.14	1.448

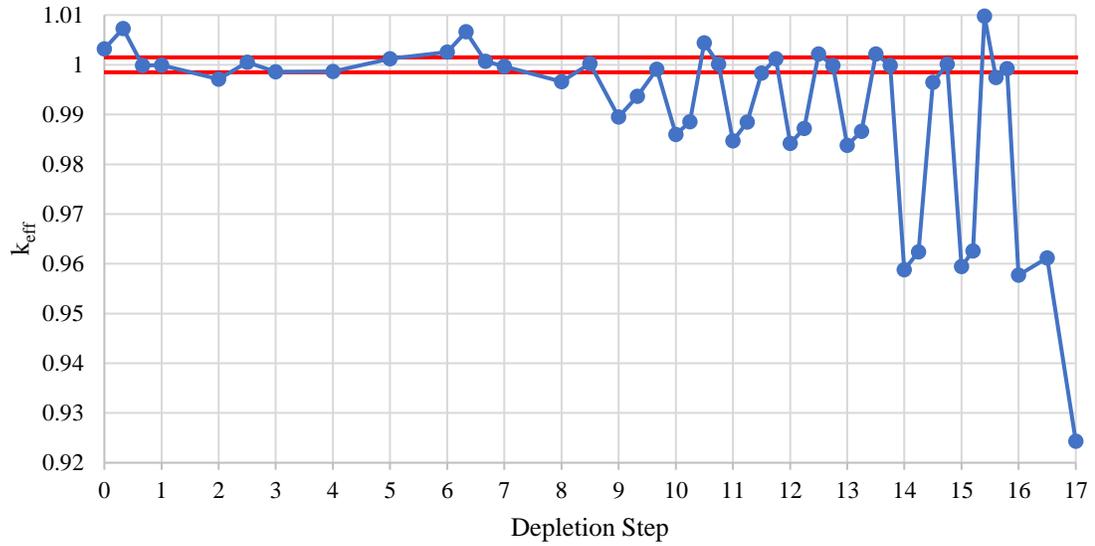


Figure 8.49. Eigenvalue over depletion sequence using criticality iteration via CB movement and equilibrium xenon treatment. Red lines correspond to the criticality tolerance for the criticality search (± 150 pcm).

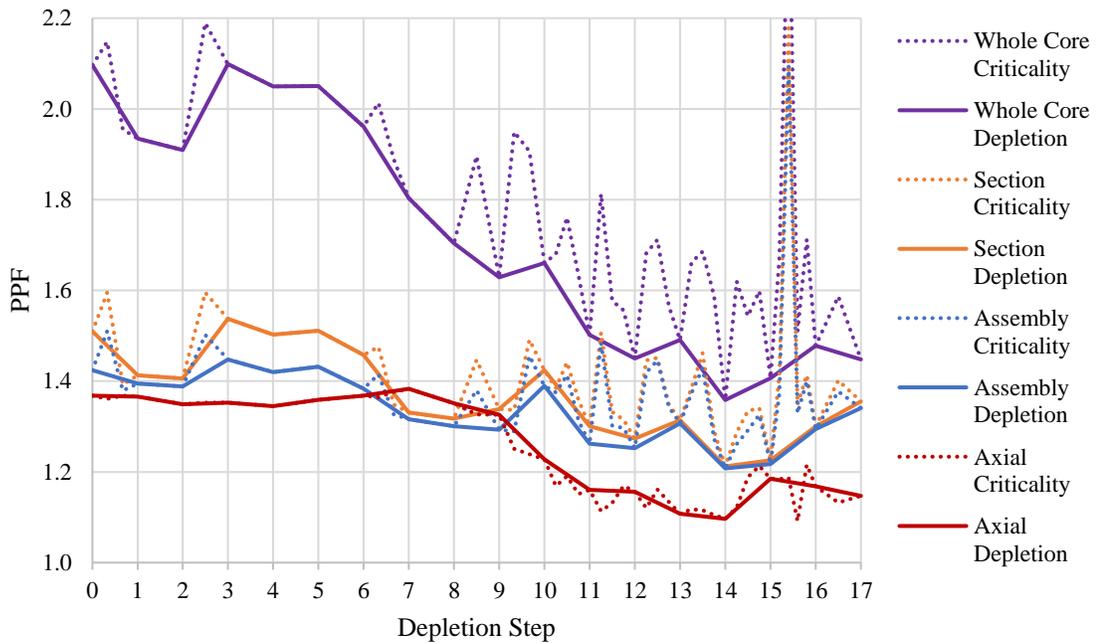


Figure 8.50. PPFs over depletion sequence using criticality iteration via CB movement and equilibrium xenon treatment.

The AO results for the case using criticality iteration via CB movement and equilibrium xenon treatment can be found in Figure 8.52.

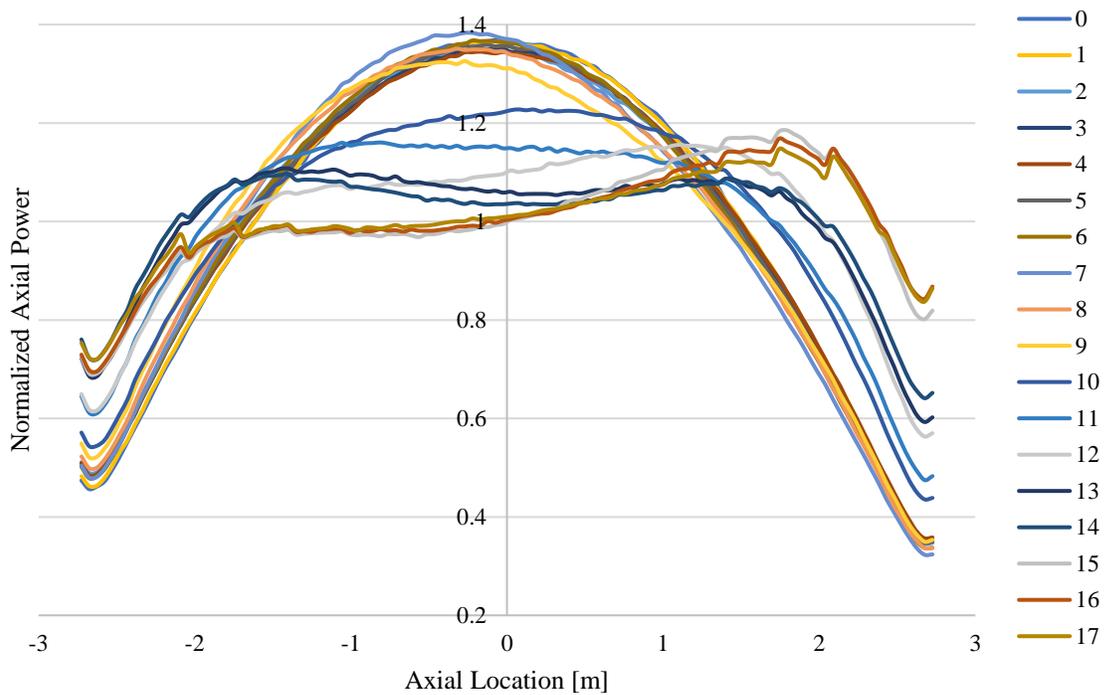


Figure 8.51. Normalized axial power profiles over depletion sequence using criticality iteration via CB movement and equilibrium xenon treatment. Legend values are for each depletion step.

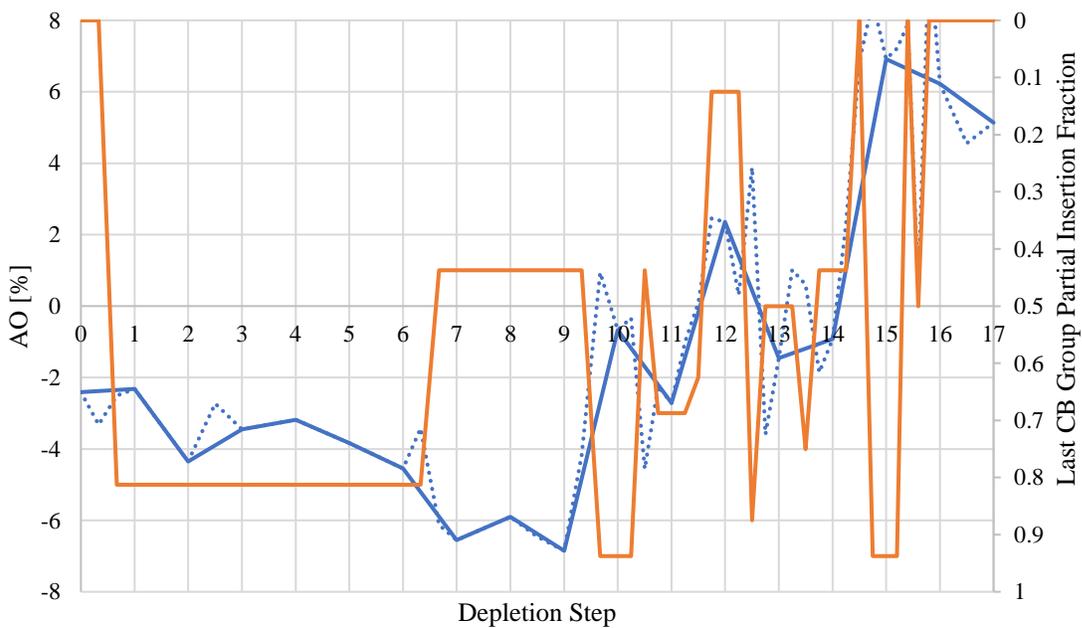


Figure 8.52. AO over depletion sequence (solid blue) using criticality iteration (dashed blue line) via CB movement and equilibrium xenon treatment. Partial insertion fraction of the last CB group (orange) is shown for reference for impact on the AO.

8.3.5.c Comparison between Xenon Treatments

This section compares some of the results between the explicit and equilibrium xenon treatments used by Serpent. As was shown in Figure 8.27, there should be very good agreement between the two xenon treatments for fine spatial tracking after the first few depletion steps once the explicit treatment has had sufficient time to allow ^{135}Xe to reach its equilibrium concentration. When comparing the number of CBs groups inserted over the depletion sequence in Figure 8.53, one can see that there is excellent agreement. After the third depletion step (1 MWd/kgHM or about 5.1 EFPD), the two xenon treatments agree extremely well. Some of the criticality search substeps might differ, but they eventually arrive to comparable critical insertion configurations (within the user-specified eigenvalue tolerance of ± 150 pcm).

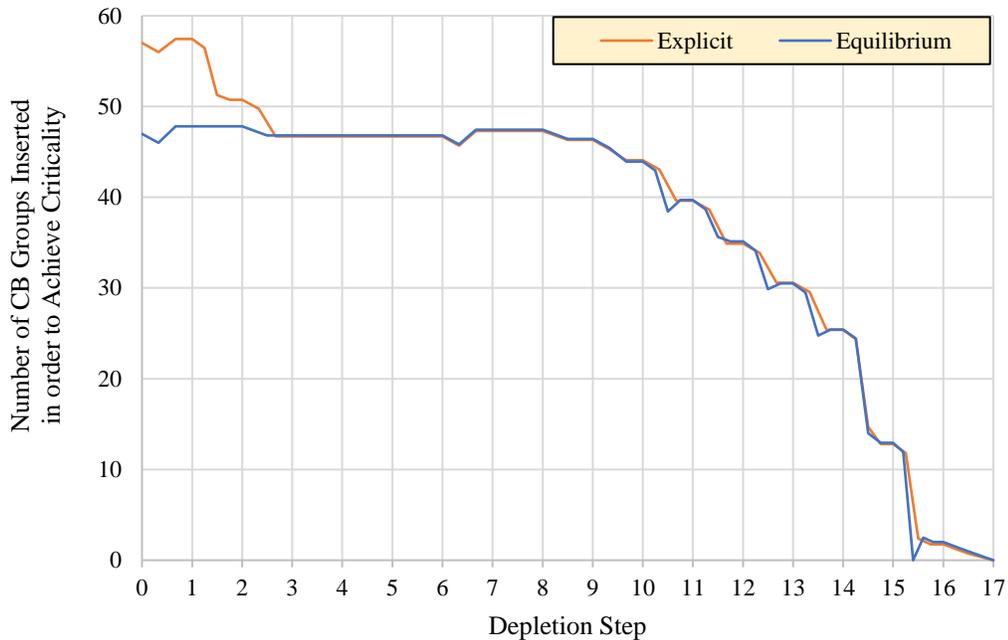


Figure 8.53. Critical CB insertion for explicit and equilibrium xenon treatments.

A comparison of the AO results for the two xenon treatments can be found in Figure 8.54. It also includes the results from the depletion case using fine spatial tracking and no CBs inserted (Figure 8.24). An important distinction for CB insertion is that CBs enter the core from the top. In the geometry model created by ATOMICS, the CB is explicitly modeled in the top axial reflector

and top axial support plate if inserted, including both fractional and complete insertions. Even for a complete CB insertion, this axial difference will shift the power toward the bottom of the core. As was observed in Figure 8.24 for fine spatial tracking without criticality iteration, AO was initially negative (due to axial support plate composition) but shifts toward the top of the core for later burnup steps. For the depletion sequences considered here, the initial axial peaking toward the bottom of the core is much larger for the first few depletion steps due to partial CB group insertion and CB presence in the top axial reflector and top axial support plate. This produces the general behavior seen in burnup steps 0 through 9 for the blue and orange lines. Due to the power being more peaked toward the bottom of the core during the initial burnup steps, the power shifts toward the top of the core sooner than for the depletion case run with all CBs withdrawn (red line). Since the latter burnup steps are also longer, more CBs are withdrawn at a time for these burnup steps. Withdrawing CBs from these assemblies with a more axially peaked burn history would promote a stronger top-peaked profile (depletion steps 10 and 12 orange; 12, 15, and 16 blue). However, partial CB insertions near the middle of the core could still shift the power lower for those depletion steps (11, 13, and 14 for both orange and blue). When all CBs are finally withdrawn from the core (depletion step 17), the results still have good agreement despite the varied depletion history between the three cases shown.

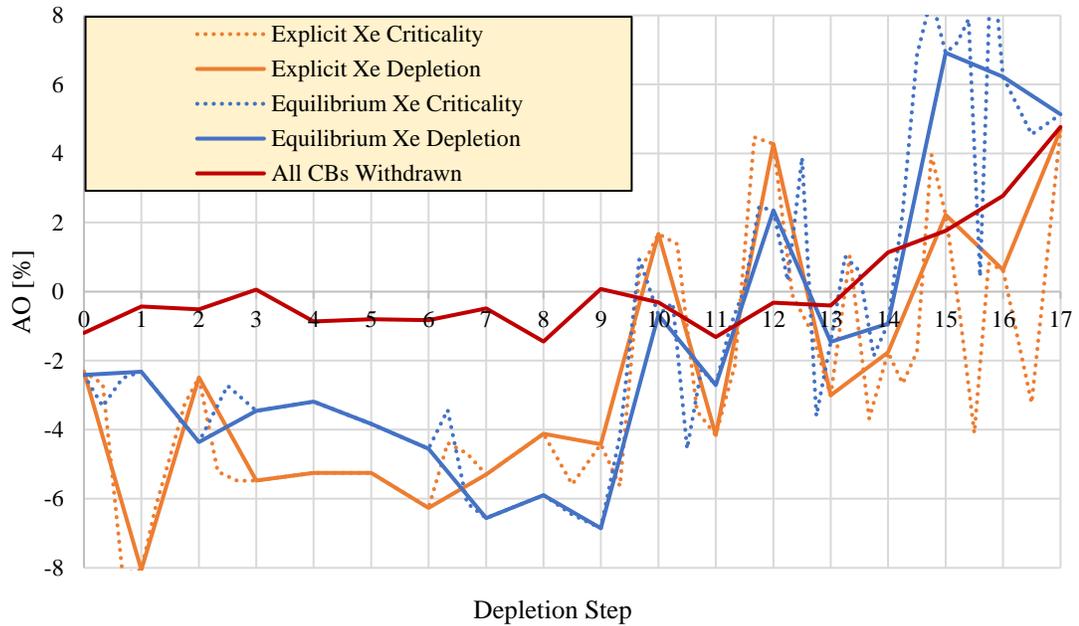


Figure 8.54. AO over depletion sequence using criticality iteration via CB movement for both explicit (orange) and equilibrium (blue) xenon treatments. Previous results from the depletion case with all CBs withdrawn (red) is shown for comparison.

8.3.6 Fine Tracking with Both Multiphysics and Criticality Control via Control Blade Movement

The final depletion testing set focuses on using both TH feedback and criticality control via CB movement. Both substep processes are executed concurrently since no stability issues were observed during testing of the capabilities. This means that for each substep simulation, CBs are moved and TH is simultaneously updated. Once the critical CB configuration is achieved, the substep process will continue for only updating TH based on the number of iterations prescribed by the user. Based on testing in section 8.3.4, it is recommended that one TH iterative substep be used beyond finding the critical insertion configuration. This section contains two subsections. The first focuses on the impact of the number of histories simulated on a full-feature implementation of ATOMICS using both multiphysics and criticality iteration. The second considers the depletion results from using the explicit xenon treatment with PC. As demonstrated in section 8.3.4, PC is necessary to manage the axial numerical instabilities which would otherwise occur without it.

8.3.6.a Impact of Number of Simulated Particles

Similar to the comparative study done in section 8.3.3.b with neither TH feedback nor CB movement, this section addresses the impact that particle statistics play in full-feature simulation of ATOMICS which uses TH feedback and CB movement. Each case uses 100 inactive cycles and either 200 active cycles for substeps or 500 active cycles for depletion steps. Cases varied based on the number of particles used per generation. Three levels of granularity were considered: 1k, 10k, and 100k particles per generation. In section 8.3.3.b, it was observed that biases existed in depletion simulations when using 25k and 2.5k particles per generation relative to using 250k particles per generation. It should be fully expected that comparable biases would be present in these simulations as well. Representative statistical uncertainties for the cases considered in this subsection are shown in Table 8.25. For each case, the eigenvalue tolerance for the critical insertion search was also relaxed to compensate for larger statistical uncertainties in cases simulating fewer particles.

Table 8.25. Example statistical uncertainties for cases using different neutron generation sizes.

Neutron Generation Size	Eigenvalue σ for Iteration Substeps [pcm]	Eigenvalue σ for Depletion Steps [pcm]	Eigenvalue Tolerance Used for Criticality Search [pcm]
1,000	300	180	450
10,000	95	57	300
100,000	30	18	150

When using criticality iteration, the number of CB groups inserted in the core over the depletion sequence is presented instead of eigenvalue in Figure 8.55. Note that a large discrepancy occurred in the case using only 1k particles per generation for depletion steps 13 and 14. This happened due to a combination of large statistical uncertainties and an oversight in possible CB insertion criteria for deciding when to conduct depletion transport, which has since been corrected⁶.

⁶ ATOMICS handles the two extreme insertion cases of all CBs inserted and all CBs withdrawn separately from the continuum between the extremes. If more CBs are estimated to be inserted than are available for insertion (84 CB groups), then ATOMICS inserts the 84 maximum possible groups and runs depletion as the best possible configuration. If fewer CBs are estimated to be inserted than zero (subcritical core), then ATOMICS withdraws all CB groups and runs depletion as the best possible configuration. What happened with the case using 1,000 particles per generation was that due to the very poor statistics, a CB movement substep case ran where a subcritical configuration *withdrew* CBs and the eigenvalue just barely *decreased*. Then, due to the critical insertion interpolation feature used by ATOMICS, it predicted that all CBs should be inserted to achieve a critical configuration. Although unphysical and a result of the poor statistics with equally poor luck, ATOMICS guessed correctly given the results provided but should not have chosen to move toward

Beyond these two depletion points when using extremely poor statistics, the agreement over the depletion sequence is quite good. Reactivity differences at EOC (70 MWd/kgHM) when all CBs are removed relative to the reference case using 100k particles per generation were 242 ± 65 pcm for the case using 10k particles per generation and 328 ± 195 pcm for the case using 1k particles per generation. These biases are comparable to those found in Figure 8.26.

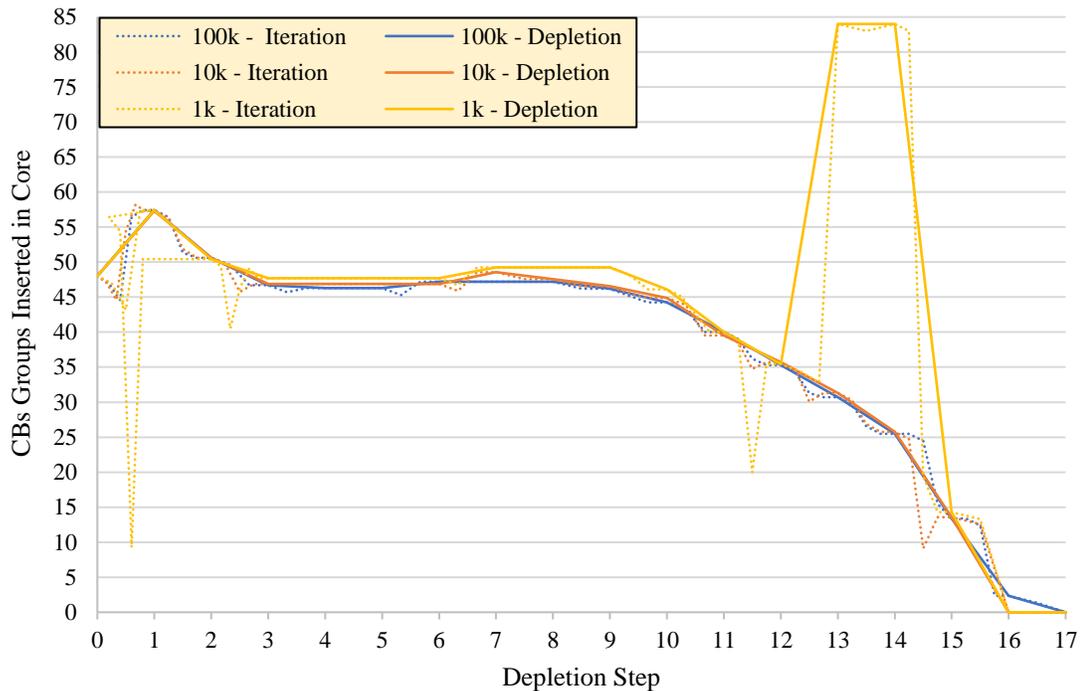


Figure 8.55. Number of control blade groups inserted to achieve critical insertion over the depletion sequence for different neutron generation sizes.

8.3.6.b Depletion using Explicit Xenon Treatment

This subsection considers a depletion sequence using the PC method, TH feedback, criticality iteration via CB movement, and explicit xenon treatment. The results come from the same simulation used for the 100,000 particles per generation case shown in the previous subsection. A summary of the results over the depletion sequence can be found in Table 8.26.

depletion. Now, there is at least a logical case when depletion only occurs for all CBs inserted when supercritical or for all CBs withdrawn when subcritical.

Table 8.26. Depletion sequence summary for using PC, TH feedback, criticality iteration via CB movement, and explicit xenon treatment.

Depletion Step	CB Movement Step	k^*_{eff}	CB Groups Inserted	Axial PPF	Axial Offset [%]	Assembly PPF	Section PPF	Whole Core PPF
0	1	1.05141	48	1.402	-3.75	1.417	1.436	2.042
0	2	1.03652	47	1.340	-6.94	1.654	1.760	2.413
0	3	1.04596	44.5625	1.354	-10.53	1.527	1.574	2.158
0	4	1.00349	56.4375	1.366	-9.03	1.398	1.441	1.955
0	5	0.99932	57.4375	1.373	-10.57	1.345	1.382	1.883
1	1	0.97672	57.4375	1.355	-8.84	1.322	1.357	1.819
1	2	0.98054	56.4375	1.341	-9.41	1.341	1.382	1.846
1	3	0.99776	51.3125	1.340	-5.23	1.326	1.383	1.894
1	4	0.99928	50.625	1.320	-8.06	1.260	1.338	1.776
2	1	0.98605	50.625	1.311	-7.04	1.251	1.310	1.752
2	2	0.98956	49.625	1.302	-8.41	1.226	1.297	1.701
2	3	1.00093	46.6875	1.311	-5.80	1.532	1.629	2.139
3	1	0.99835	46.6875	1.311	-7.96	1.315	1.398	1.809
3	2	1.00262	45.6875	1.317	-7.75	1.382	1.458	1.974
3	3	1.00089	46.3125	1.346	-6.84	1.363	1.438	1.922
4	1	1.00038	46.3125	1.334	-7.87	1.379	1.458	1.964
5	1	1.00324	46.3125	1.339	-6.77	1.393	1.458	1.979
5	2	1.00696	45.3125	1.360	-7.18	1.319	1.376	2.017
5	3	1.00004	47.1875	1.353	-6.38	1.349	1.425	2.013
6	1	1.00100	47.1875	1.351	-5.61	1.284	1.336	1.881
7	1	1.00011	47.1875	1.339	-4.91	1.268	1.334	1.818
8	1	0.99700	47.1875	1.322	-4.32	1.265	1.286	1.716
8	2	1.00137	46.1875	1.283	-3.78	1.408	1.465	1.855
9	1	0.99041	46.1875	1.294	-4.99	1.234	1.262	1.567
9	2	0.99542	45.1875	1.242	-3.31	1.569	1.604	2.160
9	3	0.99858	44.25	1.233	-4.11	1.418	1.464	1.851
10	1	0.98501	44.25	1.236	-2.63	1.353	1.387	1.625
10	2	0.98852	43.25	1.170	-1.72	1.307	1.311	1.737
10	3	0.99934	40	1.172	2.00	1.367	1.425	1.730
11	1	0.98392	40	1.167	1.40	1.232	1.263	1.432
11	2	0.98817	39	1.161	3.08	1.699	1.732	2.044
11	3	0.99724	36.1875	1.135	-0.31	1.344	1.368	1.589
11	4	1.00108	35.3125	1.127	-2.33	1.263	1.304	1.579
12	1	0.98413	35.3125	1.135	-3.24	1.202	1.253	1.383
12	2	0.98807	34.3125	1.109	-1.72	1.302	1.331	1.747
12	3	0.99798	31.3125	1.106	-1.48	1.358	1.377	1.657
12	4	1.00031	30.6875	1.158	3.38	1.377	1.380	1.656
13	1	0.98324	30.6875	1.106	-1.22	1.307	1.315	1.423
13	2	0.98730	29.6875	1.111	0.72	1.463	1.469	1.737
13	3	0.99683	26.5625	1.105	-2.10	1.262	1.293	1.447
13	4	1.00090	25.5	1.087	-0.87	1.273	1.290	1.567
14	1	0.95927	25.5	1.090	-0.91	1.189	1.195	1.361
14	1	0.95927	25.5	1.079	-0.55	1.290	1.327	1.661
14	2	0.96333	24.5	1.151	-5.67	1.326	1.358	1.642
14	3	0.99306	15.5	1.086	0.75	1.252	1.284	1.405
14	4	0.99955	13.375	1.087	-1.67	1.150	1.181	1.333
15	1	0.95924	13.375	1.090	-0.83	1.220	1.230	1.604
15	2	0.96294	12.375	1.093	-2.21	1.362	1.380	1.625

Table 8.26 (continued).

15	3	0.99902	2.375	1.107	-3.01	1.198	1.201	1.392
16	1	0.95726	2.375	1.153	-5.94	1.282	1.302	1.594
16	2	0.96097	1.375	1.169	6.72	2.049	2.111	2.435
17	1	0.92612	0	1.201	8.21	1.369	1.378	1.584

*Iteration calculations had an average statistical uncertainty of 28 pcm. Depletion had 18 pcm.

A plot of the spatial PPF values resolved at the whole core, one-third assembly section, assembly, and axial levels over the depletion sequence is presented in Figure 8.56. These profiles are comparable to those found for other depletion case considered in this chapter.

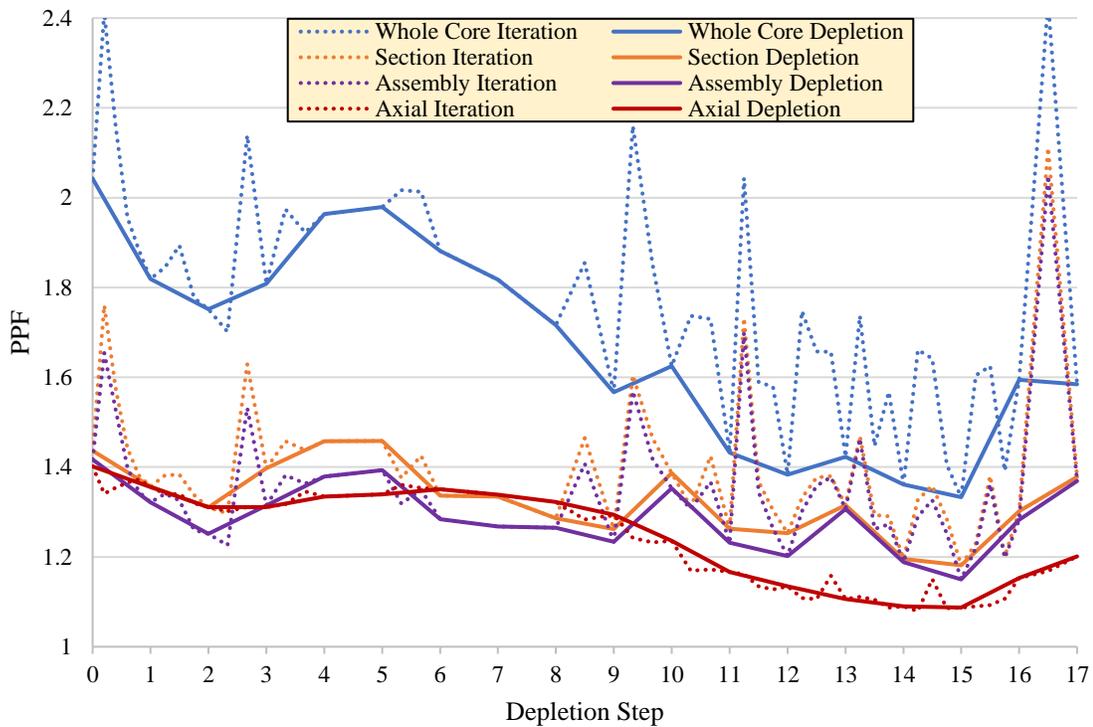


Figure 8.56. PPFs over the depletion sequence when using thermal hydraulic feedback, criticality iteration via control blade movement, and explicit xenon treatment.

AO behavior over the depletion sequence can be observed in Figure 8.57. As observed before for TH feedback and critical insertion individually, the AO is pushed lower in the core at BOC due to the negative temperature coefficient from TH feedback and CBs being inserted from the top of the core with an asymmetric partial insertion. Later in the depletion sequence, one can see that the AO is kept low by partial insertions near the middle of the core, which are the most

asymmetric insertions as discussed in a previous section. At EOC, when all CBs are withdrawn, the AO is strongly positive (about +8%), which is a slightly higher value than for other considered depletion cases. This is likely attributable to the fact that the AO was kept toward the bottom half of the core over most the depletion sequence, especially the last few depletion steps. Previous studies showed that the AO would become positive for last few depletion steps as the initial axially asymmetric burn rates depleted the bottom half of the core more than the top half. Since the CB partial insertions and TH feedback kept the power peaked to the bottom half for a longer burn time, the tendency for the power to shift to the top of the core is stronger at EOC due to these accrued asymmetric effects over the rest of the depletion sequence.

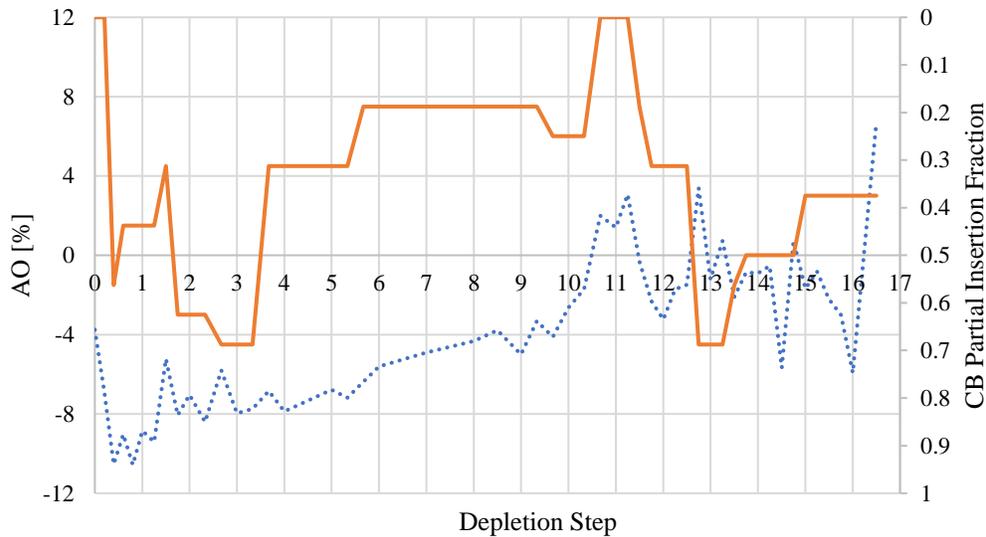


Figure 8.57. AO over depletion sequence (solid blue line) when using thermal hydraulic feedback and criticality iteration via control blade movement (dashed blue line) with explicit xenon treatment. CB partial insertion fraction (orange line) shown for reference for power shaping. AO for no TH feedback and no criticality iteration also shown for reference (black line).

8.3.7 Runtime Comparison of Depletion Simulations

A summary of the runtimes for the depletion simulations shown in this section can be found in Table 8.27. “Fine Tracking” refers to the 4032 zones used for all but the first set of depletion studies, and the cases without it use a single material. “Equilibrium” refers to using equilibrium

xenon treatment for fission products associated with ^{135}Xe and cases with “Explicit” use explicit xenon treatment. For the first four cases, the processing time per step is relatively shorter than the other cases because the geometry only needed to be generated one time since it remains static over the depletion sequence. Only isotopics need to be updated between depletion steps, which is a relatively quick process compared to rendering the full 3D AHTR geometry. The single geometry processing instance is averaged into the step-wise processing time, which still results in being much quicker than the other depletion cases.

Table 8.27. Runtime summary for depletion cases considered in this dissertation.

Case				Total Wallclock [h]	Total Time [CPU·h]	Step Type	Particles/ Cycle	Active Cycles	Inactive Cycles	# Steps	Wallclock/ Step [h]	Transport Time/ Step [h]	Processing Time/ Step [h]
Fine Tracking	CB Movement	TH Iteration	Xenon Treatment										
			Explicit	352.4	8457	Dep Step	250000	500	100	17	20.7	20.6	0.1
			Equilibrium	367.5	8820	Dep Step	250000	500	100	17	21.6	21.5	0.1
X			Explicit	359.2	8621	Dep Step	250000	500	100	17	21.1	20.9	0.2
X			Equilibrium	368.1	8835	Dep Step	250000	500	100	17	21.7	21.5	0.2
X		X	Explicit	384.0	9217	TH Substep	100000	200	100	17	5.9	2.5	3.4
						Dep Step	100000	500	100	17	16.7	13.3	3.4
X		X	Equilibrium	394.2	9460	TH Substep	100000	200	100	17	6.1	2.7	3.4
						Dep Step	100000	500	100	17	17.1	13.5	3.6
X	X		Explicit	392.3	9416	CB Substep	100000	200	100	28	5.3	1.9	3.4
						Dep Step	100000	500	100	17	14.3	10.9	3.4
X	X		Equilibrium	400.3	9608	CB Substep	100000	200	100	29	5.3	1.9	3.4
						Dep Step	100000	500	100	17	14.5	10.9	3.6
X	X	X	Explicit	452.9	10870	TH + CB Sub	100000	200	100	35	5.4	1.9	3.5
						Dep Step	100000	500	100	17	15.6	12.1	3.5

8.3.8 Spatial Power Performance of Depletion Simulations

A summary of the PPF results for all the depletion cases considered in this chapter are presented below. Each individual PPF (whole core, one-third assembly section, assembly, and axial) is plotted in its own figure. Nine cases in total are presented: single material tracking for both equilibrium and explicit xenon tracking, 4032 material tracking for both equilibrium and explicit xenon tracking, criticality iteration via CB movement for both equilibrium and explicit xenon tracking, TH iteration for both equilibrium and explicit xenon tracking, and both criticality iteration via CB movement and TH iteration for explicit xenon treatment.

Results for the whole core PPF are shown in Figure 8.58. Single material tracking cases have very little change in behavior over the depletion sequence. Cases using fine material tracking and no active criticality control (both 4032 material tracking cases and TH iteration cases) have similar behavior over cycle. Power initially peaks at the center of core due to lack of control, burning the region much faster than the periphery. Toward EOC, the PPF drops significantly. The three cases using criticality iteration have similar performance. They have the lowest BOC values due to radial peaking being reduced from use of a CB insertion schedule created to reduce radial peaking. Due to having a flatter burn history, the profiles vary less over cycle and result in higher EOC values (which are actually more realistic). The whole core PPF results are dominated by the radial profile since the axial profile has both less peaking and less variation between cases. Similar discussions for the whole core PPF are extendible to the one-third assembly section (Figure 8.59) and assembly (Figure 8.60) PPF summaries.

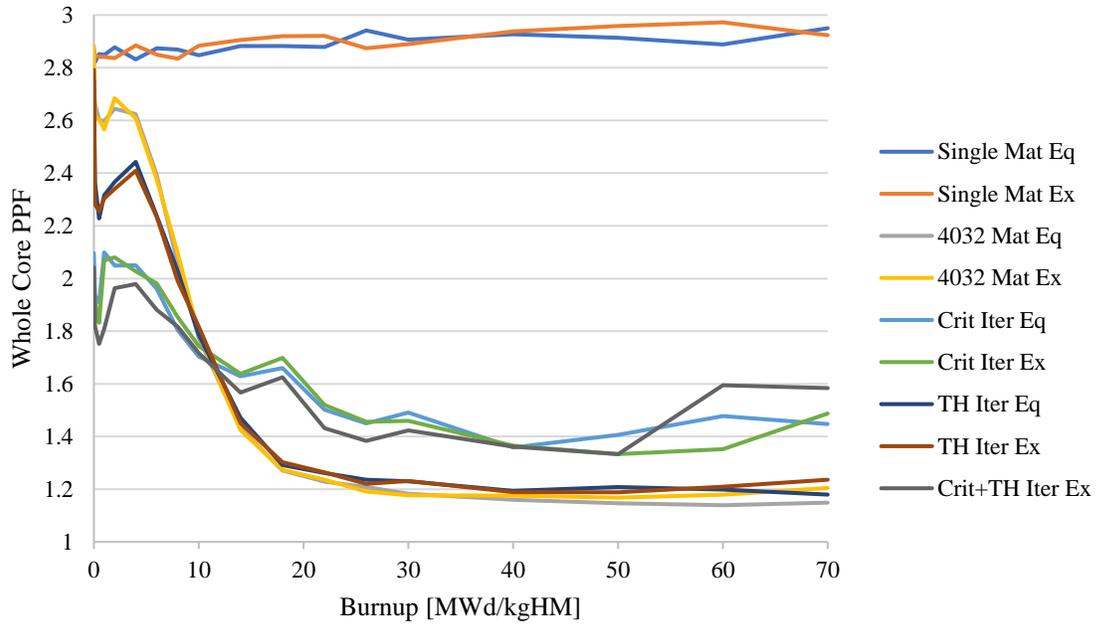


Figure 8.58. Whole Core PPFs for all depletion cases.

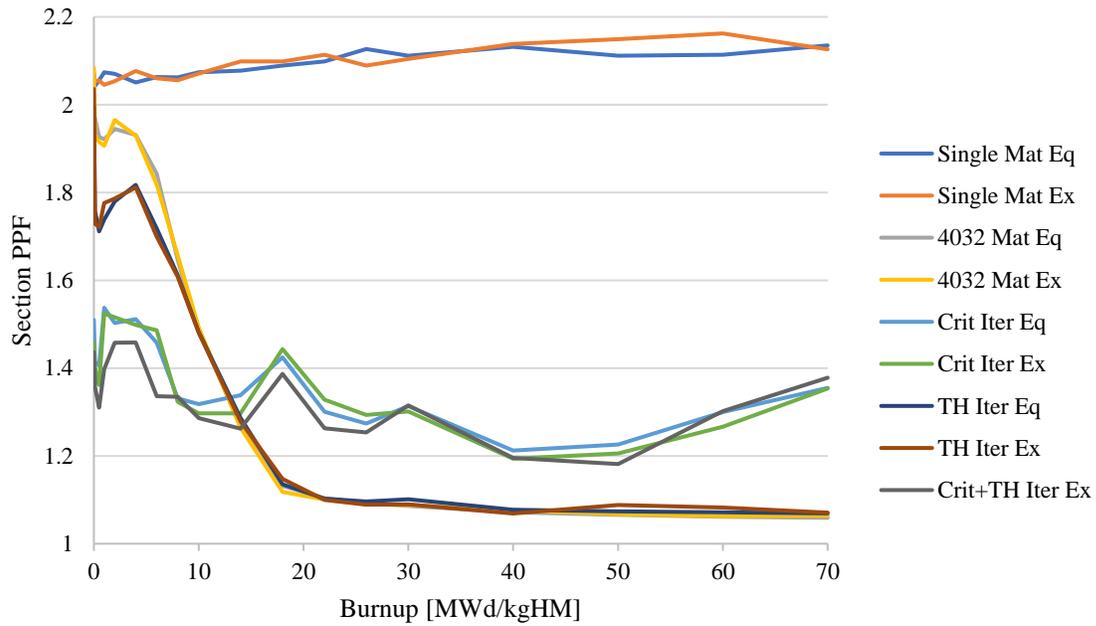


Figure 8.59. One-third assembly section PPFs for all depletion cases.

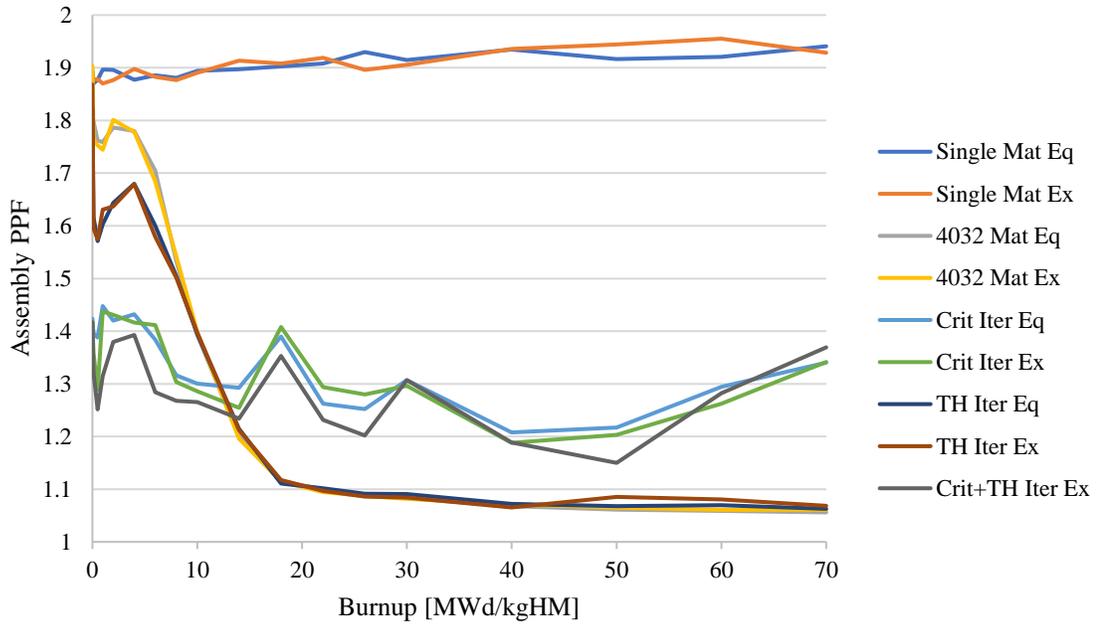


Figure 8.60. Assembly PPFs for all depletion cases.

Results for the axial PPF are shown in Figure 8.61. Once again, both single material tracking cases have little variation over the depletion sequence. The four cases with fine material tracking and no criticality iteration (both 4032 material cases and TH iteration cases) all decrease quickly to an axial PPF of about 1.1 at about 25 MWd/kgHM, which then slightly increases over the remainder of the depletion sequence. The three cases with criticality iteration decrease more slowly initially. The cause of this is the fact that these cases initially have a larger magnitude AO due to the presence of CBs in the core, which drives a slightly larger axial PPF as well since less symmetric.

Results for AO are shown in Figure 8.62. The single material cases are consistently just barely skewed toward the bottom of the core over the entire cycle. The general trend for the other cases is for the AO to be negative at BOC and positive at EOC. Cases with criticality iteration see larger shifts in AO over cycle due to different partial insertions of the last CB group to achieve criticality. These variations are expected and not driven by numerical instabilities.

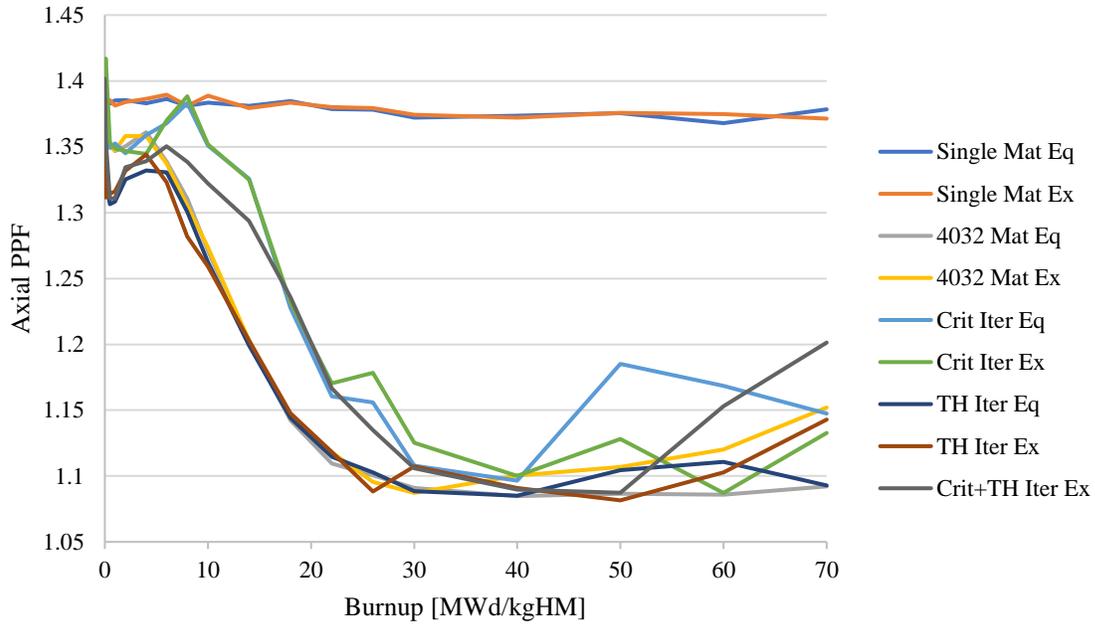


Figure 8.61. Axial PPFs for all depletion cases.

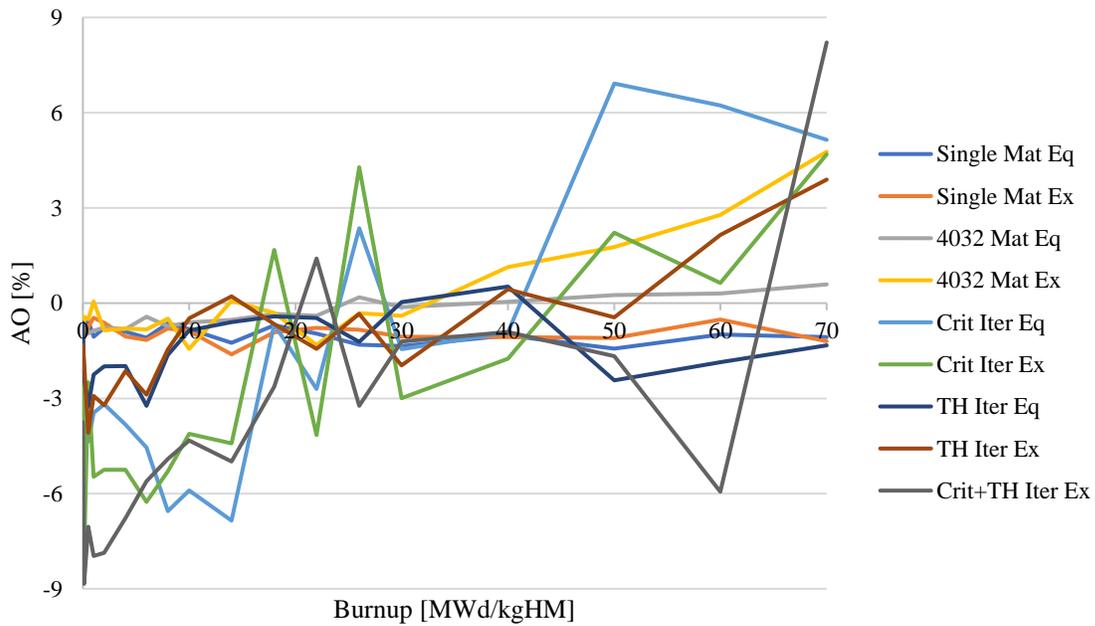


Figure 8.62. AO for all depletion cases.

CHAPTER 9

CONCLUSION

Work began by addressing the multiphysics areas of thermal expansion, heat transfer, thermal hydraulics, materials properties, and neutronics. These individually developed models were then integrated into the scripting utility ATOMICS which automatically processes results from previous simulations to create Serpent-specific input files for subsequent transport simulations. The AHTR core was analyzed in-depth using a highly detailed Serpent model with multiphysics capabilities. Many features of the complex AHTR system were parameterized to allow for easy modification of fuel element components and characteristics with respect to multiphysics coupling. ATOMICS was tested using several depletion cases to obtain cycle results as well as provide insight into the numerical stability of depletion simulations of the large AHTR design when using various reactor physics options. This work concludes by reviewing the development of ATOMICS and where future scripting work could be applied to expand the functionality of the code.

9.1 Review of Multiphysics Scripting Efforts

The motivation of this work was to create a practical tool for modeling the AHTR system with both high fidelity and multiphysics coupling. Here the individual components of multiphysics development are reviewed to capture the total effort.

9.1.1 Thermal Expansion

During operation, the components within the AHTR heat up due to power production in the TRISO fuel particles. Temperature change drives the behavior of thermal expansion, which causes components in the core to expand in size and correspondingly decrease in density. An important component of thermal expansion is parameterizing the geometric definitions of core components within Serpent so that their temperature-specific behaviors can be properly captured.

This was no small feat for the complex AHTR system due to the number of surfaces required to generate the intricate assembly layout. Additionally, components expand by different amounts due to material-specific thermal expansion coefficients and temperature gradients within the fuel assembly. This required careful treatment of core components to capture effects with the largest reactivity impact while still conserving mass of all materials.

The thermal expansion methodology used in this work was presented in Chapter 3. It began by covering assumptions used by the methodology and the scope of resolving geometric surfaces undergoing expansion. Equations were provided for relevant temperature-dependent dimensional changes and corresponding density decreases. Fine detail was paid to TRISO fuel particle, fuel stripe, and fuel plate expansions. Thermal expansion considerations axially spanned from the top and bottom axial support plates (including the active core and axial reflector regions) and radially spanned everything within the reactor pressure vessel (fuel assemblies, removable reflector assemblies, permanent reflector, boron carbide shield, core barrel, vessel liner, and vessel wall).

9.1.2 Heat Transfer and Thermal Hydraulics

Temperature feedback was the second component of multiphysics methods used in this work. As heat is produced in the TRISO fuel kernels, it ultimately is removed by the FLiBe cooling channels where the coolant warms as it flows up through the active core. Axial heat conduction was shown to be minimal in previous work [22] and heat was modeled to only move radially. Due to the fact that fuel assemblies are completely enclosed by a structural C-C composite wrapper around the hexagonal perimeter, assembly crossflow was neglected.

The heat transfer and thermal hydraulic methodology used in this work was presented in Chapter 4. Discussion began with a fuel stripe homogenization process to simplify the heat transfer to a 1D system. This produced an average temperature profile from the center of the coolant channel to the center of the fuel plank. Explicit fuel particle temperature distributions were recovered by superimposing a power modulation function onto the average temperature profile. Heat conduction

was tracked from the fuel kernel through the other TRISO particle layers (buffer, inner pyrolytic carbon, silicon carbide, and outer pyrolytic carbon layers), matrix graphite, sleeve graphite, and then ultimately transported to the FLiBe coolant. Supporting simulations demonstrated the linearity of the fuel temperature coefficient to justify the implementation of the homogenization and power modulation profile recovery methods.

9.1.3 Neutronics

Neutron simulations were executed using the 3D MC transport code Serpent. Conducting multiphysics studies of the AHTR system requires fine resolution of power production within the core. In Chapter 5, this was shown to be impractical for existing tallying methods within Serpent, so a new tally type was developed and implemented into Serpent to efficiently obtain local power results at the one-third assembly level with virtually no transport runtime penalty. The new tally was tested, and the results were shown to be consistent the results from other, drastically slower tallies previously available within Serpent.

Excess reactivity over the fuel cycle is controlled passively by BP spheres embedded in the fuel planks and actively by the insertion of CBs within the core. To maintain criticality over cycle, CBs are moved to account for depletion of fuel and other factors impacting core reactivity. In Chapter 6, CB movement schemes were discussed. Automated CB movement features were developed to find insertion and withdrawal schedules to beneficially shape the radial power distribution over the cycle as well as a method to iterate with Serpent transport simulations to find a critical CB configuration when following such a schedule. Testing of the schedule searching method showed that insertion searches were more stable than withdrawal searches, even though only withdrawal searches could be executed on-the-fly during a depletion simulation.

9.2 Use of ATOMICS

The individual multiphysics models were incorporated into a single code written in C++ and given the name ATOMICS. Its main objective is to create Serpent-specific input files based on

the results from previous Serpent simulations. To account for all multiphysics feature of the complex AHTR design, the input file is typically over a million lines long. ATOMICS reads-in power distribution results obtained from using the new tally developed as part of this work. These local powers update local temperature profiles, which are used in the next iteration's material temperatures as well as dimensions via thermal expansion. The user controls the choice of specific options and how ATOMICS runs using an external input file which is read-in each time the code is executed. ATOMICS tracks results between iterations with both ATOMICS-specific external output files and Serpent-specific output files. The ATOMICS-specific outputs are meant to be easily human-readable to inform the user on the progress of the current simulation and provide iteration-wise results from each transport simulation. ATOMICS is meant to be compiled and run on a Linux system which has Serpent version 2.31 installed. The Serpent executable needs to be recompiled with updates for the new power tally. These setup steps and example commands to run simulations were presented in Chapter 7. A visual depiction of how ATOMICS works can be seen in Figure 9.1, which is a repeat of a prior figure (Figure 7.6).

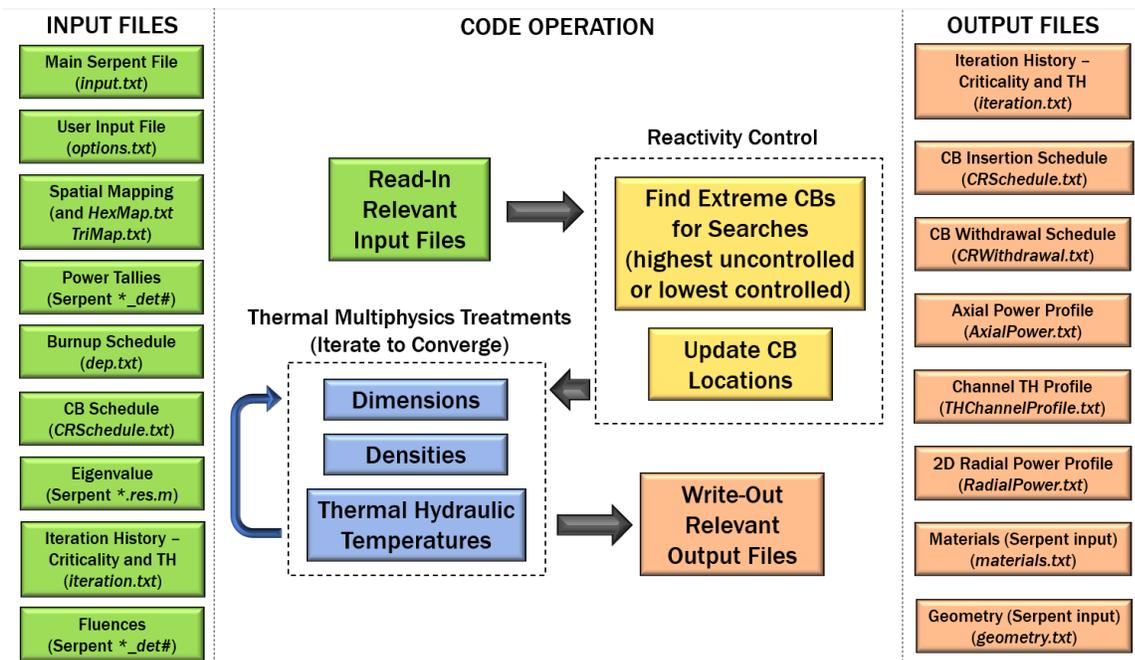


Figure 9.1. A flowchart of how ATOMICS works and a summary of important I/O files.

9.3 Summary of Results

The results found in Chapter 8 were presented in three distinct groupings: geometric sensitivity studies, thermal sensitivity studies, and depletion studies. The depletion studies fully integrate all the features of ATOMICS and are the capstone to this dissertation. The sensitivity studies highlight the benefit of having ATOMICS to make significant modifications by only modifying a few parameters in the user input file. Heat transfer and thermal hydraulic feedback allow for making local temperature and density changes which would be very time consuming to implement otherwise. Taken together, the ease with which ATOMICS makes changes to the complex AHTR model lends it to be a great tool for cross section generation. Modifications can be swiftly implemented, and Serpent can be used solely to obtain cross sections for use in other solvers.

Due to the complex geometric nature of the AHTR core, manually making changes to the geometry model is quite challenging. Fortunately, due to how the geometry model used by ATOMICS is parameterized, many features of the AHTR design can be changed by a single modification in the user input file. Results for various geometric modifications performed using ATOMICS were presented to inform any future refinements of the AHTR design.

The thermal sensitivity studies considered the impact of temperature variation within the AHTR system. Inlet coolant temperature, operating conditions, and material properties were all addressed. The simulations improved upon prior studies which only considered systems with core-averaged parameters like uniform core power and uniform coolant temperature and density. Region-specific temperature results came much closer to modeling operating conditions than simplified studies previously conducted on AHTR.

Depletion simulations are broken down into five studies of varying resolution: core average, fine spatial tracking, use of multiphysics, use of criticality iteration, and integration of multiphysics and criticality iteration. Core average results used a single fuel material for 3D depletion, which would be comparable to initial scoping studies which primarily focus on estimating integral quantities such as cycle length.

Fine spatial tracking results discretized the core into 16 axial partitions with 252 one-third assembly sections using one-third core symmetry, resulting in 4032 zones total. All subsequent depletion cases also used the same fine spatial partitioning. These results were compared with the core average results to estimate the improvement from refining geometric resolution. Reactivity differences between single and 4032 zone tracking were observed to about 2100 pcm for burnup beyond 20 MWd/kgHM (about 100 EFPD). Two simulations were run: one using explicit xenon treatment and one using equilibrium xenon treatment. Despite the large, expected differences at BOC, the results toward EOC (70 MWd/kgHM) agree very well (within 32 ± 16 pcm).

Multiphysics cases focused on determining the optimal number of thermal hydraulic substeps needed between burnup steps. Too few steps would result in ill-converged temperature distributions which do not deplete accurately. Too many steps would result in running redundant expensive transport simulations which do not improve the solution of the previous iteration for the given statistical uncertainty. Zero to four substeps per burnup step were considered, and results showed that there was no discernable benefit to using more than one substep. Therefore, it is recommended to use one thermal hydraulic substep between depletion steps when running ATOMICS and using a depletion schedule comparable to the one used in this work. Numerical instabilities were observed if no efforts were taken to reduce their occurrence. If depletion steps were executed using beginning of step isotopics, instabilities were seen for long timesteps near EOC. Use of the PC method within Serpent which assumes an estimated average isotopic profile over each depletion step mitigated this issue, eliminating observable instabilities.

Criticality iteration cases focused on CB movement over the depletion sequence. The simulations performed well, and the CB movement algorithm was demonstrated to be both efficient at finding the critical insertion configuration and stable for the number of particles used per transport calculation. Using very few particles during the transport sequence resulted in some CB placement instabilities due to large statistical uncertainties producing poor reactivity prediction

estimates, but users can avoid this by simulating a comparable number of particles to what was used in this work.

The final depletion cases focused on using both multiphysics and criticality iteration. The substep methodology used by ATOMICS first searches for the critical CB insertion, with thermal hydraulics being updated along with each transport simulation. Once the eigenvalue search is within the specified tolerance of the target value, the CB configuration is accepted, and TH iterations continue for a user-specified number of additional substeps. Results showed good performance and numerical instabilities were not observed for this most complex case.

9.4 Future Work

ATOMICS was created to update the geometry, materials and physical properties of AHTR dictated by multiphysics methods and to simulate this highly detailed 3D model within Serpent. However, only a few of the features within Serpent were leveraged and scripted to be called by ATOMICS. Future scripting efforts could target incorporating more features of Serpent to improve ATOMICS as a wrapper for the 3D MC code. Analogous to using more existing features of Serpent, ATOMICS itself could be expanded to include additional modeling capabilities to widen its breadth of features. Below are specific examples of how future work could be executed to improve the functionality of ATOMICS.

9.4.1 Further Scripting Efforts

The developers of Serpent recently released version 2.32 of the code in February 2021. ATOMICS was developed around version 2.31 of the code. Compatibility with version 2.32 was not tested, but since Serpent source code changes were necessary to implement the triangular superimposed mesh tallying capabilities, code modifications would be necessary to allow ATOMICS to function with the newest version of Serpent. These changes would not be significant though, as essentially a new tally was added to the code and a similar implementation would be performed for the new code version. The easiest way to establish compatibility would be to follow

the code changes outlined in Appendix B for the eighteen affected source files starting from the version 2.32 files a base. For files where no changes were made between versions 2.31 and 2.32, the source code modifications housed on the Github repository for ATOMICS should work fine.

One of reasons to establish compatibility with Serpent version 2.32 would be to leverage the newly implemented time integration features for determining depletion step isotopics. These were summarized in Table 8.13. These features could help with numerical stability over the fuel cycle, but the most basic improvement implementation of PC was shown to be sufficient for preventing large axial numerical oscillations over the depletion sequence considered in this work.

Although not leveraged in this work, Serpent has prebuilt multiphysics capabilities included with the code distribution [61]. One could recall from the depletion results summary seen in Table 8.27 that a nontrivial amount of time was spent for depletion simulations using fine spatial resolution to regenerate the large, detailed model geometry for each transport simulation. This requires a few hours of wallclock time before particle transport even begins; establishing a large computational overhead on simulations which cannot be overcome without either simplifying the model or changing the ATOMICS methodology. If more prebuilt features of Serpent were used by ATOMICS to make it more “integrated” within Serpent rather than “wrapping around” Serpent, it might be possible to avoid this geometry regeneration process for each transport substep. It is uncertain if this approach would be compatible with CB movement substeps due to the significant required geometric changes, but it should at the very least be possible for TH iteration substeps.

9.4.2 Verification and Validation of Results

One of the disadvantages of a novel design like AHTR is that there is limited operational data available from comparable reactor technologies. This means that validation of results is mostly impossible. The more feasible option is to verify the results by comparing them to those obtained from other codes and methodologies. For the heat transfer and TH work, comparisons could have been made to prior research using RELAP models of AHTR [22] [23]. In the future, results can be

compared with those obtained as part of currently ongoing benchmarking efforts of an AHTR-like FHR design [32] [42]. An additional consideration could be to analyze the propagation of uncertainty throughout the AHTR model. This would provide insight for the accuracy and precision of results relative to their reported statistical uncertainty. This could be further expanded to uncertainty quantification for material properties and methods used in the model to identify separate and integrated effects test facilities to support AHTR licensing.

9.4.3 Additional Modeling Capabilities

It is likely that for the fuel enrichment of 9 w% considered in this work, a two-batch refueling scheme would be used. To better simulate multiple fuel cycles, an automated fuel assembly shuffling methodology could be implemented within ATOMICS to find equilibrium cycles more readily as well as give the user ease of control for where to relocate assemblies. This point might be moot if an online refueling procedure is adopted [23], but even then, fuel assembly shuffling could be necessary.

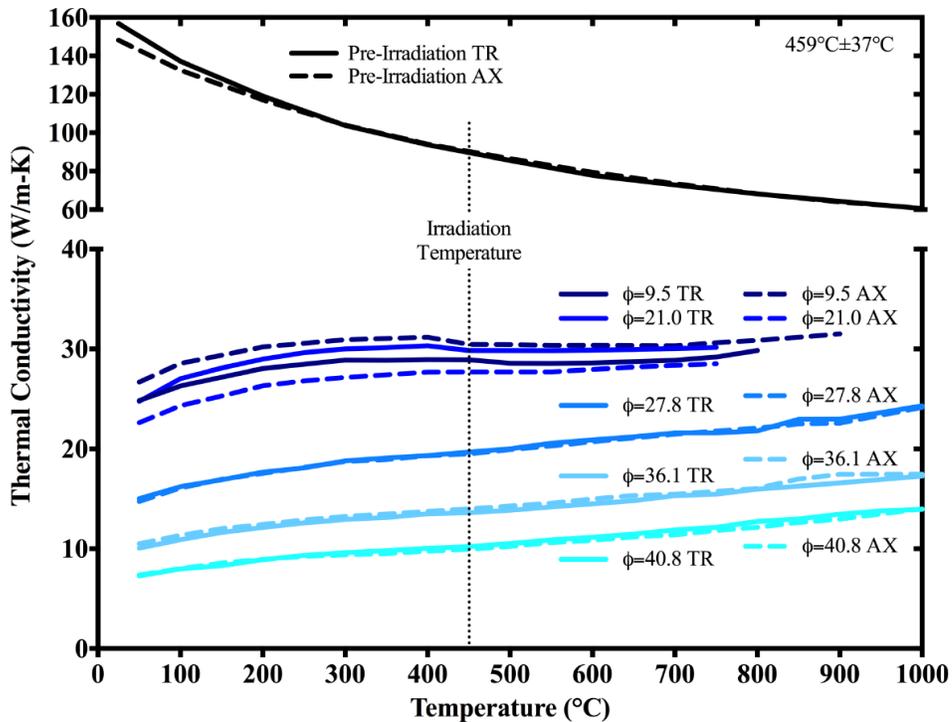
One of the final features implemented within ATOMICS was the capability to change material properties with neutron fluence. As was discussed at the end of Chapter 4, graphite density and thermal conductivity can change significantly with neutron fluence. However, these effects would likely be small as the irradiation levels experienced by the fuel assemblies over two fuel cycles would not substantially change the physical properties. Capturing the behavior could still be desired though, so adequate testing and refinement of the method used by ATOMICS would be a good candidate for additional work.

ATOMICS enables realistic analyses of AHTR using both multiphysics and criticality search. Sensitivity studies are more accessible since ATOMICS inherently makes the complex model changes for the user, which both saves engineering time and reduces the possibility of making modeling errors. Depletion studies can accurately track local changes over cycle in an automated fashion through use of iterative substeps handles entirely by ATOMICS.

APPENDIX A

DEVELOPMENT OF GRAPHITE THERMAL CONDUCTIVITY BEHAVIOR

In Figure 4.14, the temperature- and neutron fluence-dependent thermal conductivity behavior was presented for grade G347A graphite for the temperature range 400-1000 °C. The original figure taken from the reference over a larger temperature range is shown in Supplementary Figure A.1.



Supplementary Figure A.1. Measured thermal conductivity versus measurement temperature for as-received and specimens irradiated at 459 ± 37 °C. The vertical dashed line indicates the irradiation temperature. The labels indicate the specimen orientation and total neutron fluence ($\times 10^{25}$ n/m² [E > 0.1MeV]) [36].

To start, the dataset for a neutron fluence of 9.5×10^{25} n/m² will be ignored due to appearing inconsistent with the remainder of the results, as will be shown through the functional development. Next, observe that three general functional behaviors are evident:

1. The temperature dependence is essentially linear in the considered temperature range for all fluences.
2. The isothermal thermal conductivity change due to fluence is strictly decreasing.
3. This linear temperature dependence initially has a negative gradient with fluence but becomes positive at higher fluences.

This means that the relative change to the thermal conductivity of graphite can be expressed as seen in Equation A.1.

$$\frac{k(F,T)}{k_0} = f_2(F)(1 + f_1(T)f_3(F)) \quad (\text{A.1})$$

First solve for $f_1(T)$ for $F=0$. Assume that $f_2(F)$ and $f_3(F)$ are unity when $F=0$. This creates the linear temperature profile of Equation A.2.

$$\frac{k(0,T)}{k_0} = 1 + f_1(T) \quad (\text{A.2})$$

If fit over the non-irradiated temperature profile (90 W/(m · K) at 459 °C and 60 W/(m · K) at 1000 °C), the profile is obtained in Equation A.3.

$$\frac{k(0,T)}{k_{0,459}} = 1 - \frac{T-459}{1938} \quad 459 < T < 1000 \quad (\text{A.3})$$

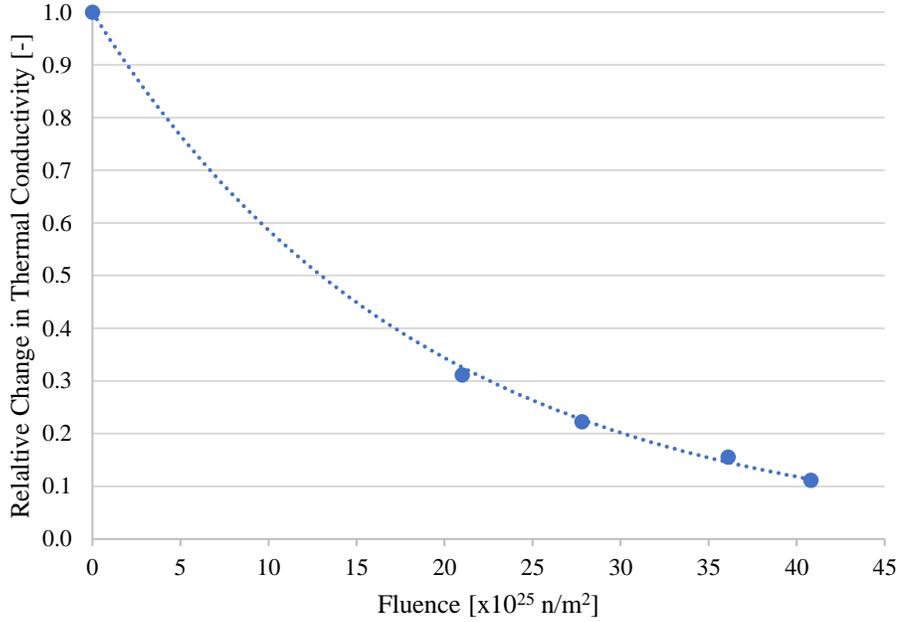
Next, the functional fit for $f_2(F)$ will be found at $T=459$. At this temperature, $f_1(T=459)$ is zero, which allows for $f_2(F)$ to be isolated since the term with $f_3(F)$ become zero as consequence of the product, giving Equation A.4.

$$\frac{k(F,459)}{k_0} = f_2(F) \quad (\text{A.4})$$

By using the actual values of the thermal conductivity of graphite inferred from Supplementary Figure A.1 and listed in the column labeled “Actual 459 °C” of Supplementary

Table A.1, the goodness of fit of $f_2(F)$ is shown in Supplementary Figure A.2 with $R^2=0.9974$. It has functional form:

$$f_2(F) = e^{-0.053364F} \quad (\text{A.5})$$



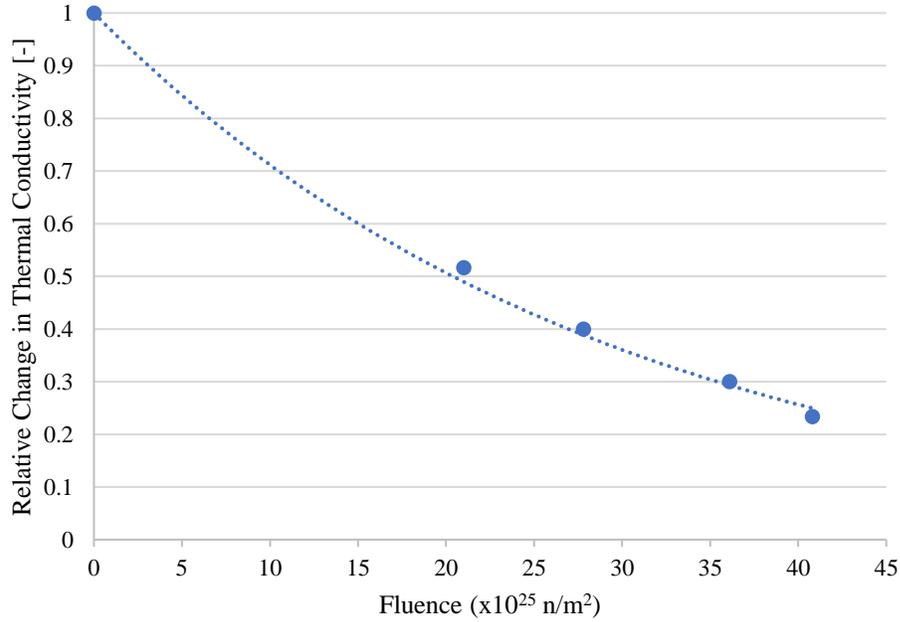
Supplementary Figure A.2. Fit of fluence-dependent function f_2 at 459 °C.

Finally, $f_3(F)$ can be found by considering the thermal conductivity due to the fluence for a temperature of 1000 °C. This gives the form:

$$\frac{k(F,1000)}{k_{0,459}} = e^{-0.053364F} \left(1 - \frac{1}{3}f_3(F)\right) \quad (\text{A.6})$$

By using the actual values of the thermal conductivity of graphite inferred from Supplementary Figure A.1 and listed in the column labeled “Actual 1000 °C” of Supplementary Table A.1, the goodness of fit of the expression shown in Equation A.6 is shown in with $R^2=0.9929$. It has the functional form seen in Equation A.7.

$$e^{-0.034000F} \quad (\text{A.7})$$



Supplementary Figure A.3. Fit of fluence-dependent function used to obtain f_3 at 1000 °C.

Equation A.6 and Equation A.7 are the same function. They can be equated to solve for f_3 , which gives Equation A.8.

$$f_3(F) = 3 - 2e^{0.019364F} \quad (\text{A.8})$$

The combined full expression for the thermal conductivity of graphite is then given by Equation A.9.

$$\frac{k(F,T)}{k_{0,459}} = e^{-0.053364F} \left(1 - \frac{T-459}{1938} (3 - 2e^{0.019364F}) \right) \quad \begin{array}{l} 0 < F < 40.8 \\ 459 < T < 1000 \end{array} \quad (\text{A.9})$$

Which is the same result given in Equation 4.14. The functional results of Equation A.9 at the fluences of interest are shown in Supplementary Table A.1 under the columns “Fit 459 °C” and “Fit 1000 °C”. After comparing the agreement with the fit of the rest of dataset and the differences observed for the fluence value of 9.5×10^{25} n/m², one can see why it is believed those values were inconsistent.

Supplementary Table A.1. List of inferred actual values from reference [36] and those obtained from the function fit of Equation A.9.

		Actual 459 °C	Fit 459 °C	Actual 1000 °C	Fit 1000 °C
Fluence ($\times 10^{25}$ n/m ²)	0.0	90	90.00	60	60.00
	9.5	30*	54.21	30*	43.44
	21.0	28	29.35	31	29.38
	27.8	20	20.42	24	23.32
	36.1	14	13.11	18	17.58
	40.8	10	10.20	14	14.99

*Values not used in development of model. Suspected to be inconsistent.

APPENDIX B

MODIFICATIONS TO SERPENT SOURCE CODE FOR TRIANGULAR MESH

The appendix contains the changes to the SERPENT source code needed to create the “dht” superimposed triangular mesh tally based upon the existing “dh” superimposed hexagonal mesh. The filenames of all modified files are highlighted in cyan. Actual changes and additions are highlighted in yellow. To save space, whole files are not included and in general line numbers with corresponding changes are listed instead. Files are grouped and ordered as they are presented in

Table 5.2. The modified files are available online at the link below.

https://github.com/KyleMRamey/ATOMICS/tree/main/Triangular_Mesh_Tally

Disclaimer: The code modifications presented below are made with respect to SERPENT version 2.1.31. Any previous versions or future releases of SERPENT may use different line numbers or more significantly omit, require additional, or change the functionality of these or other files. Thus, there is no guarantee these code modifications to create the “dht” mesh tally would be compatible with any other version of SERPENT with respect to both the tally functionality and the code at large.

B.1 Recognition of Reserved and Code-Specific Keywords

“header.h”

-Line 687

```
#define MESH_TYPE_HEXXT 41  
#define MESH_TYPE_HEXYT 51
```

-Line 1588

```
long DetIdx(long, long, long);
```

-Line 1796

```
long GetLatticeIndexes(double, double, double, double, double, double, long *, long *,  
long *, long *, long);
```

“locations.h”

-Line 2638

CELL_MESH_NT,

-Line 4307

MESH_NT,

-Line 3218

```
#define LAT_TYPE_HXT 21
#define LAT_TYPE_HYT 31
```

B.2 Initialization of New Tally from Input

“readinput.c”

Need to set the triangular number to be trivially one for other mesh types (dx, dy, dz, dh), which corresponds to line # 5300, 5364, 5428, and 5830, respectively. For both independent and compatible usage with “dht” tally.

```
WDB[loc1 + MESH_NT] = 1.0;
```

– line 5755

```
else if (!strcmp(str, "dht"))
{
    /* Hexagonal triangular mesh */

    if (j == np)
        Error(loc0, "Missing type after \"%s\"", str);

    /* Check mesh pointer */

    if ((long)RDB[loc0 + DET_PTR_MESH] > VALID_PTR)
        Error(loc0, "Multiple mesh types defined");

    /* Allocate memory for mesh structure */

    loc1 = NewItem(loc0 + DET_PTR_MESH, MESH_BLOCK_SIZE);

    /* Alloc memory for index (only seek once per collision) */

    AllocValuePair(loc1 + MESH_PREV_COL_IDX);

    WDB[loc1 + MESH_NT] = (double)6;
```

```
/* Read type */
```

```
n = (long)TestParam(pname, fname, line, params[j++],  
PTYPE_INT, 2, 3);
```

```
/* Check type */
```

```
if (n == 2)  
WDB[loc1 + MESH_TYPE] = (double)MESH_TYPE_HEXXT;  
else  
WDB[loc1 + MESH_TYPE] = (double)MESH_TYPE_HEXYT;
```

```
/* Check number of parameters */
```

```
if (j > np - 8)  
Error(loc0, "Missing mesh parameters");
```

```
/* Read central coordinates */
```

```
WDB[loc1 + MESH_MIN0] =  
TestParam(pname, fname, line, params[j++], PTYPE_REAL,  
-INFTY, INFTY);  
WDB[loc1 + MESH_MIN1] =  
TestParam(pname, fname, line, params[j++], PTYPE_REAL,  
-INFTY, INFTY);
```

```
/* Read pitch */
```

```
WDB[loc1 + MESH_MAX0] =  
TestParam(pname, fname, line, params[j++], PTYPE_REAL,  
-INFTY, INFTY);
```

```
/* Read size */
```

```
WDB[loc1 + MESH_N0] =  
TestParam(pname, fname, line, params[j++], PTYPE_INT,  
1, 1000000);  
WDB[loc1 + MESH_N1] =  
TestParam(pname, fname, line, params[j++], PTYPE_INT,  
1, 1000000);
```

```
/* Read axial binning */
```

```
WDB[loc1 + MESH_MIN2] =  
TestParam(pname, fname, line, params[j++], PTYPE_REAL,  
-INFTY, INFTY);  
WDB[loc1 + MESH_MAX2] =  
TestParam(pname, fname, line, params[j++], PTYPE_REAL,  
RDB[loc1 + MESH_MIN2], INFTY);  
WDB[loc1 + MESH_N2] =
```

```

    TestParam(pname, fname, line, params[j++], PTYPE_INT,
              1, 100000);
}

```

B.3 Creation of Tallying Structure for Recording Scores

“createmesh.c”

-line 26

```

long msh, ptr, n, m, nt;
nt = 1;

```

- line 49

```

if ((type == MESH_TYPE_CARTESIAN) || (type == MESH_TYPE_HEXX) ||
    (type == MESH_TYPE_HEXY) || (type == MESH_TYPE_HEXXT) ||
    (type == MESH_TYPE_HEXYT))
{
    /***** Cartesian and hex meshes *****/

    /* Check dimensions */

    CheckValue(FUNCTION_NAME, "nx", "", nx, 1, 100000);
    CheckValue(FUNCTION_NAME, "ny", "", ny, 1, 100000);
    CheckValue(FUNCTION_NAME, "nz", "", nz, 1, 100000);

    /* Get parameters */

    if ((type == MESH_TYPE_HEXXT) || (type == MESH_TYPE_HEXYT))
        {nt=6;}
    xmin = params[0];
    xmax = params[1];
    ymin = params[2];
    ymax = params[3];
    zmin = params[4];
    zmax = params[5];

    /* Check */

    if (type == MESH_TYPE_CARTESIAN)
    {
        CheckValue(FUNCTION_NAME, "xmin", "", xmin, -INFTY, INFTY);
        CheckValue(FUNCTION_NAME, "xmax", "", xmax, xmin, INFTY);
        CheckValue(FUNCTION_NAME, "ymin", "", ymin, -INFTY, INFTY);
        CheckValue(FUNCTION_NAME, "ymax", "", ymax, ymin, INFTY);
        CheckValue(FUNCTION_NAME, "zmin", "", zmin, -INFTY, INFTY);
        CheckValue(FUNCTION_NAME, "zmax", "", zmax, zmin, INFTY);
    }
}

```

```

else
{
  CheckValue(FUNCTION_NAME, "xmin", "", xmin, -INFTY, INFTY);
  CheckValue(FUNCTION_NAME, "xmax", "", xmax, 0.0, INFTY);
  CheckValue(FUNCTION_NAME, "ymin", "", ymin, -INFTY, INFTY);
  CheckValue(FUNCTION_NAME, "ymax", "", ymax, 0.0, 0.0);
  CheckValue(FUNCTION_NAME, "zmin", "", zmin, -INFTY, INFTY);
  CheckValue(FUNCTION_NAME, "zmax", "", zmax, zmin, INFTY);
}

/*****
*/
}

```

- line 414

```
WDB[msh + MESH_NT] = nt;
```

-line 421

```

if (cont == MESH_CONTENT_RES)
{
  /* Allocate memory for results */

  ptr = AllocPrivateData(nt*nx*ny*nz + 1, RES2_ARRAY);

  /* Put pointer */

  WDB[msh + MESH_PTR_RES2] = (double)ptr;
}
else if (cont == MESH_CONTENT_DAT)
{
  /* Allocate memory for data */

  ptr = ReallocMem(DATA_ARRAY, nt*nx*ny*nz + 1);

  /* Put pointer */

  WDB[msh + MESH_PTR_DATA] = (double)ptr;
}
else if (cont == MESH_CONTENT_PTR)
{
  /* Allocate memory for pointer */

  ptr = ReallocMem(DATA_ARRAY, nt*nx*ny*nz);
}

```

“processdetectors.c”

-Line 28

```
long ebins, ubins, cbins, mbins, lbins, rbins, zbins, ybins, xbins, tribins, tbins;
```

-Line 117

```
if ((msh1 = (long)RDB[det + DET_PTR_MESH]) > VALID_PTR)
{
    tribins = (long)RDB[msh1+MESH_NT];
    xbins = (long)RDB[msh1 + MESH_N0];
    ybins = (long)RDB[msh1 + MESH_N1];
    zbins = (long)RDB[msh1 + MESH_N2];
}
else
{
    tribins = 1;
    xbins = 1;
    ybins = 1;
    zbins = 1;
}
```

-Line 1017

```
tot = ebins*ubins*cbins*mbins*lbins*zbins*ybins*xbins*tribins*tbins;
```

-Line 1174

```
if (RDB[msh1 + MESH_NT] != RDB[msh2 + MESH_NT])
    Error(det, "Mismatch in triangle mesh bins of detector %s",
          GetText(ptr + DET_PTR_NAME));
```

B.4 Bin Index Searching to Find Where a Score Should be Recorded

From “collectdet.c”

-Line 25

```
long rbins0, zbins0, ybins0, xbins0, tribins0, tbins0, eb0, ub0, cb0, mb0, lb0;
long rb0, zb0, yb0, xb0, trib0, tb0, idx0, idx2, rb2, tb, tme;
```

-Line 101

```
if ((ptr = (long)RDB[det0 + DET_PTR_MESH]) > VALID_PTR)
{
    tribins0 = (long)RDB[ptr + MESH_NT];
    xbins0 = (long)RDB[ptr + MESH_N0];
    ybins0 = (long)RDB[ptr + MESH_N1];
    zbins0 = (long)RDB[ptr + MESH_N2];
}
else
{
    tribins0 = 1;
    xbins0 = 1;
    ybins0 = 1;
}
```

```

    zbins0 = 1;
}

```

-Line 150

```

for (trib0 = 0; trib0 < tribins0; trib0++)
for (tb0 = 0; tb0 < tbins0; tb0++)
{
    /* Get index */

    idx0 = DetIdx(det0, eb0, ub0, cb0, mb0, lb0, zb0, yb0, xb0, trib0, tb0);

```

“detbin.c”

-Line 49

```

/* Mesh bins */

if ((ptr = (long)RDB[det + DET_PTR_MESH]) > VALID_PTR)
    ni = (long)(RDB[ptr + MESH_NT]*RDB[ptr + MESH_N0]*RDB[ptr + MESH_N1]*RDB[ptr
+ MESH_N2]);
else
    ni = 1;

```

“detidx.c”

-Line 22

```

long DetIdx(long det, long ebin, long ubin, long cbin, long mbin, long lbin,
    long zbin, long ybin, long xbin, long tribin, long tbin)
{
    long ne, nu, nc, nm, nl, nz, ny, nx, ntri, nt, nmax, idx, ptr;

    /* Get number of bins */

    ne = (long)RDB[det + DET_N_EBINS];
    nu = (long)RDB[det + DET_N_UBINS];
    nc = (long)RDB[det + DET_N_CBINS];
    nm = (long)RDB[det + DET_N_MBINS];
    nl = (long)RDB[det + DET_N_LBINS];
    nt = (long)RDB[det + DET_N_TBINS];

    /* Mesh bins */

    if ((ptr = (long)RDB[det + DET_PTR_MESH]) > VALID_PTR)
    {
        ntri = (long)RDB[ptr + MESH_NT];
        nx = (long)RDB[ptr + MESH_N0];
        ny = (long)RDB[ptr + MESH_N1];
        nz = (long)RDB[ptr + MESH_N2];
    }
}

```

```

else
{
  ntri = 1;
  nx = 1;
  ny = 1;
  nz = 1;
}

```

-Line 97

```

idx = idx + tribin*nmax;
nmax = nmax*ntri;

```

“getlatticeindexes.c”

- Line 23

```

double z0, long *tri, long *i, long *j, long *k, long type)

```

- line 39

```

else if ((type == LAT_TYPE_HXT) || (type == LAT_TYPE_HYT))
{
  if (type == LAT_TYPE_HYT)
  {
    y = x0/px;
    x = y0/py;
  }
  else
  {
    x = x0/px;
    y = y0/py;
  }
}

```

```

n1 = 2.0*x - 0.5;
n2 = -x + SQRT3*y - 0.5;
n3 = -x - SQRT3*y - 0.5;

```

```

mid1 = rint(n1);
mid2 = rint(n2);
mid3 = rint(n3);

```

```

*i = (long)floor(0.5 + (mid1 - mid2)/3.0);
*j = (long)floor(0.5 + (mid2 - mid3)/3.0);
*k = (long)rint(z0/pz);

```

```

n1 = floor(0.5 + (mid1 - mid2)/3.0);
n2 = floor(0.5 + (mid2 - mid3)/3.0);

```

```

x = x0 - (n1 + 0.5 * n2)*px;
y = y0 - (0.5*SQRT3)*py*n2;

```

```

if ( (y/x) > (1/SQRT3) )
{
    if ( x > 0) { *tri = 0;}
    else { *tri = 3;}
}
else if ( (y/x) < (-1/SQRT3) )
{
    if (x > 0) { *tri = 2;}
    else { *tri = 5;}
}
else
{
    if ( x > 0) { *tri=1;}
    else { *tri=4;}
}

```

```

return 0;
}

```

“meshindex.c”

-Line 25

```

long nt, n0, n1, n2, tri, i, j, k, idx, ptr;

```

-Line 34

```

nt = (long)RDB[msh + MESH_NT];

```

- Line 242

```

else if (((long)RDB[msh + MESH_TYPE] == MESH_TYPE_HEXXT) ||
        ((long)RDB[msh + MESH_TYPE] == MESH_TYPE_HEXYT))
{
    /******
    /**** Hexagonal mesh *****/
    /* Variables: min0 = x0, min1 = y0, max0 = pitch */
    /* Coordinate transformation relative to origin */
    x = x - min0;
    y = y - min1;
    /* Adjust if even number of cells */
    x = x - (1 - (n0 % 2))*0.5*max0;
    y = y - (1 - (n1 % 2))*0.5*max0;
    /* Adjust axial coordinate */

```

```

if (max2 - min2 > 0.0)
    z = (z - min2)/(max2 - min2);
else
    z = 0.0;

/* Check */

if ((z < 0.0) || (z >= 1.0))
    return -1;

/* Get hex indexes */

if ((long)RDB[msh + MESH_TYPE] == MESH_TYPE_HEXXT)
    GetLatticeIndexes(max0, max0, 1.0, x, y, 0.0, &tri, &i, &j, &k, LAT_TYPE_HXT);
else
    GetLatticeIndexes(max0, max0, 1.0, x, y, 0.0, &tri, &i, &j, &k, LAT_TYPE_HYT);

/* Calculate indexes */

i = i + (long)(((double)n0 - 1.0)/2.0);
j = j + (long)(((double)n1 - 1.0)/2.0);
k = (long)(z*n2);

/* Check */

if ((i < 0) || (i > n0 - 1))
    return -1;
else if ((j < 0) || (j > n1 - 1))
    return -1;
else if ((k < 0) || (k > n2 - 1))
    return -1;
else if ((tri < 0) || (tri > 5))
    Die(FUNCTION_NAME, "Indexing error");

/* Calculate index */

idx = tri + i*6 + j*6*n0 + k*6*n0*n1;

/*****
}

```

B.5 Modifying Writing of Results to Detector-Specific Output File

“detectoroutput.c”

-Line 24

```

long det0, erg, ptr, n0, n1, n2, n, m, l, i, j, k, loc0, tmp;
long ebins0, ubins0, cbins0, mbins0, lbins0, rbins0, zbins0, ybins0, xbins0, tribins0;
long tbins0, eb0, ub0, cb0, mb0, lb0, rb0, zb0, yb0, xb0, trib0, tb0, idx0;

```

```
double fetValue, fetRelUnc;  
double min0, max0, min1, max1, min2, max2, tri, x, y, x0, y0, pitch;
```

-Line 109

```
tribins0 = (long)RDB[ptr + MESH_NT];
```

-Line 115

```
tribins0 = 1;
```

-Line 166 and 250

```
for (trib0 = 0; trib0 < tribins0; trib0++)
```

-Lines 178 and 264

```
if ((long)RDB[ptr + MESH_NT] == 6)  
{fprintf(fp, "%4ld ", trib0 + 1);}
```

-Lines 197 and 268

```
        idx0 = DetIdx(det0, eb0, ub0, cb0, mb0, lb0, zb0, yb0,  
                    xb0, trib0, tb0);
```

- Line 405

```
else if (((long)RDB[ptr + MESH_TYPE] == MESH_TYPE_HEXXT) ||  
        ((long)RDB[ptr + MESH_TYPE] == MESH_TYPE_HEXYT))  
{  
    /* Print cell center coordinates */
```

```
    if ((n0 > 0) && (n1 > 0))  
    {  
        /* Get pitch */
```

```
        pitch = RDB[ptr + MESH_MAX0];
```

```
        /* Get center coordinates */
```

```
        x0 = RDB[ptr + MESH_MIN0] + (1 - (n0 % 2))*0.5*pitch;  
        y0 = RDB[ptr + MESH_MIN1] + (1 - (n1 % 2))*0.5*pitch;
```

```
        /* Print */
```

```
        fprintf(fp, "\nDET%sCOORD = [\n",  
              GetText(det0 + DET_PTR_NAME));
```

```
        /* Avoid compiler warning */
```

```
        tri = 0.0;
```

```
x = 0.0;
y = 0.0;
```

```
/* Loop over lattice */
```

```
j = -(long)((double)n1/2.0);
for (n = 0; n < n1; n++)
{
    i = -(long)((double)n0/2.0);
    for (m = 0; m < n0; m++)
    {
        k = 0;
        for (l = 0; l < 6; l++)
        {
            if ((long)RDB[ptr + MESH_TYPE] == MESH_TYPE_HEXXT)
            {
                tri = (PI/180)*60*(1-k);
                x = x0 + (i + COS60*j)*pitch + (pitch/3)*cos(tri);
                y = y0 + j*SIN60*pitch + (pitch/3)*sin(tri);
            }
            else if ((long)RDB[ptr + MESH_TYPE] == MESH_TYPE_HEXYT)
            {
                tri = (PI/180)*60*(0.5 + k);
                x = x0 + j*SIN60*pitch + (pitch/3)*cos(tri);
                y = y0 + (i + COS60*j)*pitch + (pitch/3)*sin(tri);
            }
            k++;
            /* If x,y are close to zero (due to precision), set them to zero. */
            if (abs(x) < pow(1,-10)) {x=0;}
            if (abs(y) < pow(1,-10)) {y=0;}
            fprintf(fp, "%E %E\n", x, y);
        }
        i++;
    }
    j++;
}
```

```
fprintf(fp, "];\n");
}
```

```
/* z-direction */
```

```
if (n2 > 0)
{
    fprintf(fp, "\nDET%sZ = [\n", GetText(det0 + DET_PTR_NAME));
```

```
for (n = 0; n < n2; n++)
    fprintf(fp, "%12.5E %12.5E %12.5E\n",
        ((double)n)/((double)n2)*(max2 - min2) + min2,
        ((double)n + 1.0)/((double)n2)*(max2 - min2) + min2,
        ((double)n + 0.5)/((double)n2)*(max2 - min2) + min2);
```

```
        fprintf(fp, "];\n");
    }
}
```

B.6 Minor Modifications due to Argument Changes in Functional Calls

Some source code files only required very minor modifications due to function call changes with GetLatticeIndexes(). The function needed to add additional phase space to account for the triangular mesh index. The change and the affected files are listed below.

Change:

```
GetLatticeIndexes(pr, pr, pz, x, y, z, &tri, &i, &j, &k, type);
```

Affected files (7 total), number of instances, and corresponding line numbers:

boundaryconditions.c (1) – Line 118

dfpos.c (14) – Lines 212, 333, 408, 500, 600, 680, 822, 907, 1052, 1123, 1217, 1343, 1408, and 1534

findlatticeregion.c (5) – Lines 313, 428, 430, 432, and 559

icmidx.c (1) – Line 72

nearestmeshboundary.c (2) – Lines 329 and 332

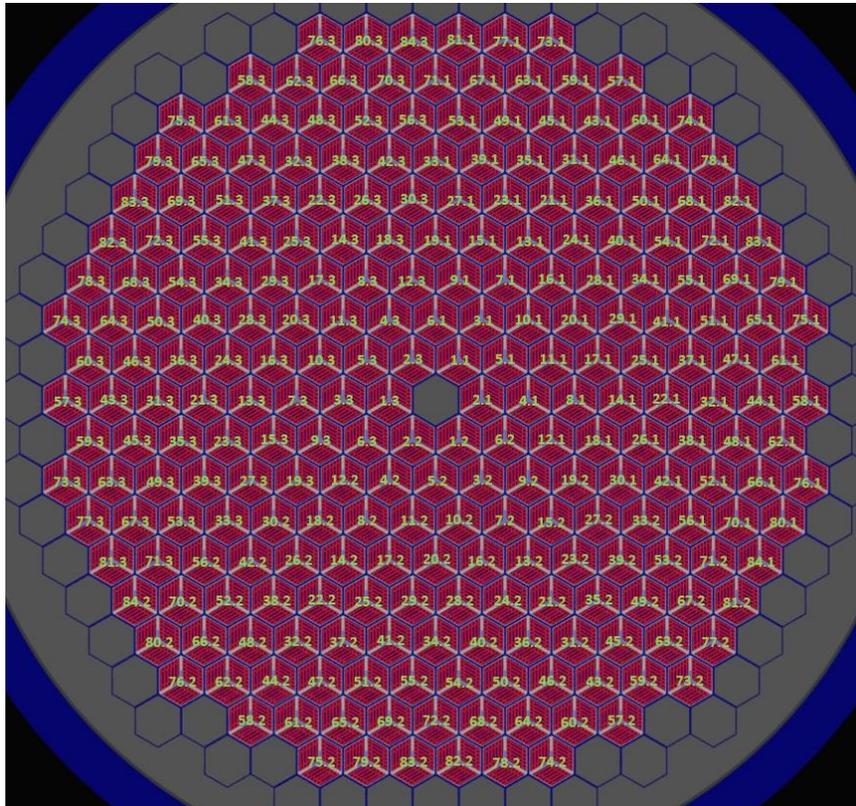
scoreufs.c (1) – Line 133

uvsfactor.c (1) – Line 156

APPENDIX C

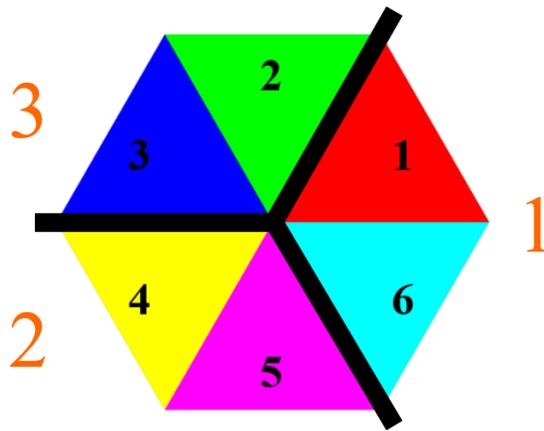
MAPPING FROM TRIANGULAR MESH TALLY BINNING TO 3D AHTR SERPENT MODEL LOCATION NOTATION

As shown in Figure 2.2 and repeated again for convenience as Supplementary Figure C.1, the 3D AHTR SERPENT model uses a one-third assembly-wise naming convention for tracking results as part of simulations. This convention generally uses an in-out radial assembly progression with a clockwise rotational symmetry progression. It is important to establish the mapping of results from the triangular mesh tally to the corresponding location in the naming convention scheme. This appendix documents that mapping.



Supplementary Figure C.1. Physical tallying mapping scheme for the first and third columns of Supplementary Table C.1. The first number corresponds to the “Hex Group”, which is for one-third core symmetric assembly locations. The second number corresponds to the “Radial Group”, which loosely starts in the positive y-direction and rotates clockwise.

As discussed in the Supplementary Figure C.1 caption, assemblies are designated at an assembly group- and rotationally indexed-level. Within each of these assemblies are three one-third assembly sections, given the name “Section Group” in the second column of Supplementary Table C.1. In the 3D AHTR SERPENT model, section progression proceeds rotationally clockwise. When combining triangular mesh elements of the “dht” tally, the results should match those of these sections. This was in fact the case in testing, and visualization of the section numbering is shown in Supplementary Figure C.2.



Supplementary Figure C.2. Triangular element mapping scheme for the sixth column of Supplementary Table C.1. Note that triangular elements (black numbers) have been combined into one-third assembly sections (designated by orange numbers and divisions shown by thick black lines). The notation correspondence is: tri section 1 for triangular regions 1 and 6, tri section 2 for triangular regions 4 and 5, and tri section 3 for triangular regions 2 and 3.

With these notations established, one can obtain the one-to-one mapping between the naming convention used in the 3D AHTR SERPENT model and the “dht” superimposed triangular mesh tally, as shown in Supplementary Table C.1. The first three columns cover the model naming conventions. The last three columns correspond to which elements of a 19x19 “dht” tally produce the matching results. Note that the “tri” index (column six) designates the *pair* of triangular regions needed to obtain the matching result. The pairing is the same as discussed in Supplementary Figure C.2. Another note is that small rounding errors can occur when comparing the results and are expected. All SERPENT tally bins report results to six significant digits and conducting a postscript

sum of two elements might differ due to rounding in the last digit from the combined binning in the code. These rounding deviations were observed when constructing the mapping shown in Supplementary Table C.1 but are not a cause for concern.

Supplementary Table C.1. One-to-one mapping of results from the 3D AHTR SERPENT model using both a physical tally and the “dht” superimposed triangular mesh tally. Note that the “tri” index corresponds to the sum of two adjacent triangular regions as discussed above.

Hex Group	Section Group	Radial Group	Tri Element (y x tri)			
1	1	1	11	10	1	
		2	9	11	3	
		3	10	9	2	
	2	2	1	11	10	2
			2	9	11	1
			3	10	9	3
	3	3	1	11	10	3
			2	9	11	2
			3	10	9	1
2	1	1	10	11	1	
		2	9	10	3	
		3	11	9	2	
	2	2	1	10	11	2
			2	9	10	1
			3	11	9	3
	3	3	1	10	11	3
			2	9	10	2
			3	11	9	1
3	1	1	12	10	1	
		2	8	12	3	
		3	10	8	2	
	2	2	1	12	10	2
			2	8	12	1
			3	10	8	3
	3	3	1	12	10	3
			2	8	12	2
			3	10	8	1
4	1	1	10	12	1	
		2	8	10	3	
		3	12	8	2	
	2	2	1	10	12	2
			2	8	10	1
			3	12	8	3
	3	3	1	10	12	3
			2	8	10	2
			3	12	8	1
5	1	1	11	11	1	
		2	8	11	3	
		3	11	8	2	
	2	2	1	11	11	2
			2	8	11	1
			3	11	8	3
	3	3	1	11	11	3
			2	8	11	2
			3	11	8	1
6	1	1	12	9	1	
		2	9	12	3	
		3	9	9	2	
	2	2	1	12	9	2
			2	9	12	1
			3	9	9	3

	3	1	12	9	3
		2	9	12	2
		3	9	9	1
7	1	1	13	10	1
		2	7	13	3
		3	10	7	2
	2	1	13	10	2
		2	7	13	1
		3	10	7	3
	3	1	13	10	3
		2	7	13	2
		3	10	7	1
8	1	1	10	13	1
		2	7	10	3
		3	13	7	2
	2	1	10	13	2
		2	7	10	1
		3	13	7	3
	3	1	10	13	3
		2	7	10	2
		3	13	7	1
9	1	1	13	9	1
		2	8	13	3
		3	9	8	2
	2	1	13	9	2
		2	8	13	1
		3	9	8	3
	3	1	13	9	3
		2	8	13	2
		3	9	8	1
10	1	1	12	11	1
		2	7	12	3
		3	11	7	2
	2	1	12	11	2
		2	7	12	1
		3	11	7	3
	3	1	12	11	3
		2	7	12	2
		3	11	7	1
11	1	1	11	12	1
		2	7	11	3
		3	12	7	2
	2	1	11	12	2
		2	7	11	1
		3	12	7	3
	3	1	11	12	3
		2	7	11	2
		3	12	7	1
12	1	1	9	13	1
		2	8	9	3
		3	13	8	2
	2	1	9	13	2
		2	8	9	1
		3	13	8	3
	3	1	9	13	3
		2	8	9	2
		3	13	8	1
13	1	1	14	10	1
		2	6	14	3
		3	10	6	2
	2	1	14	10	2
		2	6	14	1
		3	10	6	3
	3	1	14	10	3
		2	6	14	2
		3	10	6	1

14	1	1	10	14	1
		2	6	10	3
		3	14	6	2
	2	1	10	14	2
		2	6	10	1
		3	14	6	3
	3	1	10	14	3
		2	6	10	2
		3	14	6	1
15	1	1	14	9	1
		2	7	14	3
		3	9	7	2
	2	1	14	9	2
		2	7	14	1
		3	9	7	3
	3	1	14	9	3
		2	7	14	2
		3	9	7	1
16	1	1	13	11	1
		2	6	13	3
		3	11	6	2
	2	1	13	11	2
		2	6	13	1
		3	11	6	3
	3	1	13	11	3
		2	6	13	2
		3	11	6	1
17	1	1	11	13	1
		2	6	11	3
		3	13	6	2
	2	1	11	13	2
		2	6	11	1
		3	13	6	3
	3	1	11	13	3
		2	6	11	2
		3	13	6	1
18	1	1	9	14	1
		2	7	9	3
		3	14	7	2
	2	1	9	14	2
		2	7	9	1
		3	14	7	3
	3	1	9	14	3
		2	7	9	2
		3	14	7	1
19	1	1	14	8	1
		2	8	14	3
		3	8	8	2
	2	1	14	8	2
		2	8	14	1
		3	8	8	3
	3	1	14	8	3
		2	8	14	2
		3	8	8	1
20	1	1	12	12	1
		2	6	12	3
		3	12	6	2
	2	1	12	12	2
		2	6	12	1
		3	12	6	3
	3	1	12	12	3
		2	6	12	2
		3	12	6	1
21	1	1	15	10	1
		2	5	15	3
		3	10	5	2

	2	1	15	10	2
		2	5	15	1
		3	10	5	3
	3	1	15	10	3
		2	5	15	2
		3	10	5	1
22	1	1	10	15	1
		2	5	10	3
		3	15	5	2
	2	1	10	15	2
		2	5	10	1
		3	15	5	3
	3	1	10	15	3
		2	5	10	2
		3	15	5	1
23	1	1	15	9	1
		2	6	15	3
		3	9	6	2
	2	1	15	9	2
		2	6	15	1
		3	9	6	3
	3	1	15	9	3
		2	6	15	2
		3	9	6	1
24	1	1	14	11	1
		2	5	14	3
		3	11	5	2
	2	1	14	11	2
		2	5	14	1
		3	11	5	3
	3	1	14	11	3
		2	5	14	2
		3	11	5	1
25	1	1	11	14	1
		2	5	11	3
		3	14	5	2
	2	1	11	14	2
		2	5	11	1
		3	14	5	3
	3	1	11	14	3
		2	5	11	2
		3	14	5	1
26	1	1	9	15	1
		2	6	9	3
		3	15	6	2
	2	1	9	15	2
		2	6	9	1
		3	15	6	3
	3	1	9	15	3
		2	6	9	2
		3	15	6	1
27	1	1	15	8	1
		2	7	15	3
		3	8	7	2
	2	1	15	8	2
		2	7	15	1
		3	8	7	3
	3	1	15	8	3
		2	7	15	2
		3	8	7	1
28	1	1	13	12	1
		2	5	13	3
		3	12	5	2
	2	1	13	12	2
		2	5	13	1
		3	12	5	3

	3	1	13	12	3
		2	5	13	2
		3	12	5	1
29	1	1	12	13	1
		2	5	12	3
		3	13	5	2
	2	1	12	13	2
		2	5	12	1
		3	13	5	3
	3	1	12	13	3
		2	5	12	2
		3	13	5	1
30	1	1	8	15	1
		2	7	8	3
		3	15	7	2
	2	1	8	15	2
		2	7	8	1
		3	15	7	3
	3	1	8	15	3
		2	7	8	2
		3	15	7	1
31	1	1	16	10	1
		2	4	16	3
		3	10	4	2
	2	1	16	10	2
		2	4	16	1
		3	10	4	3
	3	1	16	10	3
		2	4	16	2
		3	10	4	1
32	1	1	10	16	1
		2	4	10	3
		3	16	4	2
	2	1	10	16	2
		2	4	10	1
		3	16	4	3
	3	1	10	16	3
		2	4	10	2
		3	16	4	1
33	1	1	16	7	1
		2	7	16	3
		3	7	7	2
	2	1	16	7	2
		2	7	16	1
		3	7	7	3
	3	1	16	7	3
		2	7	16	2
		3	7	7	1
34	1	1	13	13	1
		2	4	13	3
		3	13	4	2
	2	1	13	13	2
		2	4	13	1
		3	13	4	3
	3	1	13	13	3
		2	4	13	2
		3	13	4	1
35	1	1	16	9	1
		2	5	16	3
		3	9	5	2
	2	1	16	9	2
		2	5	16	1
		3	9	5	3
	3	1	16	9	3
		2	5	16	2
		3	9	5	1

36	1	1	15	11	1
		2	4	15	3
		3	11	4	2
	2	1	15	11	2
		2	4	15	1
		3	11	4	3
	3	1	15	11	3
		2	4	15	2
		3	11	4	1
37	1	1	11	15	1
		2	4	11	3
		3	15	4	2
	2	1	11	15	2
		2	4	11	1
		3	15	4	3
	3	1	11	15	3
		2	4	11	2
		3	15	4	1
38	1	1	9	16	1
		2	5	9	3
		3	16	5	2
	2	1	9	16	2
		2	5	9	1
		3	16	5	3
	3	1	9	16	3
		2	5	9	2
		3	16	5	1
39	1	1	16	8	1
		2	6	16	3
		3	8	6	2
	2	1	16	8	2
		2	6	16	1
		3	8	6	3
	3	1	16	8	3
		2	6	16	2
		3	8	6	1
40	1	1	14	12	1
		2	4	14	3
		3	12	4	2
	2	1	14	12	2
		2	4	14	1
		3	12	4	3
	3	1	14	12	3
		2	4	14	2
		3	12	4	1
41	1	1	12	14	1
		2	4	12	3
		3	14	4	2
	2	1	12	14	2
		2	4	12	1
		3	14	4	3
	3	1	12	14	3
		2	4	12	2
		3	14	4	1
42	1	1	8	16	1
		2	6	8	3
		3	16	6	2
	2	1	8	16	2
		2	6	8	1
		3	16	6	3
	3	1	8	16	3
		2	6	8	2
		3	16	6	1
43	1	1	17	10	1
		2	3	17	3
		3	10	3	2

	2	1	17	10	2
		2	3	17	1
		3	10	3	3
	3	1	17	10	3
		2	3	17	2
		3	10	3	1
44	1	1	10	17	1
		2	3	10	3
		3	17	3	2
	2	1	10	17	2
		2	3	10	1
		3	17	3	3
	3	1	10	17	3
		2	3	10	2
		3	17	3	1
45	1	1	17	9	1
		2	4	17	3
		3	9	4	2
	2	1	17	9	2
		2	4	17	1
		3	9	4	3
	3	1	17	9	3
		2	4	17	2
		3	9	4	1
46	1	1	16	11	1
		2	3	16	3
		3	11	3	2
	2	1	16	11	2
		2	3	16	1
		3	11	3	3
	3	1	16	11	3
		2	3	16	2
		3	11	3	1
47	1	1	11	16	1
		2	3	11	3
		3	16	3	2
	2	1	11	16	2
		2	3	11	1
		3	16	3	3
	3	1	11	16	3
		2	3	11	2
		3	16	3	1
48	1	1	9	17	1
		2	4	9	3
		3	17	4	2
	2	1	9	17	2
		2	4	9	1
		3	17	4	3
	3	1	9	17	3
		2	4	9	2
		3	17	4	1
49	1	1	17	8	1
		2	5	17	3
		3	8	5	2
	2	1	17	8	2
		2	5	17	1
		3	8	5	3
	3	1	17	8	3
		2	5	17	2
		3	8	5	1
50	1	1	15	12	1
		2	3	15	3
		3	12	3	2
	2	1	15	12	2
		2	3	15	1
		3	12	3	3

	3	1	15	12	3
		2	3	15	2
		3	12	3	1
51	1	1	12	15	1
		2	3	12	3
		3	15	3	2
	2	1	12	15	2
		2	3	12	1
		3	15	3	3
	3	1	12	15	3
		2	3	12	2
		3	15	3	1
52	1	1	8	17	1
		2	5	8	3
		3	17	5	2
	2	1	8	17	2
		2	5	8	1
		3	17	5	3
	3	1	8	17	3
		2	5	8	2
		3	17	5	1
53	1	1	17	7	1
		2	6	17	3
		3	7	6	2
	2	1	17	7	2
		2	6	17	1
		3	7	6	3
	3	1	17	7	3
		2	6	17	2
		3	7	6	1
54	1	1	14	13	1
		2	3	14	3
		3	13	3	2
	2	1	14	13	2
		2	3	14	1
		3	13	3	3
	3	1	14	13	3
		2	3	14	2
		3	13	3	1
55	1	1	13	14	1
		2	3	13	3
		3	14	3	2
	2	1	13	14	2
		2	3	13	1
		3	14	3	3
	3	1	13	14	3
		2	3	13	2
		3	14	3	1
56	1	1	7	17	1
		2	6	7	3
		3	17	6	2
	2	1	7	17	2
		2	6	7	1
		3	17	6	3
	3	1	7	17	3
		2	6	7	2
		3	17	6	1
57	1	1	18	10	1
		2	2	18	3
		3	10	2	2
	2	1	18	10	2
		2	2	18	1
		3	10	2	3
	3	1	18	10	3
		2	2	18	2
		3	10	2	1

58	1	1	10	18	1
		2	2	10	3
		3	18	2	2
	2	1	10	18	2
		2	2	10	1
		3	18	2	3
	3	1	10	18	3
		2	2	10	2
		3	18	2	1
59	1	1	18	9	1
		2	3	18	3
		3	9	3	2
	2	1	18	9	2
		2	3	18	1
		3	9	3	3
	3	1	18	9	3
		2	3	18	2
		3	9	3	1
60	1	1	17	11	1
		2	2	17	3
		3	11	2	2
	2	1	17	11	2
		2	2	17	1
		3	11	2	3
	3	1	17	11	3
		2	2	17	2
		3	11	2	1
61	1	1	11	17	1
		2	2	11	3
		3	17	2	2
	2	1	11	17	2
		2	2	11	1
		3	17	2	3
	3	1	11	17	3
		2	2	11	2
		3	17	2	1
62	1	1	9	18	1
		2	3	9	3
		3	18	3	2
	2	1	9	18	2
		2	3	9	1
		3	18	3	3
	3	1	9	18	3
		2	3	9	2
		3	18	3	1
63	1	1	18	8	1
		2	4	18	3
		3	8	4	2
	2	1	18	8	2
		2	4	18	1
		3	8	4	3
	3	1	18	8	3
		2	4	18	2
		3	8	4	1
64	1	1	16	12	1
		2	2	16	3
		3	12	2	2
	2	1	16	12	2
		2	2	16	1
		3	12	2	3
	3	1	16	12	3
		2	2	16	2
		3	12	2	1
65	1	1	12	16	1
		2	2	12	3
		3	16	2	2

	2	1	12	16	2
		2	2	12	1
		3	16	2	3
	3	1	12	16	3
		2	2	12	2
		3	16	2	1
66	1	1	8	18	1
		2	4	8	3
		3	18	4	2
	2	1	8	18	2
		2	4	8	1
		3	18	4	3
	3	1	8	18	3
		2	4	8	2
		3	18	4	1
67	1	1	18	7	1
		2	5	18	3
		3	7	5	2
	2	1	18	7	2
		2	5	18	1
		3	7	5	3
	3	1	18	7	3
		2	5	18	2
		3	7	5	1
68	1	1	15	13	1
		2	2	15	3
		3	13	2	2
	2	1	15	13	2
		2	2	15	1
		3	13	2	3
	3	1	15	13	3
		2	2	15	2
		3	13	2	1
69	1	1	13	15	1
		2	2	13	3
		3	15	2	2
	2	1	13	15	2
		2	2	13	1
		3	15	2	3
	3	1	13	15	3
		2	2	13	2
		3	15	2	1
70	1	1	7	18	1
		2	5	7	3
		3	18	5	2
	2	1	7	18	2
		2	5	7	1
		3	18	5	3
	3	1	7	18	3
		2	5	7	2
		3	18	5	1
71	1	1	18	6	1
		2	6	18	3
		3	6	6	2
	2	1	18	6	2
		2	6	18	1
		3	6	6	3
	3	1	18	6	3
		2	6	18	2
		3	6	6	1
72	1	1	14	14	1
		2	2	14	3
		3	14	2	2
	2	1	14	14	2
		2	2	14	1
		3	14	2	3

	3	1	14	14	3
		2	2	14	2
		3	14	2	1
73	1	1	19	8	1
		2	3	19	3
		3	8	3	2
	2	1	19	8	2
		2	3	19	1
		3	8	3	3
	3	1	19	8	3
		2	3	19	2
		3	8	3	1
74	1	1	17	12	1
		2	1	17	3
		3	12	1	2
	2	1	17	12	2
		2	1	17	1
		3	12	1	3
	3	1	17	12	3
		2	1	17	2
		3	12	1	1
75	1	1	12	17	1
		2	1	12	3
		3	17	1	2
	2	1	12	17	2
		2	1	12	1
		3	17	1	3
	3	1	12	17	3
		2	1	12	2
		3	17	1	1
76	1	1	8	19	1
		2	3	8	3
		3	19	3	2
	2	1	8	19	2
		2	3	8	1
		3	19	3	3
	3	1	8	19	3
		2	3	8	2
		3	19	3	1
77	1	1	19	7	1
		2	4	19	3
		3	7	4	2
	2	1	19	7	2
		2	4	19	1
		3	7	4	3
	3	1	19	7	3
		2	4	19	2
		3	7	4	1
78	1	1	16	13	1
		2	1	16	3
		3	13	1	2
	2	1	16	13	2
		2	1	16	1
		3	13	1	3
	3	1	16	13	3
		2	1	16	2
		3	13	1	1
79	1	1	13	16	1
		2	1	13	3
		3	16	1	2
	2	1	13	16	2
		2	1	13	1
		3	16	1	3
	3	1	13	16	3
		2	1	13	2
		3	16	1	1

80	1	1	7	19	1
		2	4	7	3
		3	19	4	2
	2	1	7	19	2
		2	4	7	1
		3	19	4	3
	3	1	7	19	3
		2	4	7	2
		3	19	4	1
81	1	1	19	6	1
		2	5	19	3
		3	6	5	2
	2	1	19	6	2
		2	5	19	1
		3	6	5	3
	3	1	19	6	3
		2	5	19	2
		3	6	5	1
82	1	1	15	14	1
		2	1	15	3
		3	14	1	2
	2	1	15	14	2
		2	1	15	1
		3	14	1	3
	3	1	15	14	3
		2	1	15	2
		3	14	1	1
83	1	1	14	15	1
		2	1	14	3
		3	15	1	2
	2	1	14	15	2
		2	1	14	1
		3	15	1	3
	3	1	14	15	3
		2	1	14	2
		3	15	1	1
84	1	1	6	19	1
		2	5	6	3
		3	19	5	2
	2	1	6	19	2
		2	5	6	1
		3	19	5	3
	3	1	6	19	3
		2	5	6	2
		3	19	5	1

APPENDIX D

SUPPLEMENTARY RADIAL POWER PROFILES

These results supplement those presented in Chapter 6. A complete table of results and all radial profiles are presented here for both the full- and quarter-density BP design considerations.

D.1 Full Density Burnable Poison Loading Control Blade Insertion

Additional results are provided in Supplementary Table D.1 from those already shown in Chapter 6. The 23 one-third radially symmetric assembly groups of Scheme 4 from Figure 6.9 were inserted sequentially according to the order given by the column labelled “Assembly Group Inserted” of Supplementary Table D.1.

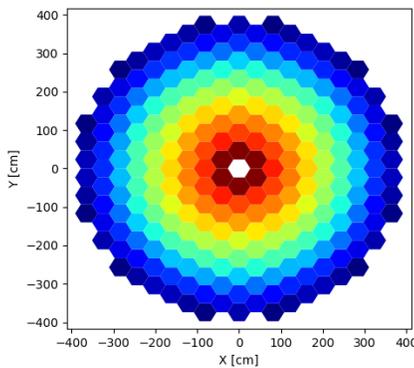
A note on the insertion order is that one-sixth radially symmetric assembly group pairs were inserted sequentially (as noted by the gray rows of Supplementary Table D.1). Pairwise insertions were conducted serially but are ordered together. In an actual implementation, these pairs could be moved as a single grouping instead of two distinct groups. The advantage of this is the ability to better maintain power symmetry and to reduce power peaking. When the first individual group of a symmetric pair is inserted, the power peaks in the location of the non-inserted second group. Pairwise group insertion (i.e. moving six CBs at once instead of just three) could lead to improved radial performance. The impact would depend on the pair in question, where power would be radially pushed outside of the group pair, and other cycle-related parameters. The axial impact of this type of pairwise movement was not considered, but it is likely that moving more CBs simultaneously would have a more significant impact on the AO and a higher axial PPF than from only moving three CBs at a time.

Supplementary Table D.1. Complete results to accompany the CB insertion analysis for Scheme 4. Rows highlighted in gray indicate one-sixth radially symmetric assembly group pairs.

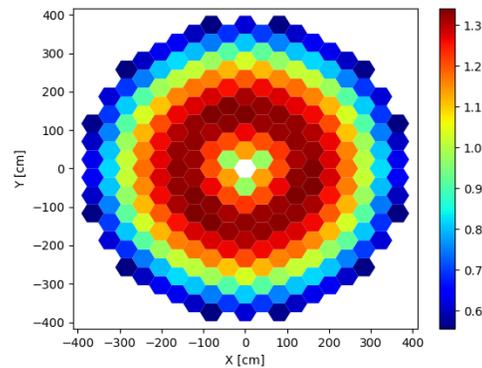
Groups Inserted	Assembly Group Inserted	k_{eff}	σ [pcm]	Δk [pcm]	Maximum Assembly Group	Maximum Assembly PPF	Maximum Section Group	Maximum Section PPF
0	(none)	1.07532	5	-	2	1.880	1.2	2.050
1	1	1.06986	5	-546	8	1.342	18.2	1.350
2	17	1.06574	5	-412	15	1.346	15.1	1.353
3	18	1.06200	5	-374	13	1.344	13.3	1.349
4	13	1.05783	5	-417	34	1.206	40.3	1.217
5	53	1.05492	5	-291	34	1.287	34.1	1.291
6	54	1.05145	5	-347	22	1.225	22.1	1.244
7	22	1.04793	5	-352	10	1.282	7.2	1.295
8	9	1.04454	5	-339	36	1.256	31.2	1.271
9	10	1.04206	5	-248	31	1.301	31.2	1.308
10	35	1.03843	5	-363	51	1.307	51.3	1.315
11	36	1.03566	5	-277	47	1.417	32.3	1.440
12	47	1.03125	5	-441	52	1.351	52.2	1.365
13	48	1.02778	5	-347	42	1.203	33.2	1.230
14	27	1.02460	5	-318	41	1.297	41.1	1.324
15	28	1.02125	5	-335	43	1.191	42.3	1.218
16	41	1.01831	5	-294	56	1.292	45.1	1.321
17	42	1.01497	5	-334	43	1.398	43.3	1.411
18	43	1.01101	5	-396	2	1.300	2.2	1.370
19	4	1.00826	5	-275	69	1.309	69.1	1.330
20	69	1.00500	5	-326	71	1.461	70.2	1.491
21	70	1.00105	5	-395	63	1.319	63.2	1.338
22	63	0.99760	5	-345	64	1.321	64.2	1.357
23	64	0.99431	5	-329	2	1.309	2.2	1.398

Visualization of the radial power profiles for the CB insertions of Scheme 4 can be seen in Supplementary Table D.2. Note that the scale differs for each profile, so colors do not necessarily signify the same peaking factor from one case to another.

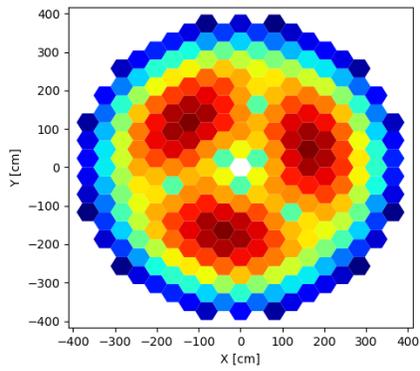
Supplementary Table D.2. Radial power profiles for Scheme 4.



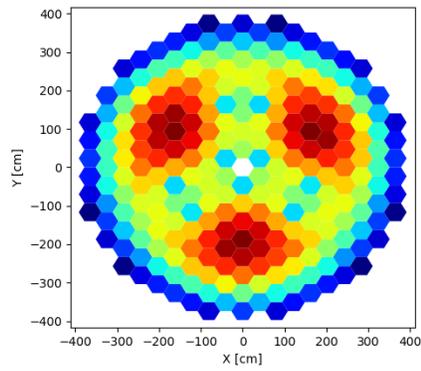
Case 0. No assembly groups inserted.



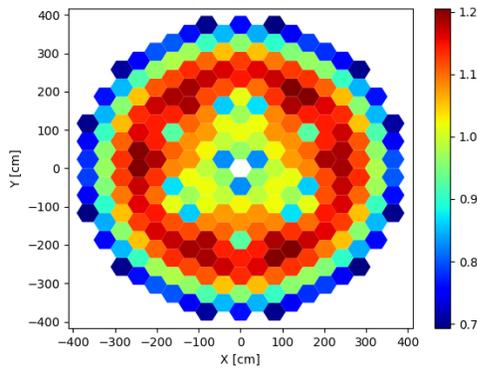
Case 1. Assembly Group 1 inserted.



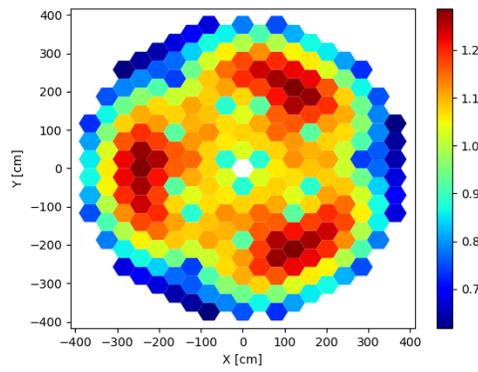
Case 2. Assembly Group 17 inserted.



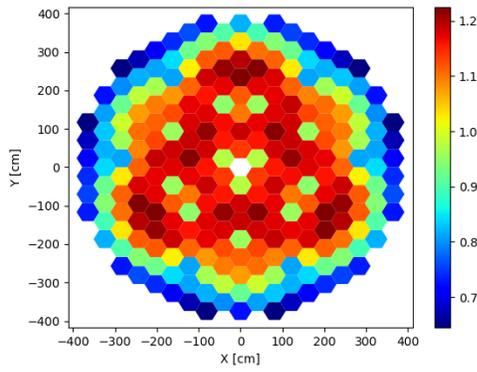
Case 3. Assembly Group 18 inserted.



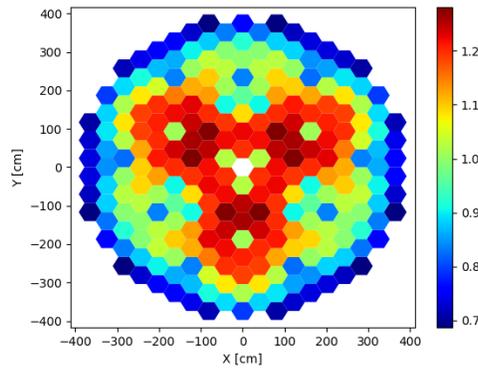
Case 4. Assembly Group 13 inserted.



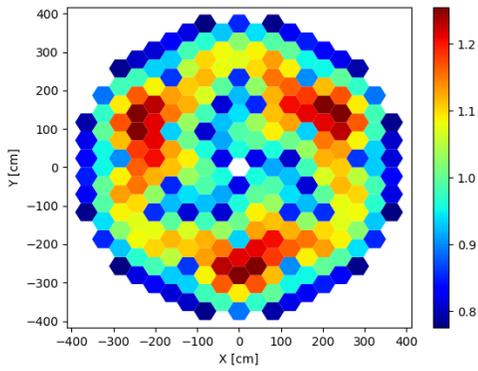
Case 5. Assembly Group 53 inserted.



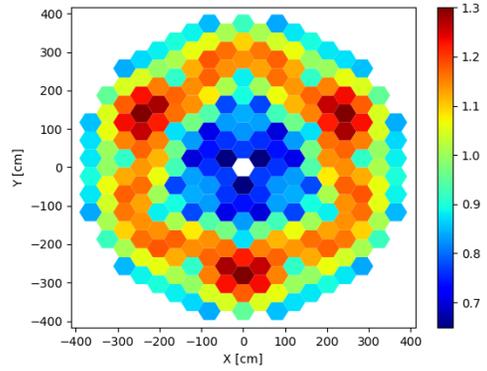
Case 6. Assembly Group 54 inserted.



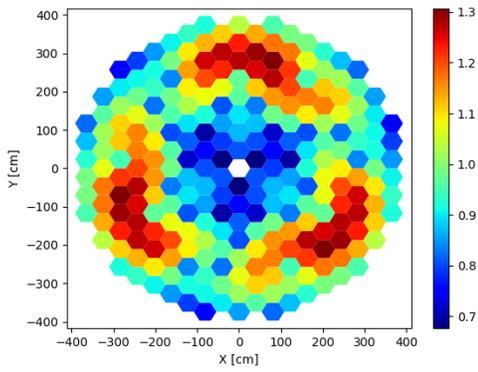
Case 7. Assembly Group 22 inserted.



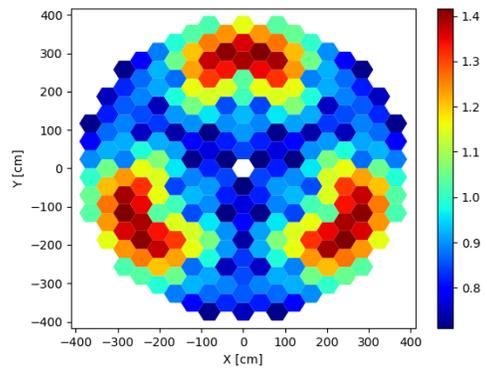
Case 8. Assembly Group 9 inserted.



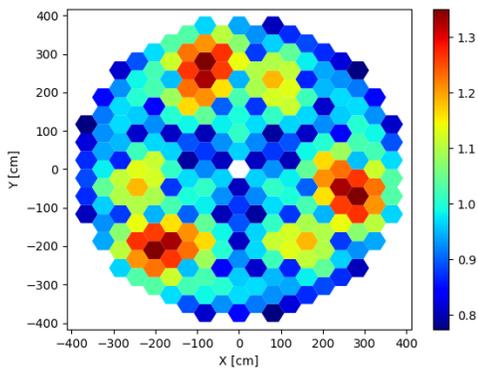
Case 9. Assembly Group 10 inserted.



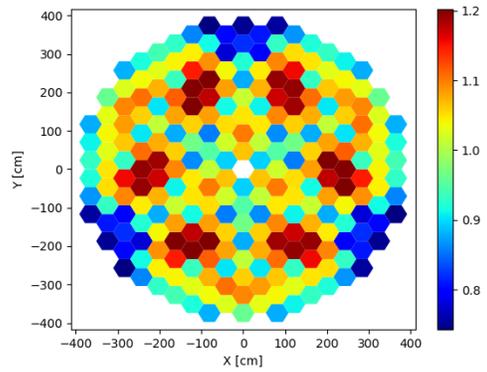
Case 10. Assembly Group 35 inserted.



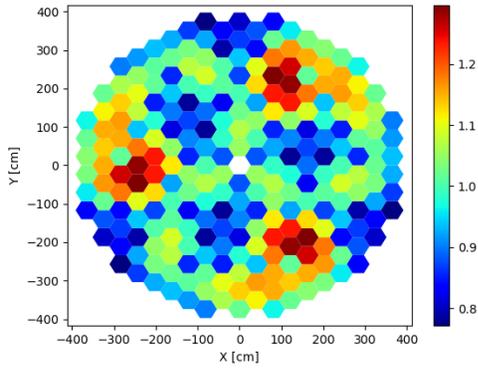
Case 11. Assembly Group 36 inserted.



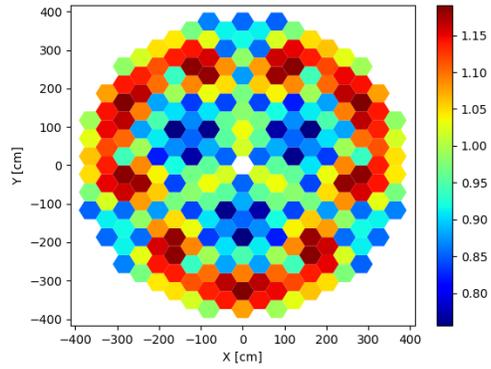
Case 12. Assembly Group 47 inserted.



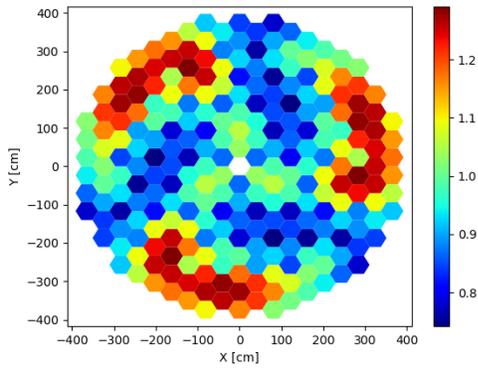
Case 13. Assembly Group 48 inserted.



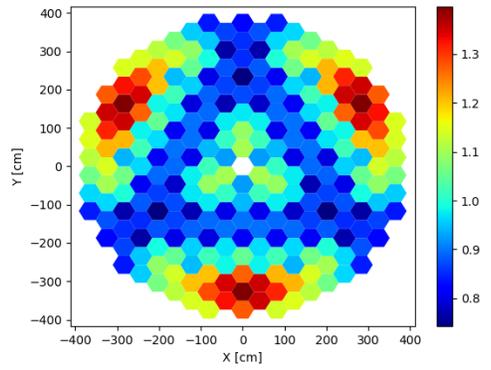
Case 14. Assembly Group 27 inserted.



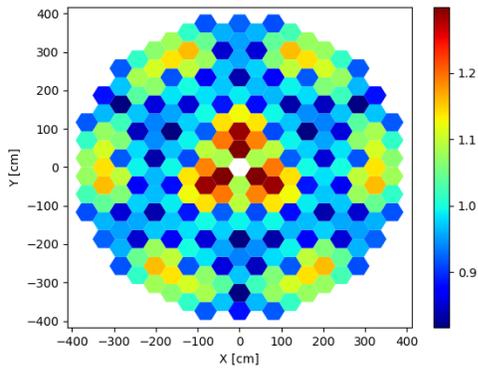
Case 15. Assembly Group 28 inserted.



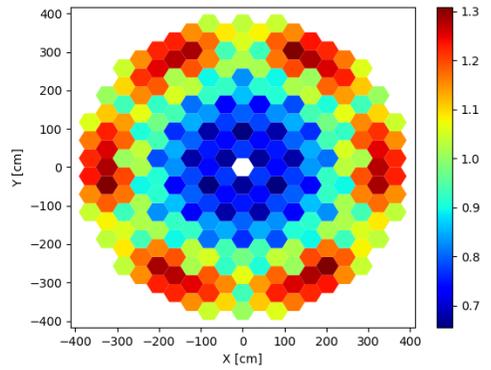
Case 16. Assembly Group 41 inserted.



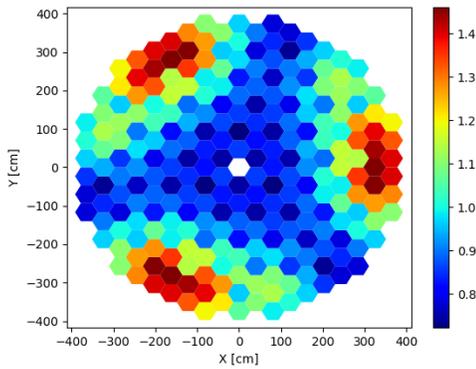
Case 17. Assembly Group 42 inserted.



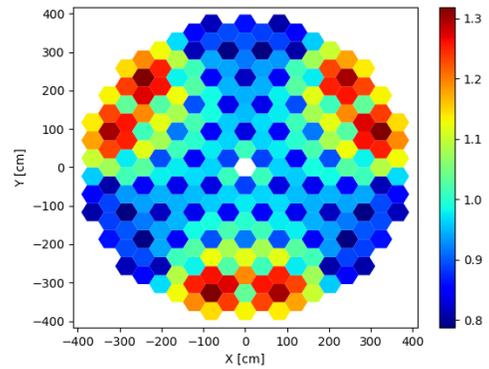
Case 18. Assembly Group 43 inserted.



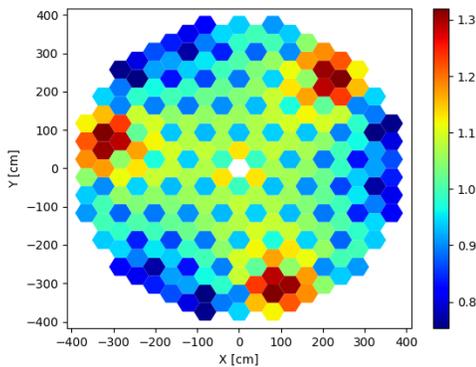
Case 19. Assembly Group 4 inserted.



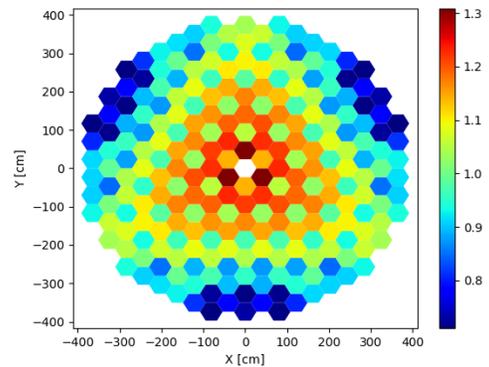
Case 20. Assembly Group 69 inserted.



Case 21. Assembly Group 70 inserted.



Case 22. Assembly Group 63 inserted.



Case 23. Assembly Group 64 inserted.

D.2 Quarter Density Burnable Poison Loading Control Blade Insertion

Recall that the insertion schedule for the quarter density BP loading was generated using an automated search procedure of inserting CBs into the assembly group with the highest peaking factor. There were no other restrictions – placement could be next to previously-inserted groupings and no regard for one-sixth symmetric pairings. This is why there are no highlighted rows in Supplementary Table D.3 like there are in Supplementary Table D.1. Additionally, one might note that there are occasional discrepancies where the entry in the column “Maximum Assembly Group” does not match the entry in the next row of the column “Assembly Group Inserted”. There are two reasons for this. First, especially toward the end of the search when most of the CB groups are inserted, the assembly group with the highest power may already have its CBs inserted. In this case, the search procedure looks for the non-inserted assembly group with the highest PPF. Second,

simulations were conducted in two stages. The first stages ran a quick simulation to find the where the next highest PPF group is located. The second stage ran a longer transport simulation to better converge on the power profile for the results shown in Supplementary Table D.3 and the power profiles shown in Supplementary Table D.4. Occasionally, the more accurate results of the longer simulation (column “Maximum Assembly Group”) could differ from those of the shorter simulation (column “Assembly Group Inserted”). This is not too concerning though, since the two assembly groups had very similar peaking factors; especially in cases where the two assembly groups are pairwise one-sixth radially symmetric groupings. In this case, selecting one group over the other should have no meaningful impact since the two have equal peaking factors and simply only one could be selected by the search procedure. In other cases, the peaking factors in non-symmetric competing groups should still have been very close and the difference in selection should not have drastically impacted the results.

Supplementary Table D.3. Complete results to accompany the quarter density europa CB insertion scheme.

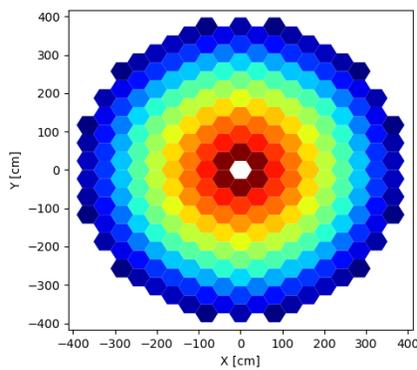
Groups Inserted	Assembly Group Inserted	k_{eff}	σ [pcm]	Δk [pcm]	Maximum Assembly Group	Maximum Assembly PPF	Maximum Section Group	Maximum Section PPF
0	-	1.22208	5	-	2	1.878	2.2	2.044
1	1	1.21546	5	-662	8	1.363	8.3	1.374
2	20	1.21055	4	-491	19	1.346	19.1	1.355
3	30	1.20555	4	-500	13	1.283	13.1	1.295
4	13	1.20105	4	-450	25	1.270	25.3	1.282
5	25	1.19650	5	-455	35	1.199	35.2	1.209
6	12	1.19304	4	-346	36	1.264	36.1	1.278
7	36	1.18873	4	-431	56	1.272	38.3	1.299
8	38	1.18435	4	-438	39	1.266	39.1	1.283
9	39	1.18017	4	-418	55	1.340	55.2	1.348
10	55	1.17539	5	-478	56	1.154	45.2	1.169
11	56	1.17191	5	-348	5	1.217	5.1	1.245
12	11	1.16841	5	-350	45	1.241	31.1	1.273
13	31	1.16452	5	-389	47	1.315	47.2	1.325
14	47	1.16020	5	-432	40	1.197	50.3	1.232
15	3	1.15708	5	-312	54	1.280	68.2	1.322
16	68	1.15305	5	-403	67	1.301	49.1	1.325
17	63	1.14936	4	-369	48	1.237	48.3	1.259
18	14	1.14567	5	-369	48	1.267	48.3	1.308
19	40	1.14218	5	-349	66	1.408	66.2	1.432
20	48	1.13748	5	-470	15	1.238	15.1	1.252
21	53	1.13397	4	-351	15	1.262	15.2	1.302

22	15	1.13016	5	-381	69	1.308	69.3	1.319
23	69	1.12605	5	-411	2	1.190	2.2	1.247
24	2	1.12381	5	-224	43	1.302	60.1	1.323
25	60	1.12044	5	-337	66	1.363	70.1	1.388
26	66	1.11610	5	-434	41	1.287	34.2	1.325
27	41	1.11216	5	-394	35	1.274	35.1	1.300
28	35	1.10805	5	-411	16	1.202	16.3	1.216
29	19	1.10502	5	-303	61	1.258	61.3	1.304
30	61	1.10175	4	-327	81	1.288	71.1	1.316
31	81	1.09866	5	-309	28	1.395	28.2	1.401
32	28	1.09396	5	-470	50	1.203	50.1	1.249
33	54	1.09075	5	-321	26	1.260	26.3	1.275
34	42	1.08714	4	-361	43	1.271	43.1	1.284
35	43	1.08369	5	-345	4	1.352	10.1	1.386
36	8	1.07966	5	-403	32	1.287	32.1	1.315
37	32	1.07590	5	-376	7	1.290	7.3	1.339
38	16	1.07223	5	-367	70	1.324	70.3	1.351
39	70	1.06879	5	-344	82	1.433	82.3	1.462
40	82	1.06512	5	-367	9	1.244	7.2	1.278
41	45	1.06185	5	-327	6	1.382	6.1	1.410
42	4	1.05833	5	-352	65	1.326	65.2	1.337
43	65	1.05466	5	-367	23	1.274	23.2	1.288
44	67	1.05163	5	-303	7	1.336	7.2	1.369
45	7	1.04779	4	-384	64	1.393	64.2	1.429
46	64	1.04412	5	-367	62	1.451	62.3	1.506
47	22	1.03970	5	-442	62	1.429	62.3	1.504
48	62	1.03617	5	-353	23	1.399	23.2	1.416
49	23	1.03201	5	-416	29	1.341	29.1	1.344
50	29	1.02807	5	-394	59	1.214	59.1	1.298
51	50	1.02525	5	-282	6	1.319	6.1	1.347
52	84	1.02248	5	-277	6	1.476	6.1	1.511
53	5	1.01900	5	-348	79	1.346	79.1	1.396
54	59	1.01627	5	-273	79	1.460	79.1	1.502
55	83	1.01307	4	-320	18	1.479	26.2	1.527
56	26	1.00856	5	-451	27	1.359	27.1	1.393
57	33	1.00482	5	-374	51	1.415	51.2	1.452
58	51	1.00085	4	-397	9	1.484	9.2	1.526
59	9	0.99693	5	-392	58	1.375	58.3	1.407
60	58	0.99418	5	-275	21	1.455	21.3	1.512
61	24	0.98993	4	-425	77	1.311	77.1	1.381
62	49	0.98663	4	-330	74	1.261	74.1	1.289
63	74	0.98430	5	-233	17	1.344	17.2	1.352
64	17	0.98075	5	-355	80	1.428	80.3	1.493
65	80	0.97766	5	-309	6	1.308	6.3	1.314
66	6	0.97494	5	-272	79	1.417	79.3	1.468
67	79	0.97190	5	-304	73	1.397	77.1	1.463
68	73	0.96925	5	-265	37	1.370	37.1	1.376
69	34	0.96568	5	-357	52	1.451	52.1	1.459
70	27	0.96174	5	-394	52	1.469	44.2	1.497
71	52	0.95788	5	-386	46	1.506	46.2	1.523
72	46	0.95393	5	-395	37	1.522	37.3	1.549
73	37	0.94955	5	-438	10	1.462	10.2	1.467
74	10	0.94611	5	-344	44	1.380	71.3	1.395

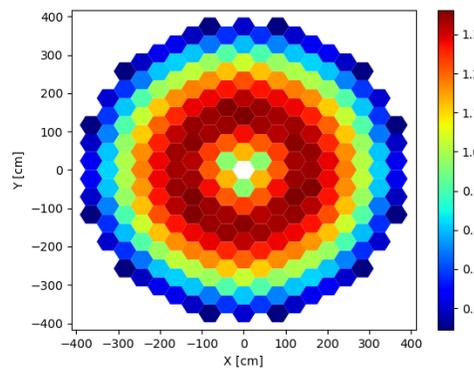
75	71	0.94281	5	-330	72	1.476	72.1	1.488
76	72	0.93906	5	-375	18	1.493	18.2	1.502
77	18	0.93518	5	-388	44	1.548	44.3	1.553
78	21	0.93141	5	-377	44	1.861	44.3	1.872
79	44	0.92597	4	-544	57	1.589	57.1	1.658
80	57	0.92291	4	-306	2	1.350	2.2	1.513
81	76	0.92113	5	-178	2	1.523	2.2	1.705
82	78	0.91920	5	-193	2	1.700	2.2	1.905
83	77	0.91762	5	-158	1	1.880	2.2	2.126
84	75	0.91616	5	-146	1	2.047	2.2	2.297

The radial power profiles for the automated CB insertion schedule search procedure can be seen in Supplementary Table D.4. Note that the scale differs for each profile, so colors do not necessarily signify the same peaking factor from one case to another.

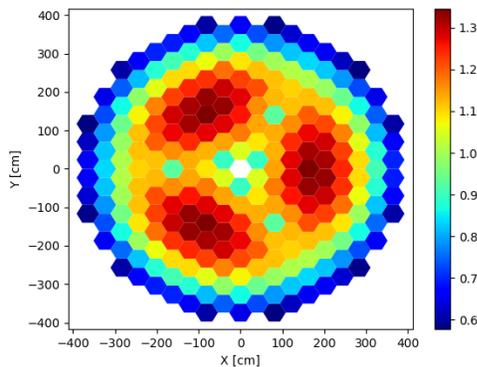
Supplementary Table D.4. Radial power profiles for the quarter density europa automated CB insertion scheme.



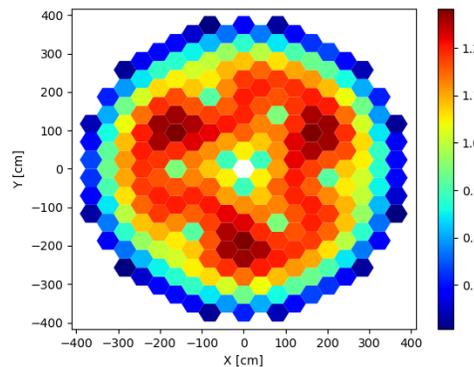
Case 0. No Assembly Groups Inserted.



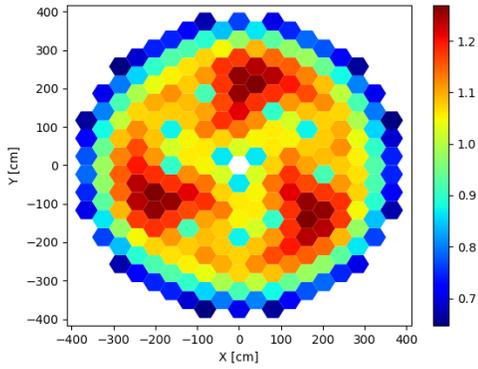
Case 1. Assembly Group 1 inserted.



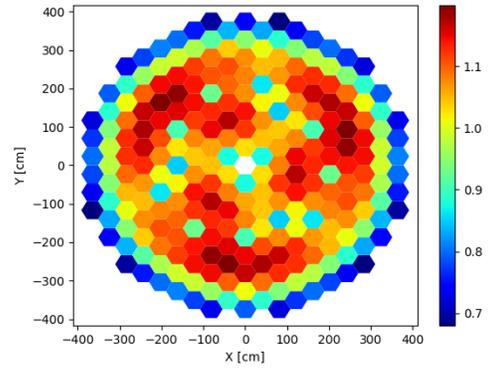
Case 2. Assembly Group 20 inserted.



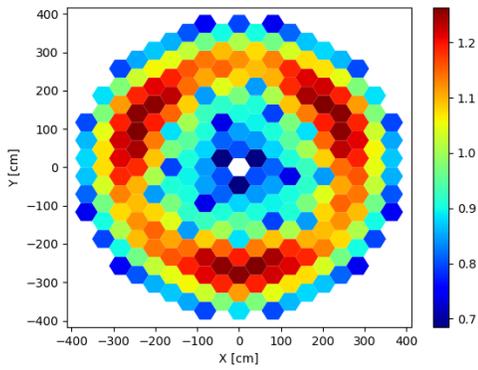
Case 3. Assembly Group 30 inserted.



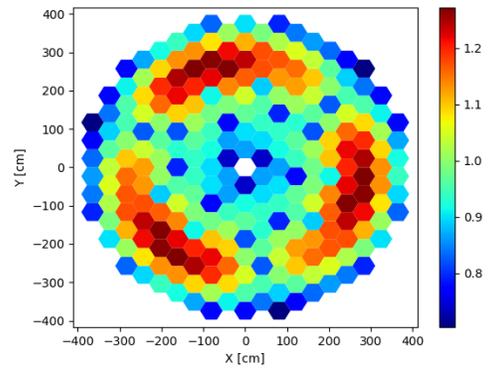
Case 4. Assembly Group 13 inserted.



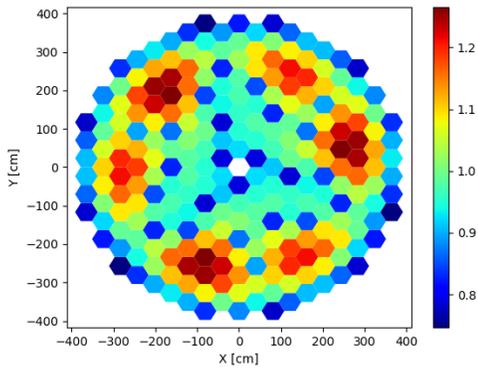
Case 5. Assembly Group 25 inserted.



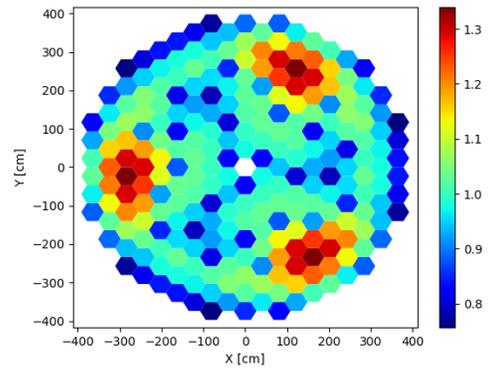
Case 6. Assembly Group 12 inserted.



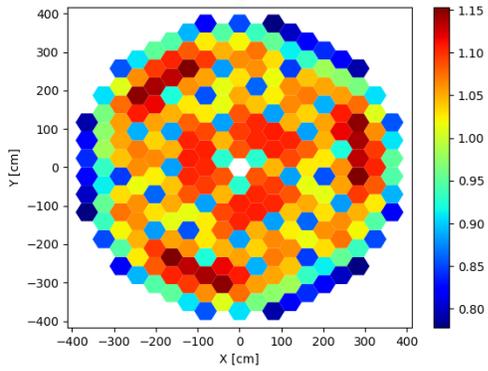
Case 7. Assembly Group 36 inserted.



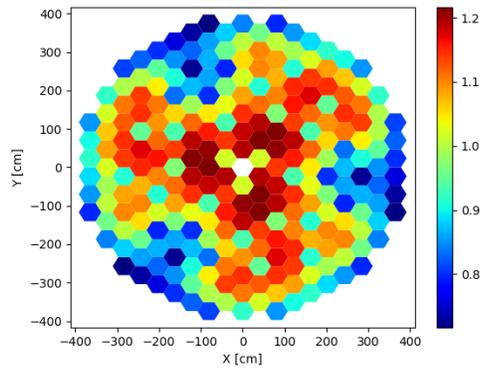
Case 8. Assembly Group 38 inserted.



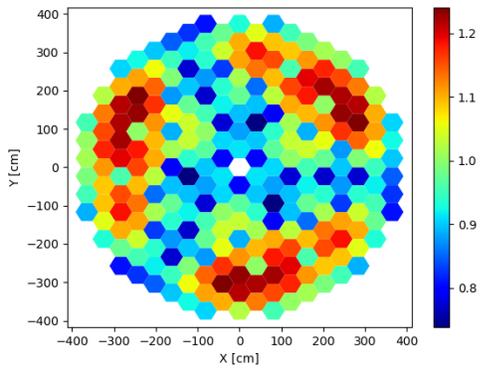
Case 8. Assembly Group 39 inserted.



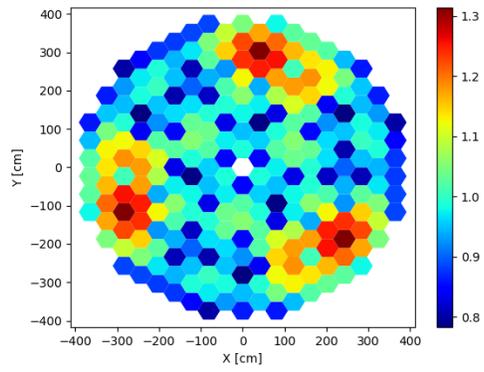
Case 10. Assembly Group 55 inserted.



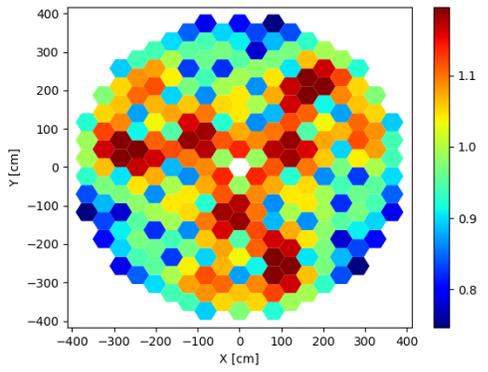
Case 11. Assembly Group 56 inserted.



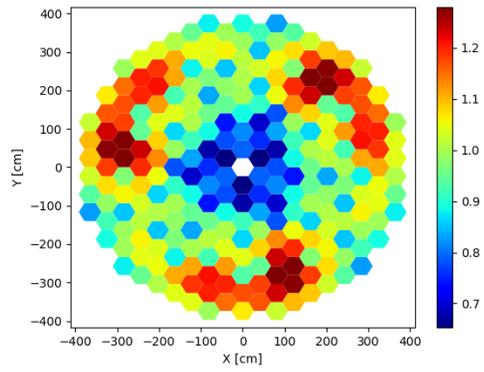
Case 12. Assembly Group 11 inserted.



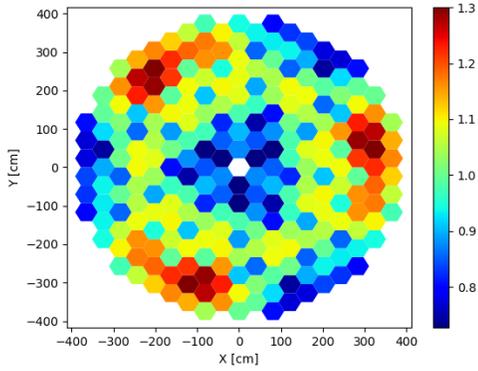
Case 13. Assembly Group 31 inserted.



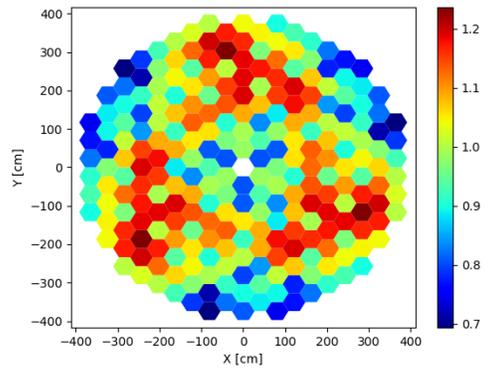
Case 14. Assembly Group 47 inserted.



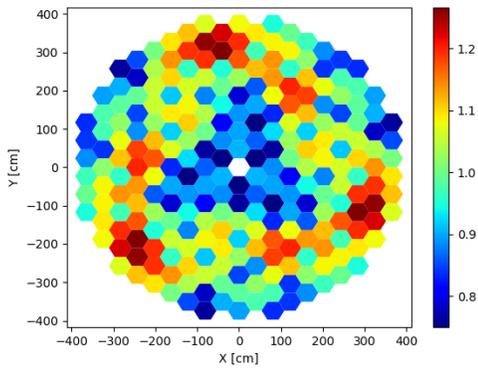
Case 15. Assembly Group 3 inserted.



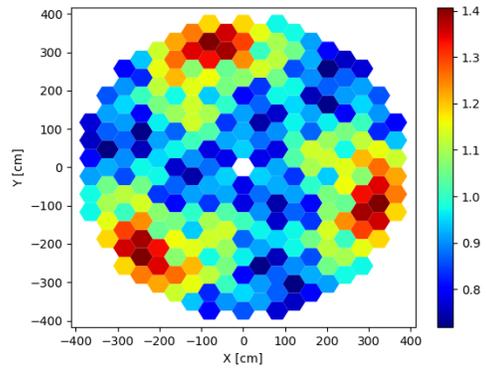
Case 16. Assembly Group 68 inserted.



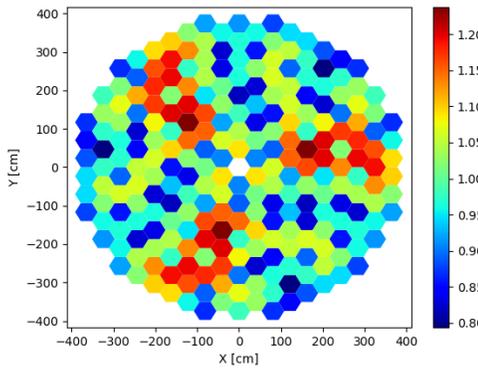
Case 17. Assembly Group 63 inserted.



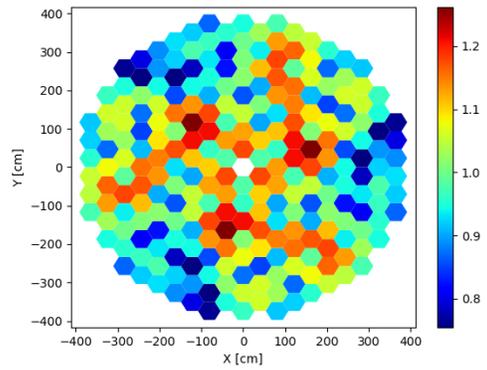
Case 18. Assembly Group 14 inserted.



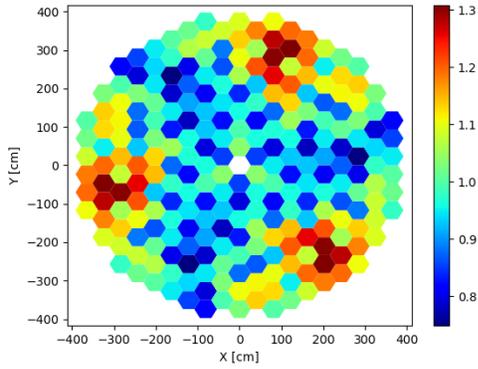
Case 19. Assembly Group 40 inserted.



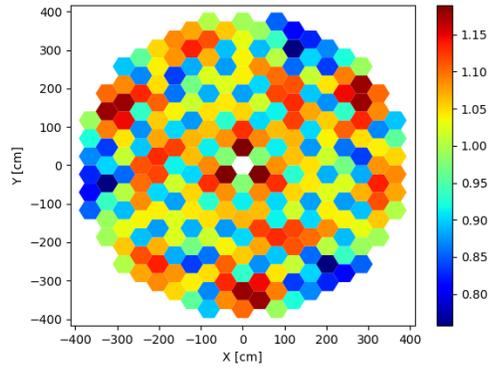
Case 20. Assembly Group 48 inserted.



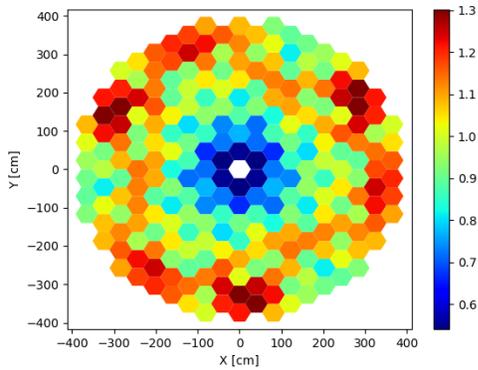
Case 21. Assembly Group 53 inserted.



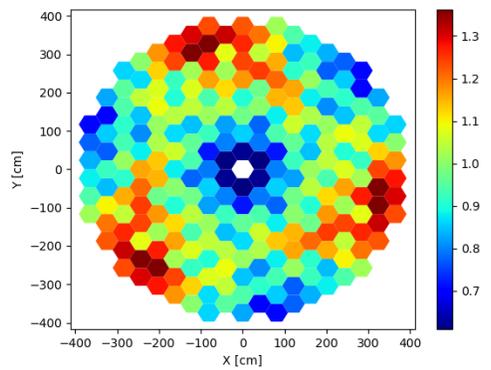
Case 22. Assembly Group 15 inserted.



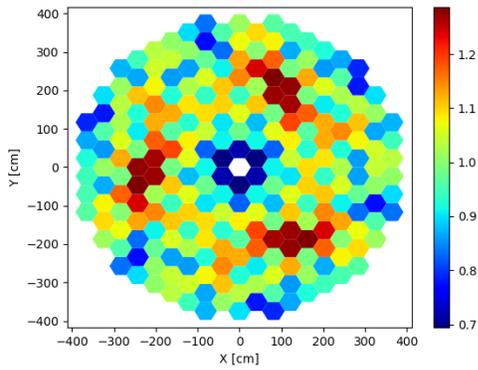
Case 23. Assembly Group 69 inserted.



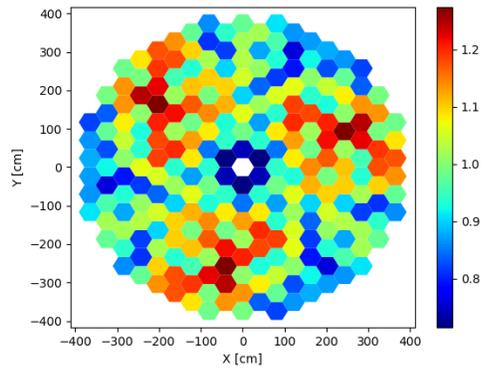
Case 24. Assembly Group 2 inserted.



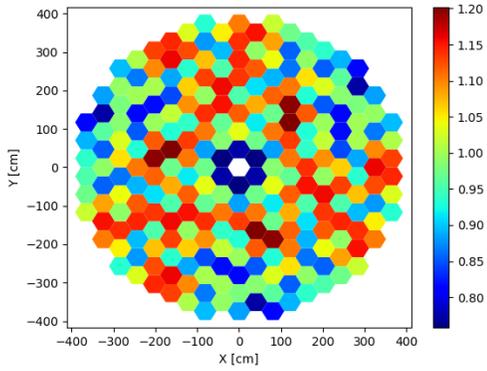
Case 25. Assembly Group 60 inserted.



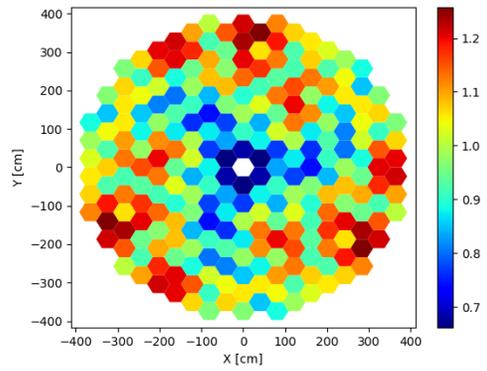
Case 26. Assembly Group 66 inserted.



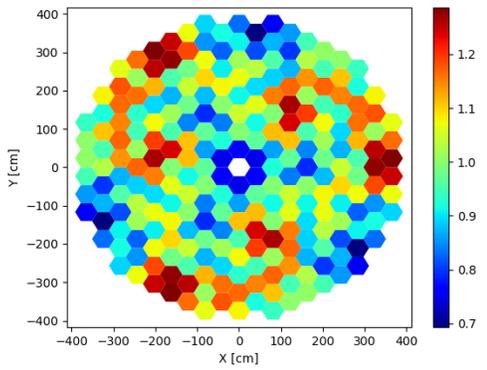
Case 27. Assembly Group 41 inserted.



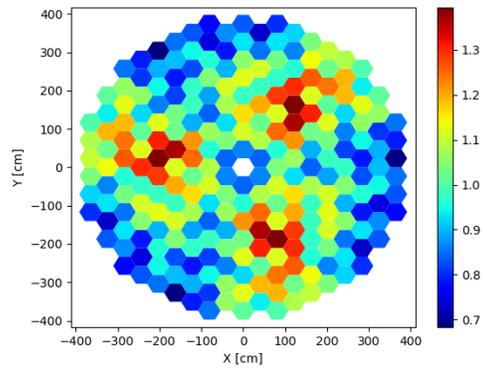
Case 28. Assembly Group 35 inserted.



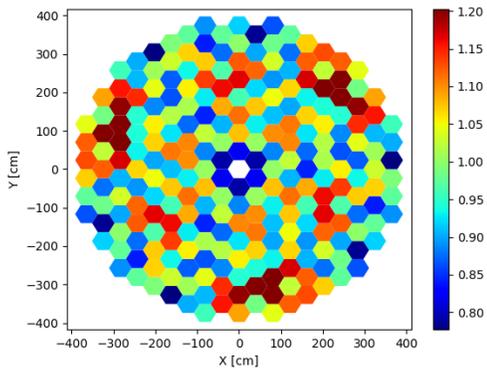
Case 29. Assembly Group 19 inserted.



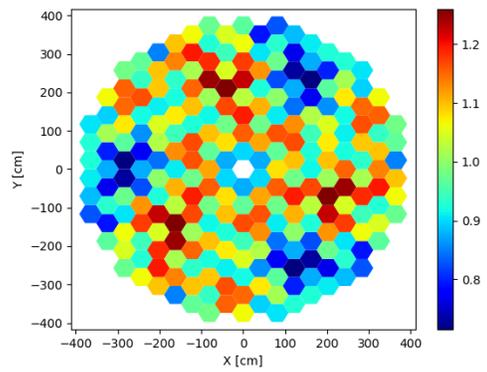
Case 30. Assembly Group 61 inserted.



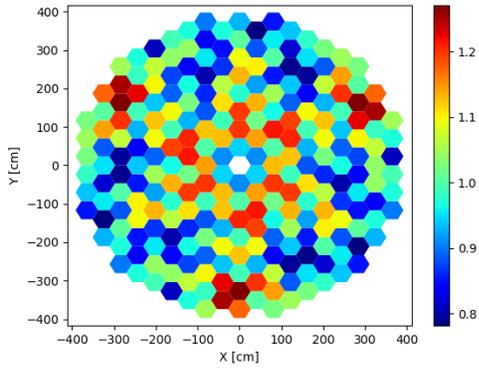
Case 31. Assembly Group 81 inserted.



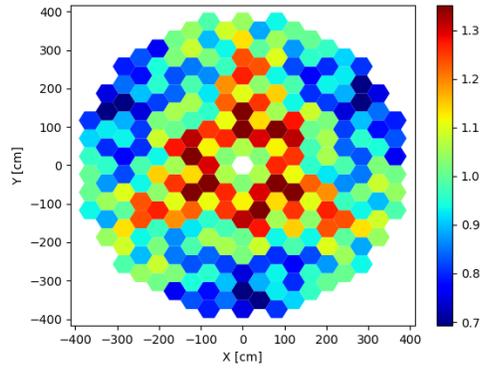
Case 32. Assembly Group 28 inserted.



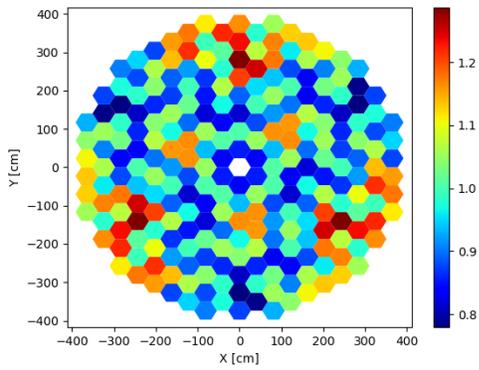
Case 33. Assembly Group 54 inserted.



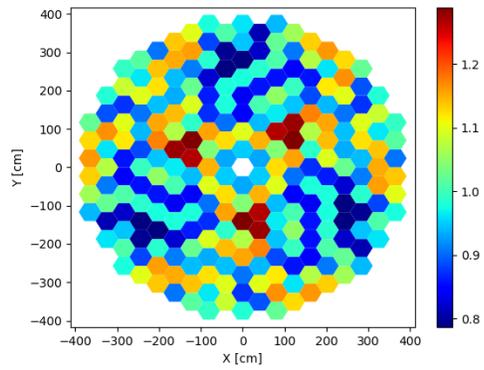
Case 34. Assembly Group 42 inserted.



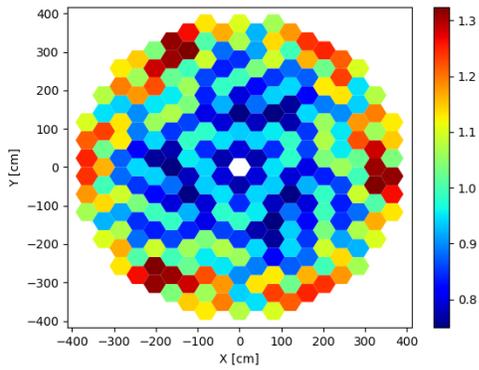
Case 35. Assembly Group 43 inserted.



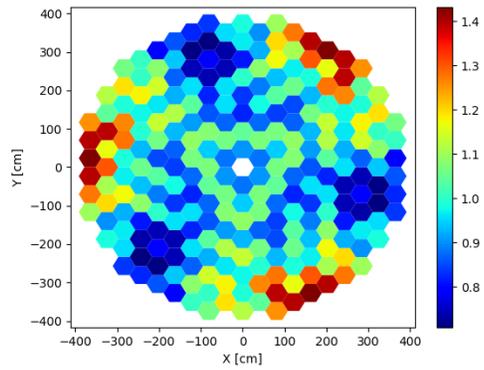
Case 36. Assembly Group 8 inserted.



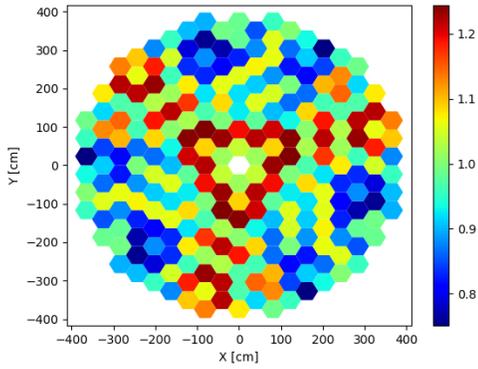
Case 37. Assembly Group 32 inserted.



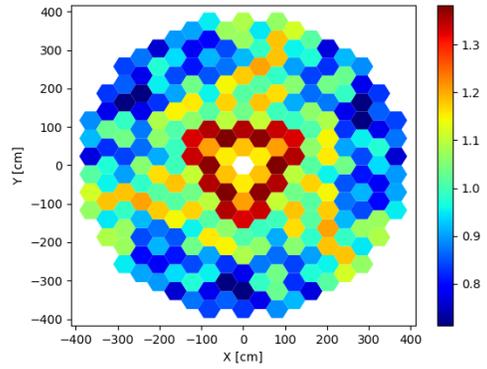
Case 38. Assembly Group 16 inserted.



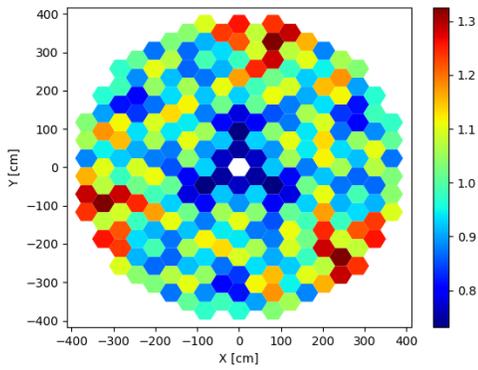
Case 39. Assembly Group 70 inserted.



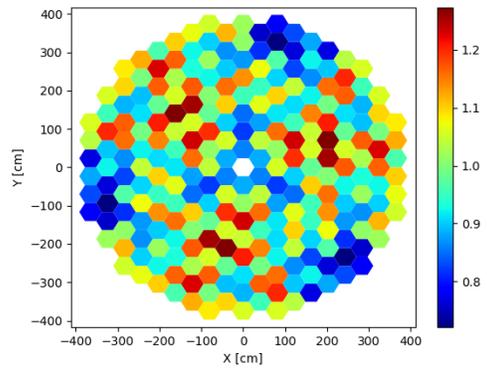
Case 40. Assembly Group 82 inserted.



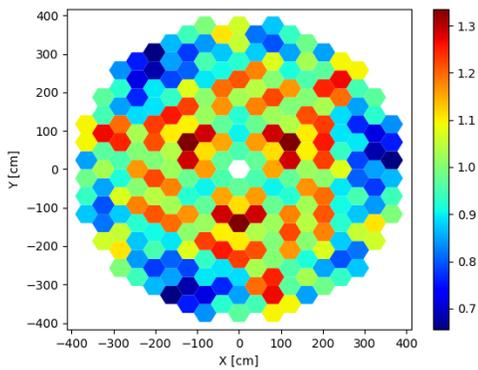
Case 41. Assembly Group 45 inserted.



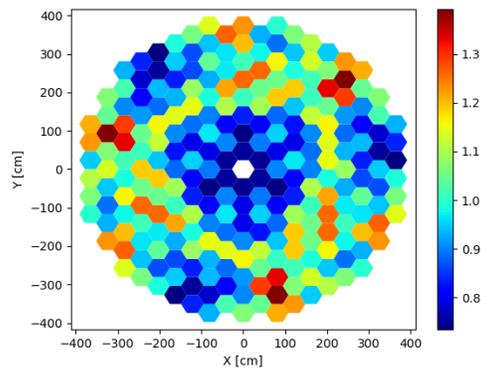
Case 42. Assembly Group 4 inserted.



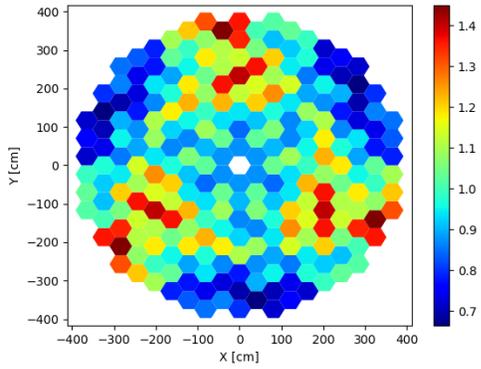
Case 43. Assembly Group 65 inserted.



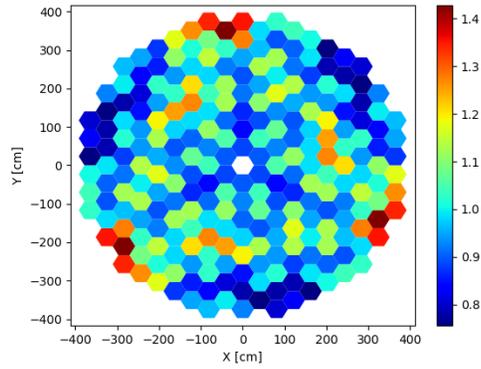
Case 44. Assembly Group 67 inserted.



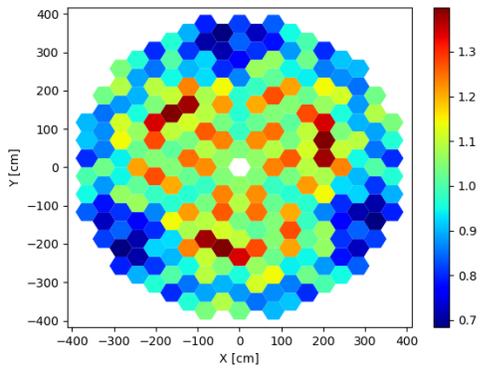
Case 45. Assembly Group 7 inserted.



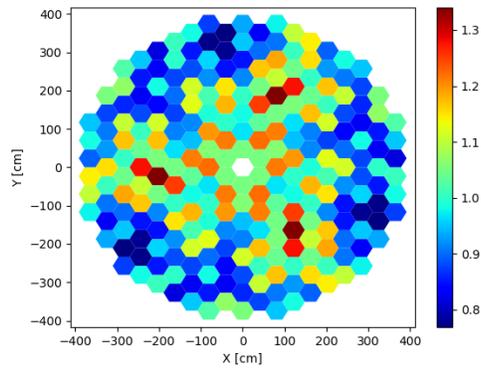
Case 46. Assembly Group 64 inserted.



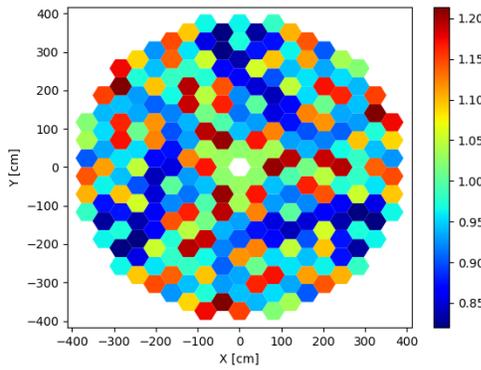
Case 47. Assembly Group 22 inserted.



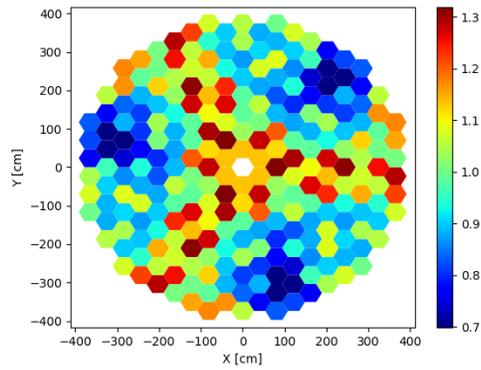
Case 48. Assembly Group 62 inserted.



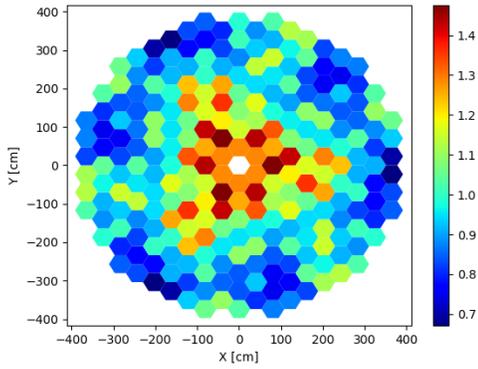
Case 49. Assembly Group 23 inserted.



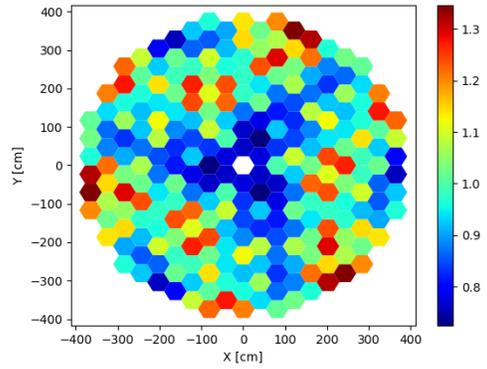
Case 50. Assembly Group 29 inserted.



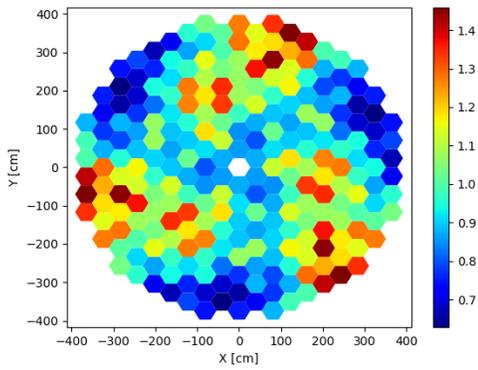
Case 51. Assembly Group 50 inserted.



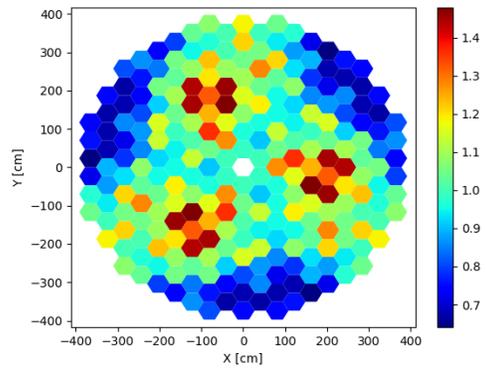
Case 52. Assembly Group 84 inserted.



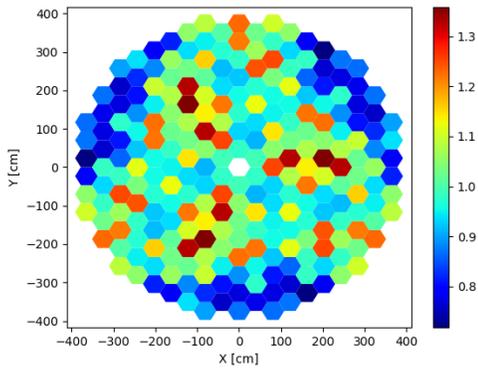
Case 53. Assembly Group 5 inserted.



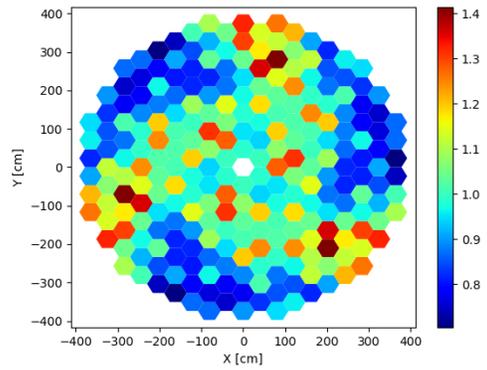
Case 54. Assembly Group 59 inserted.



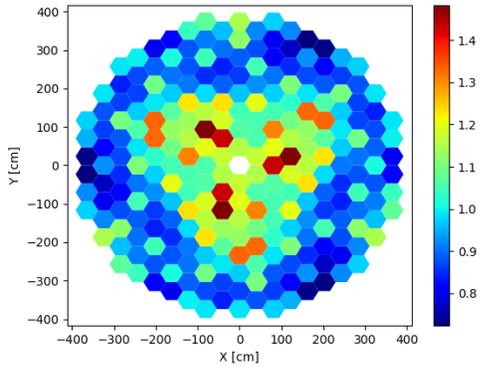
Case 55. Assembly Group 83 inserted.



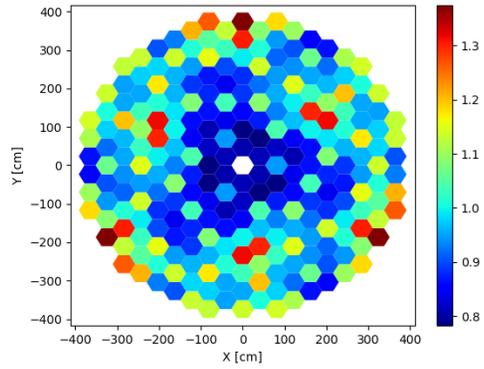
Case 56. Assembly Group 26 inserted.



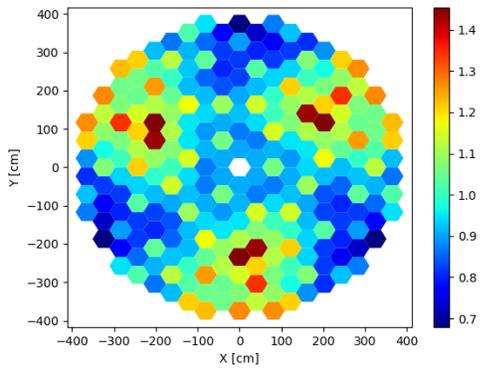
Case 57. Assembly Group 33 inserted.



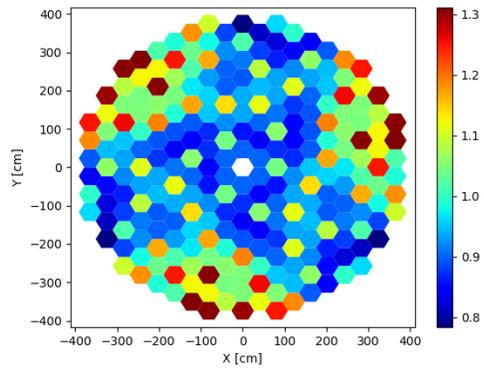
Case 58. Assembly Group 51 inserted.



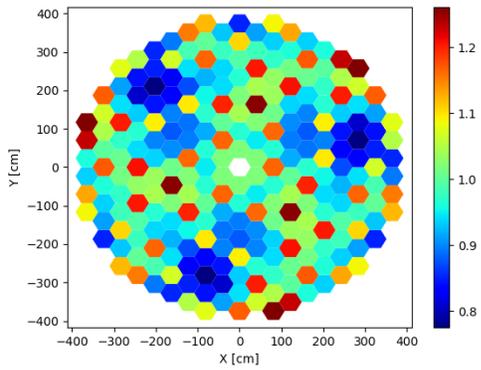
Case 59. Assembly Group 9 inserted.



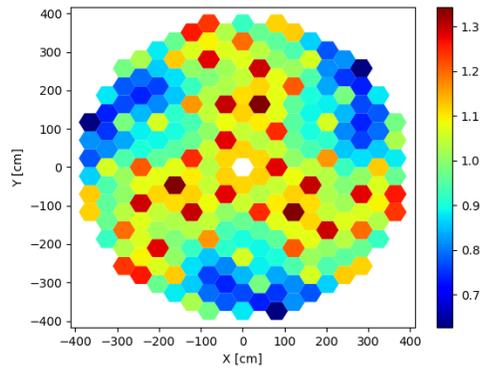
Case 60. Assembly Group 58 inserted.



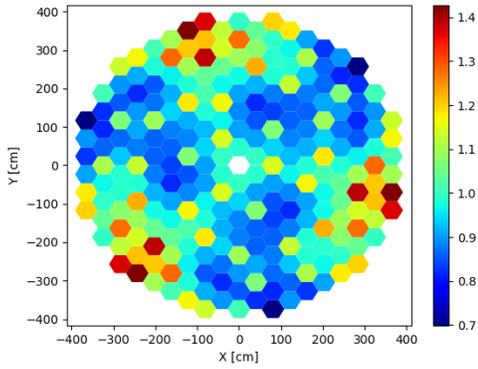
Case 61. Assembly Group 24 inserted.



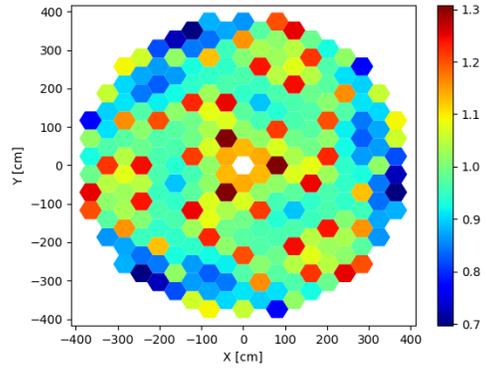
Case 62. Assembly Group 49 inserted.



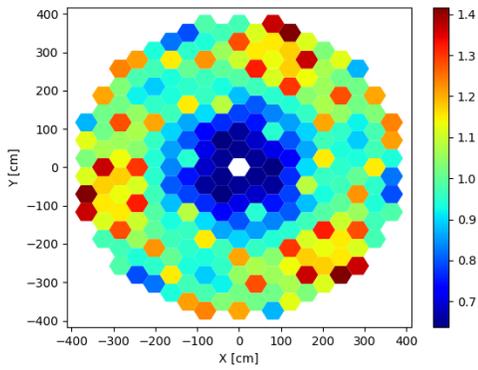
Case 63. Assembly Group 74 inserted.



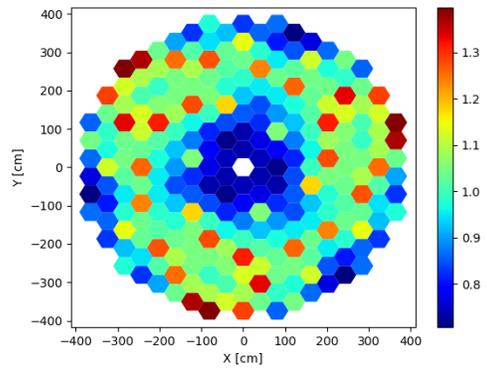
Case 64. Assembly Group 17 inserted.



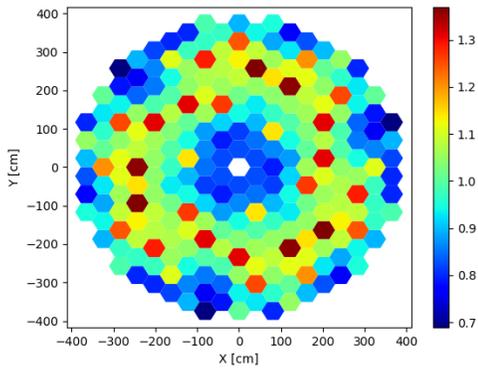
Case 65. Assembly Group 80 inserted.



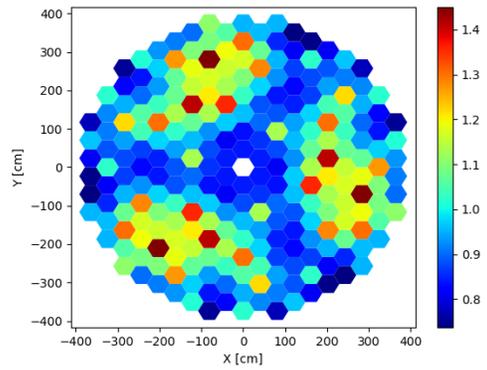
Case 66. Assembly Group 6 inserted.



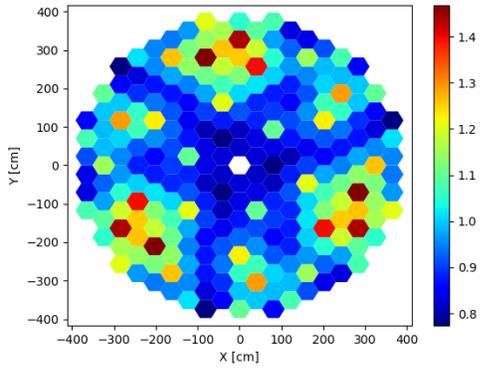
Case 67. Assembly Group 79 inserted.



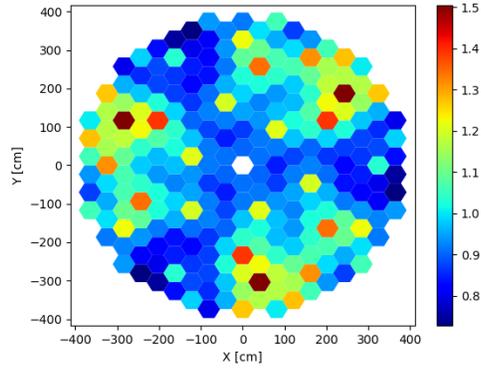
Case 68. Assembly Group 73 inserted.



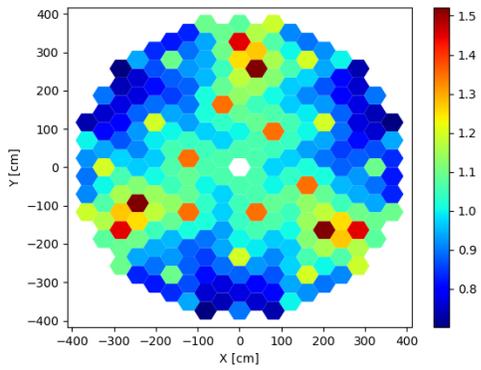
Case 69. Assembly Group 34 inserted.



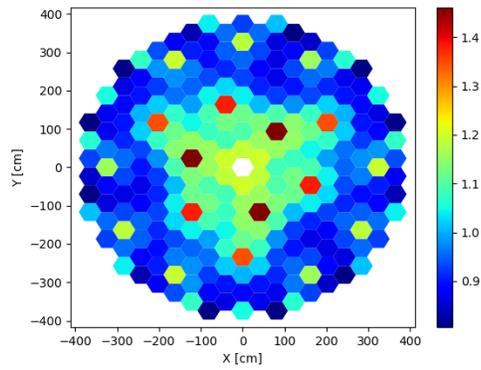
Case 70. Assembly Group 27 inserted.



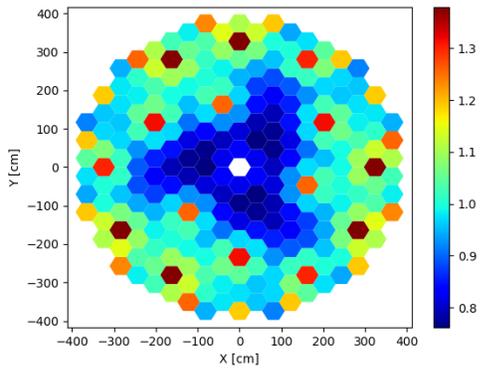
Case 71. Assembly Group 52 inserted.



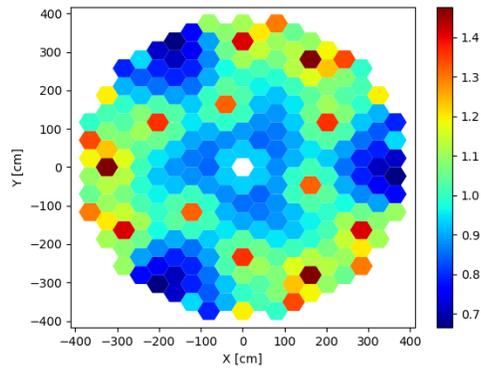
Case 72. Assembly Group 46 inserted.



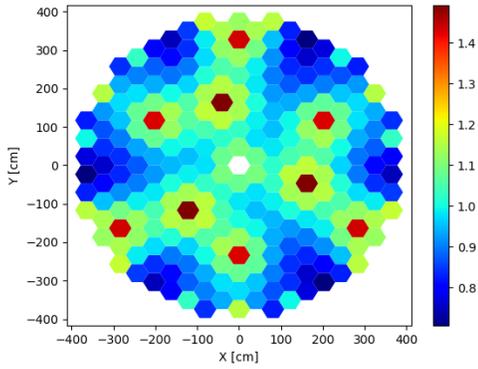
Case 73. Assembly Group 37 inserted.



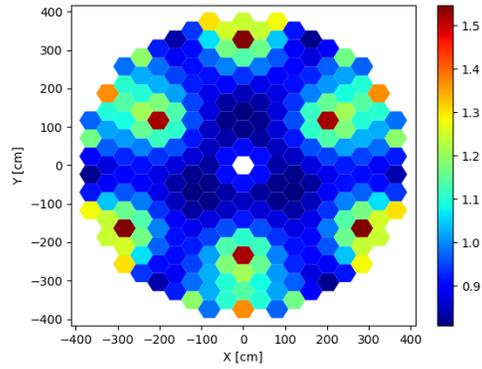
Case 74. Assembly Group 10 inserted.



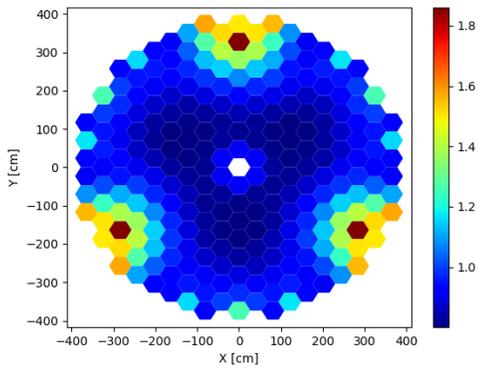
Case 75. Assembly Group 71 inserted.



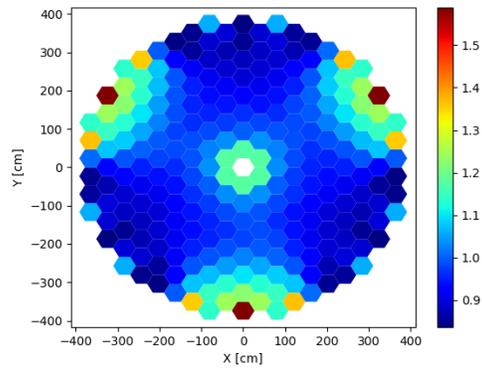
Case 76. Assembly Group 72 inserted.



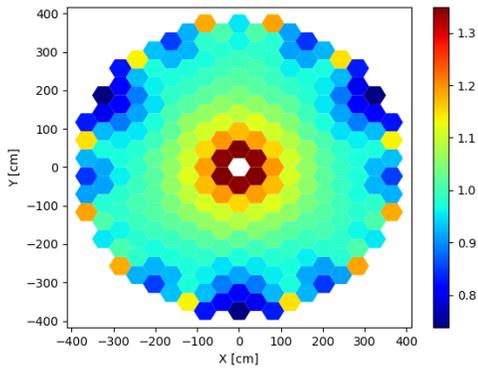
Case 77. Assembly Group 18 inserted.



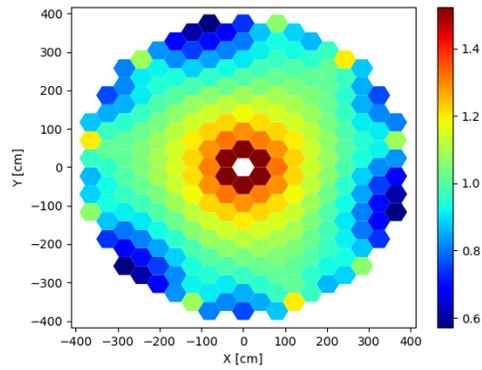
Case 78. Assembly Group 21 inserted.



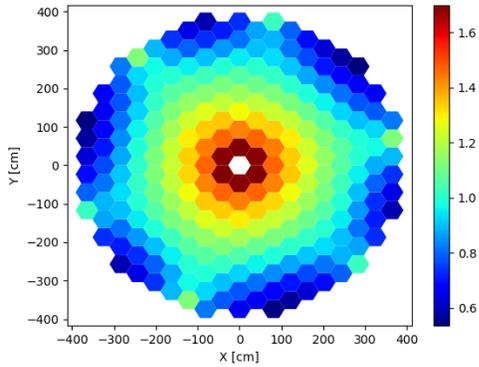
Case 79. Assembly Group 44 inserted.



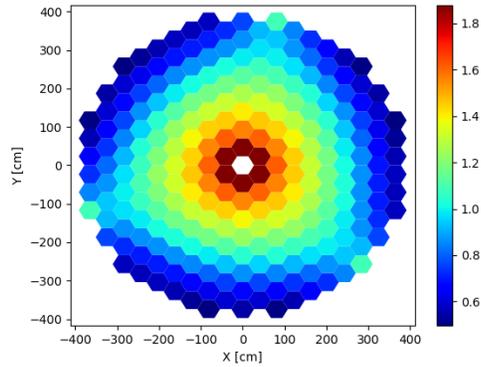
Case 80. Assembly Group 57 inserted.



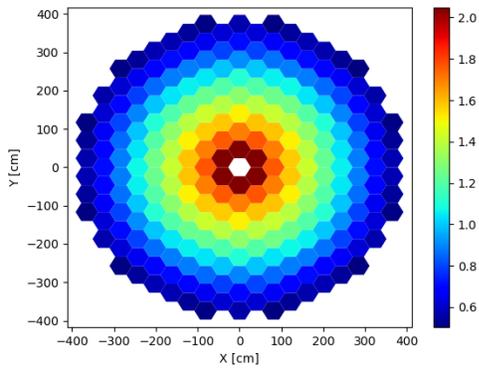
Case 81. Assembly Group 76 inserted.



Case 82. Assembly Group 78 inserted.



Case 83. Assembly Group 77 inserted.



Case 84. Assembly Group 75 inserted.

D.3 Quarter Density Burnable Poison Loading Control Blade Withdrawal

The eigenvalue and PPF results for each one-third symmetric CB group withdrawal can be seen in Supplementary Table D.5. The results are similar to those presented in section D.2, except CBs are withdrawn from full insertion instead of inserted from full withdrawal. The search procedure for each differs and can explain why in general the results of this section perform worse than those of the previous section.

Supplementary Table D.5. Complete results to accompany the quarter density europa CB withdrawal schedule.

Groups Withdrawn	Assembly Group Withdrawn	k_{eff}	σ [pcm]	Δk [pcm]	Maximum Assembly Group	Maximum Assembly PPF	Maximum Section Group	Maximum Section PPF
0	-	0.91612	16	-	1	2.095	1.2	2.357
1	76	0.91756	17	144	1	1.848	2.2	2.062
2	74	0.91859	17	103	2	1.716	1.2	1.934
3	73	0.92084	16	225	1	1.519	2.2	1.704
4	75	0.92222	17	138	1	1.461	2.2	1.636
5	84	0.92629	18	407	84	1.710	84.3	1.818

6	83	0.92903	17	274	84	1.502	84.3	1.603
7	20	0.93368	17	465	20	1.634	20.2	1.643
8	57	0.93673	17	305	57	1.773	57.1	1.853
9	44	0.94163	16	490	44	1.597	44.1	1.615
10	19	0.94460	17	297	44	1.441	44.2	1.456
11	68	0.94898	16	438	68	1.658	68.3	1.671
12	53	0.95169	17	271	57	1.477	57.1	1.538
13	5	0.95637	17	468	5	1.891	5.2	1.928
14	51	0.96028	16	391	5	1.668	5.2	1.685
15	77	0.96378	17	350	73	1.554	77.1	1.658
16	21	0.96682	16	304	73	1.544	77.1	1.634
17	22	0.96966	17	284	73	1.362	77.1	1.429
18	40	0.97310	17	344	40	1.368	40.2	1.382
19	66	0.97713	17	403	66	1.462	66.3	1.495
20	7	0.98066	16	353	7	1.430	5.1	1.461
21	78	0.98497	16	431	78	1.680	78.1	1.713
22	42	0.98765	17	268	78	1.512	78.2	1.526
23	45	0.99119	17	354	78	1.506	78.1	1.540
24	58	0.99455	16	336	58	1.534	58.3	1.560
25	8	0.99694	17	239	58	1.383	58.3	1.436
26	79	1.00134	18	440	79	1.682	79.1	1.729
27	6	1.00439	17	305	6	1.542	6.2	1.574
28	81	1.00757	17	318	81	1.446	81.1	1.557
29	46	1.01053	17	296	77	1.378	81.1	1.498
30	41	1.01367	16	314	79	1.309	79.1	1.348
31	39	1.01701	18	334	81	1.333	81.1	1.431
32	26	1.02051	17	350	26	1.337	26.3	1.373
33	24	1.02418	16	367	24	1.342	24.1	1.376
34	80	1.02905	17	487	80	1.648	80.3	1.726
35	72	1.03231	16	326	80	1.437	80.3	1.502
36	1	1.03844	17	613	1	2.436	1.2	2.679
37	59	1.04130	17	286	1	2.103	1.2	2.324
38	61	1.04436	16	306	1	1.706	1.2	1.883
39	27	1.04766	17	330	1	1.762	1.2	1.933
40	34	1.05087	17	321	1	1.655	1.2	1.827
41	60	1.05491	17	404	60	1.487	60.1	1.571
42	52	1.05785	17	294	60	1.365	60.1	1.443
43	17	1.06070	17	285	1	1.431	1.2	1.567
44	67	1.06485	17	415	77	1.459	81.1	1.545
45	65	1.06978	17	493	61	1.575	61.3	1.627
46	9	1.07229	17	251	1	1.430	1.2	1.536
47	31	1.07590	17	361	59	1.328	59.1	1.386
48	32	1.07962	16	372	61	1.420	61.3	1.452
49	10	1.08281	17	319	1	1.539	1.2	1.652
50	82	1.08655	16	374	82	1.371	1.2	1.461
51	70	1.08973	17	318	61	1.323	66.2	1.357
52	18	1.09351	16	378	1	1.295	1.2	1.417
53	50	1.09689	16	338	68	1.329	68.3	1.360
54	23	1.10027	17	338	68	1.225	1.2	1.297
55	62	1.10445	16	418	62	1.452	62.3	1.516
56	4	1.10846	17	401	1	1.592	1.2	1.709
57	63	1.11173	17	327	1	1.357	1.2	1.458
58	55	1.11501	16	328	1	1.233	1.2	1.342

59	25	1.11821	17	320	1	1.207	1.2	1.290
60	33	1.12183	17	362	33	1.222	39.2	1.264
61	64	1.12615	16	432	64	1.421	64.2	1.443
62	16	1.12913	16	298	64	1.365	64.3	1.384
63	48	1.13319	16	406	48	1.302	48.3	1.349
64	49	1.13779	16	460	49	1.355	49.1	1.387
65	69	1.14217	17	438	69	1.383	69.2	1.395
66	3	1.14576	16	359	1	1.363	1.2	1.421
67	71	1.14948	17	372	53	1.335	53.1	1.369
68	28	1.15334	17	386	3	1.265	3.3	1.279
69	43	1.15763	17	429	45	1.339	45.1	1.372
70	47	1.16063	17	300	49	1.216	49.1	1.245
71	14	1.16400	16	337	51	1.189	51.3	1.213
72	15	1.16749	17	349	3	1.248	3.1	1.267
73	56	1.17247	16	498	53	1.358	53.2	1.376
74	54	1.17592	16	345	53	1.228	54.2	1.243
75	2	1.18201	16	609	1	1.976	1.2	2.150
76	38	1.18563	16	362	1	1.769	2.2	1.951
77	36	1.18903	16	340	1	1.642	1.2	1.798
78	37	1.19300	17	397	1	1.438	1.2	1.574
79	35	1.19693	16	393	1	1.328	1.2	1.421
80	12	1.20167	16	474	2	1.715	2.2	1.882
81	29	1.20617	16	450	2	1.650	2.2	1.813
82	30	1.21108	16	491	1	1.602	1.2	1.755
83	13	1.21563	16	455	1	1.619	1.2	1.769
84	11	1.22208	5	645	2	1.878	2.2	2.044

APPENDIX E

EXAMPLE INPUT FILE

An example user input file to be used with the C++ script can be found below. The script reads in the first 18 characters of each line relevant to user input, which excludes the blank and header text lines meant to help organize the file. Note that in the text file editor lines can span as far as necessary but are shown with word wrap and tabbing here for readability. Each new line is indicated by the line number column shown on the left of the text. The line numbers shown are illustrative only; each new line is meant to start with the input value in column position one. Screenshots of individual code sections can be found in Chapter 7 if additional visual reference is desired.

```
1   === User Input File for AHTR Script ===
2
3   -SCRIPT MODE AND BASIC PHYSICS
4   8           // Script Mode. 0 - Statepoint, 1 - CB
      Insertion, 2 - CB Withdrawal, 3 - Criticality Search, 4 - TH
      Search, 5 - Depletion (without criticality search, no TH), 6 -
      Depletion (with criticality search, no TH), 7 - Depletion
      (without criticality search, TH), 8 - Depletion (with criticality
5   0           // Use cold dimensions. 0 - False (Thermal
      Expansion), 1 - True (No Thermal Expansion)
6   10000       // Statepoint Particles per Cycle
7   20          // Statepoint Active Cycles
8   20          // Statepoint Inactive Cycles
9
10  -GEOMETRY
11  --Fuel Particle
12  0.02135     // Cold Fuel Kernel Radius [cm]
13  0.03135     // Cold Buffer Layer Radius [cm]
14  0.03485     // Cold Inner Pyrolytic Carbon Layer Radius
      [cm]
15  0.03835     // Cold Silicon Carbide Layer Radius [cm]
16  0.04235     // Cold Outer Pyrolytic Carbon Layer Radius
      [cm]
17  --Fuel Lattice
18  1           // Particle Lattice type. 0 -
      cuboidal, 1 - cubic (wrt z, overrides x and y pitches)
19  4           // Width of Fuel Stripe (in Layers)
20  202        // Length of Fuel Stripe
21  0.09406     // Cold X Lattice Pitch [cm]
22  0.09128     // Cold Y Lattice Pitch [cm]
```

```

23 0.09266          // Cold Z Lattice Pitch [cm]
24 --Burnable Poison Lattice
25 1                // Burnable Poison Sphere Usage. 0 - None, 1 -
    Use
26 0.035           // Cold Poison Kernel Radius [cm]
27 0.09936        // Cold Poison Kernel Axial Pitch [cm]
28 5               // Number of BP columns (must be odd) [cm]
29 40              // Integral pitch spacing along the fuel plank
30 --Planks and Assembly
31 2.55           // Cold Plank Width [cm]
32 0.1            // Cold Sleeve Width (distance from coolant to
    fuel stripe) [cm]
33 14             // Distance Between Spacers [cm]
34 22.5           // Cold Assembly Apothem [cm]
35 46.8           // Cold Assembly Pitch [cm]
36 --Reflector Assemblies
37 1              // Reflector Assembly Central Cooling Hole
    Usage. 0 - False, 1 - True
38 2              // Reflector Assembly Central Cooling Hole
    Cold Radius [cm]
39 22.5           // Reflector Assembly Cold Apothem [cm]
40 --Axial Discretization
41 1              // Number of Modeled Axial Partitions in Active
    Core (up to 16) (fewer requires using cubic triso lattice)
42 550.02976      // Cold Active Core Height [cm]
43 25             // Cold Height of Top/Bottom Axial Reflectors
    [cm]
44 35             // Cold Height of Top/Bottom Axial Support
    Plates [cm]
45 1              // Axial Flowering / Interassembly Expansion.
    0 - False (uses axial average expansion). 1 - True (uses axial
    partition specific expansion)
46 --Beyond Active Core
47 478            // Cold Permanent Radial Reflector Outer
    Radius [cm]
48 479            // Cold Boron Carbide Layer Outer Radius [cm]
49 481            // Cold Core Barrel Outer Radius [cm]
50 519            // Cold Downcomer Outer Radius [cm]
51 520            // Cold Alloy N Outer Radius [cm]
52 525            // Cold Pressure Vessel Outer Radius [cm]
53
54 -MATERIALS
55 --Uniform Definitions (For statepoint or desiring uniformity, no
    need for unique material definitions. Uses axial average
    temperature and properties.)
56 0              // Fuel. 0 - Unique for each 1/3 assembly
    section , 1 - Only one, over whole geometry
57 0              // Other TRISO Particle Layers. 0 - Unique for
    each 1/3 assembly section, 1 - Only one material per layer
58 0              // Structural Graphite Components. 0 - Unique
    for each 1/3 assembly section, 1 - Just one graphite and C-C
    composite
59 0              // Burnable Poison Material. 0 - Unique, 1 -
    Uniform
60 0              // Single Control Blade Material. 0 - Unique,
    1 - Uniform
61 0              // Flibe. 0 - Unique, 1 - Uniform

```

```

62  --Densities
63  10.9          // Cold Fuel Density [g/cc]
64  1.0           // Cold Carbon Buffer Density [g/cc]
65  1.9           // Cold Inner Pyrolytic Carbon Density [g/cc]
66  3.1           // Cold Silicon Carbide Density [g/cc]
67  1.87          // Cold Outer Pyrolytic Carbon Density [g/cc]
68  1.75          // Cold Graphite Density [g/cc]
69  1.95          // Cold Carbon-Carbon Composite Density [g/cc]
70  1.25          // Cold Europa (Burnable Poison) Density
    [g/cc]
71  10.28         // Cold MHC (Control Blade) Density [g/cc]
72  2.37          // Cold Boron Carbide Density [g/cc]
73  8.93          // Cold Alloy N Density [g/cc]
74  7.92          // Cold Hastelloy 800 Density [g/cc]
75  --Thermal Expansion Coefficients (x 10^-6)
76  7.6           // Fuel
77  5.5           // Buffer
78  5.5           // Inner Pyrolytic Carbon
79  5             // Silicon Carbide
80  5.5           // Outer Pyrolytic Carbon
81  5             // Graphite
82  5             // Carbon-Carbon Composite
83  7.5           // Europa (BP)
84  4.8           // MHC (CB)
85  5             // Boron Carbide
86  13.6          // Alloy N (RPV Liner)
87  17.3          // Hastelloy 800 (RPV)
88  --Thermal Conductivities
89  3.7           // Fuel [W/(m*K)]
90  0.5           // Buffer Graphite [W/(m*K)]
91  4             // Inner Pyrolytic Carbon [W/(m*K)]
92  16            // Silicon Carbide [W/(m*K)]
93  4             // Outer Pyrolytic Carbon [W/(m*K)]
94  15            // Unirradiated Graphite [W/(m*K)]
95  15            // Fuel Stripe Matrix [W/(m*K)]
96
97  -THERMAL HYDRAULIC PARAMETERS
98  --Core Power and Flow Properties
99  1             // Number of Thermal Hydraulic Iterations
100  1.953376E-01 // Core Average Power Density [kW/g]
101  26750         // Mass Flow Rate through Core [kg/s]
102  293           // Cold Component Reference Temperature [K]
103  923           // Core Inlet Temperature [K]
104  1             // Print TH Profile for Highest Power Zone. 0
    - No, 1 - Yes
105  1             // Fuel Stripe Temperature Profile. 0 -
    Homogenized, 1 - Particle Reconstructed
106  --Flibe Properties
107  0.0056        // Viscosity [Pa*s]
108  13.525        // Prandtl Number
109  2415          // Heat Capacity [J/(kg*K)]
110  1             // Thermal Conductivity [W/(m*K)]
111
112  -DEPLETION SIMULATION OPTIONS
113  0             // Use Fluence-Dependent Thermal Conductivity
    and Thermal Expansion for Graphite. 0 - False (Invariant with
    Burnup), 1 - True

```

```
114 0 // Use Equilibrium Xenon Treatment for Fuel. 0
    - False, 1 - True
115 1 // Number of BP Burnable Zones. (Serpent
    supports up to 10)
116 1 // Target eigenvalue for Control Blade
    Movement
117 150 // Eigenvalue tolerance for Control Blade
    movement. [pcm]
118 0 // Initial Guess for Number of CB Groups to
    Insert
119 11000 // Depletion Particles per Cycle
120 20 // Depletion Active Cycles
121 20 // Depletion Inactive Cycles
122
```

APPENDIX F

SCRIPTS FOR WORK CONDUCTED

Below are a collection of additional input/output files and scripts used for the work conducted.

F.1 Hex Map Input File

Below are the contexts of the *HexMap.txt* file used in this work. It is used as a support file so that ATOMICS knows how to map the 84 symmetric assembly groups to hexagonal array locations for outputs results for the *RadialResults.txt* file. The first column corresponds to the y -index of the hexagonal array. The second column corresponds to the x -index of the hexagonal array. The third column corresponds with the assembly group number used with ATOMICS. Note that there are 252 lines present, which reflects back to the one-third core symmetry assumed since each of the 84 assembly groups are listed three times.

Supplementary Table F.1. Hexagon mapping to recreate a 19 x 19 array from ATOMICS phase space.

1	12	75
1	13	79
1	14	83
1	15	82
1	16	78
1	17	74
2	10	58
2	11	61
2	12	65
2	13	69
2	14	72
2	15	68
2	16	64
2	17	60
2	18	57
3	8	76
3	9	62
3	10	44
3	11	47
3	12	51
3	13	55
3	14	54

3	15	50
3	16	46
3	17	43
3	18	59
3	19	73
4	7	80
4	8	66
4	9	48
4	10	32
4	11	37
4	12	41
4	13	34
4	14	40
4	15	36
4	16	31
4	17	45
4	18	63
4	19	77
5	6	84
5	7	70
5	8	52
5	9	38
5	10	22
5	11	25
5	12	29
5	13	28
5	14	24
5	15	21
5	16	35
5	17	49
5	18	67
5	19	81
6	5	81
6	6	71
6	7	56
6	8	42
6	9	26
6	10	14
6	11	17
6	12	20
6	13	16
6	14	13
6	15	23
6	16	39
6	17	53
6	18	71
6	19	84
7	4	77
7	5	67
7	6	53
7	7	33
7	8	30
7	9	18

7	10	8
7	11	11
7	12	10
7	13	7
7	14	15
7	15	27
7	16	33
7	17	56
7	18	70
7	19	80
8	3	73
8	4	63
8	5	49
8	6	39
8	7	27
8	8	19
8	9	12
8	10	4
8	11	5
8	12	3
8	13	9
8	14	19
8	15	30
8	16	42
8	17	52
8	18	66
8	19	76
9	3	59
9	4	45
9	5	35
9	6	23
9	7	15
9	8	9
9	9	6
9	10	2
9	11	1
9	12	6
9	13	12
9	14	18
9	15	26
9	16	38
9	17	48
9	18	62
10	2	57
10	3	43
10	4	31
10	5	21
10	6	13
10	7	7
10	8	3
10	9	1
10	11	2
10	12	4

10	13	8
10	14	14
10	15	22
10	16	32
10	17	44
10	18	58
11	2	60
11	3	46
11	4	36
11	5	24
11	6	16
11	7	10
11	8	5
11	9	2
11	10	1
11	11	5
11	12	11
11	13	17
11	14	25
11	15	37
11	16	47
11	17	61
12	1	74
12	2	64
12	3	50
12	4	40
12	5	28
12	6	20
12	7	11
12	8	4
12	9	6
12	10	3
12	11	10
12	12	20
12	13	29
12	14	41
12	15	51
12	16	65
12	17	75
13	1	78
13	2	68
13	3	54
13	4	34
13	5	29
13	6	17
13	7	8
13	8	12
13	9	9
13	10	7
13	11	16
13	12	28
13	13	34
13	14	55

13	15	69
13	16	79
14	1	82
14	2	72
14	3	55
14	4	41
14	5	25
14	6	14
14	7	18
14	8	19
14	9	15
14	10	13
14	11	24
14	12	40
14	13	54
14	14	72
14	15	83
15	1	83
15	2	69
15	3	51
15	4	37
15	5	22
15	6	26
15	7	30
15	8	27
15	9	23
15	10	21
15	11	36
15	12	50
15	13	68
15	14	82
16	1	79
16	2	65
16	3	47
16	4	32
16	5	38
16	6	42
16	7	33
16	8	39
16	9	35
16	10	31
16	11	46
16	12	64
16	13	78
17	1	75
17	2	61
17	3	44
17	4	48
17	5	52
17	6	56
17	7	53
17	8	49
17	9	45

17	10	43
17	11	60
17	12	74
18	2	58
18	3	62
18	4	66
18	5	70
18	6	71
18	7	67
18	8	63
18	9	59
18	10	57
19	3	76
19	4	80
19	5	84
19	6	81
19	7	77
19	8	73

F.2 ATOMICS C++ Source Code

The ATOMICS C++ source code can be found in the GitHub directory below. It is quite lengthy at nearly 4000 lines (would require dozens of pages to present here).

<https://github.com/KyleMRamey/ATOMICS>

REFERENCES

- [1] M. S. Sohal, M. A. Ebner, P. Sabharwall and P. Sharpe, "Engineering Database of Liquid Salt Thermophysical and Thermochemical Properties," Idaho National Laboratory, Report INL/EXT-10-18297, Idaho Falls, Idaho, 2013.
- [2] V. K. Varma, D. E. Holcomb, F. J. Peretz, E. C. Bradley, D. Ilas, A. L. Qualls and N. M. Zaharia, "AHTR Mechanical, Structural, and Neutronic Preconceptual Design," Oak Ridge National Laboratory, Report ORNL/TM-2012/320, Oak Ridge, Tennessee, 2012.
- [3] D. E. Holcomb, D. Ilas, V. K. Varma, A. T. Cisneros, R. P. Kelly and J. C. Gehin, "Core and Refueling Design Studies for the Advanced High Temperature Reactor," Oak Ridge National Laboratory, Report ORNL/TM-2011/365, Oak Ridge, Tennessee, 2011.
- [4] D. T. Ingersoll, C. W. Forsberg, L. J. Ott, D. F. Williams, J. P. Renier, D. Wilson, S. J. Ball, L. Reid, W. R. Corwin, G. D. Del Cul, P. F. Peterson, H. Zhao, P. S. Pickard, E. J. Parma and M. Vernon, "Status of Preconceptual Design of the Advanced High-Temperature Reactor (AHTR)," Oak Ridge National Laboratory, Report ORNL/TM-2004/104, Oak Ridge, Tennessee, 2004.
- [5] P. N. Haubenreich and J. R. Engel, "Experience with the Molten-Salt Reactor Experiment," *Nuclear Applications and Technology*, vol. 8, no. 2, pp. 118-136, 1970.
- [6] S. R. Greene, J. C. Gehin, D. E. Holcomb, J. J. Carbajo, D. Ilas, A. T. Cisneros, V. K. Varma, W. R. Corwin, D. F. Wilson, G. L. Yoder Jr., A. L. Qualls, F. J. Peretz, G. F. Flanagan, D. A. Clayton, E. C. Bradley, G. L. Bell and J. D. Hunn, "Pre-Conceptual Design of a Fluoride-Salt-Cooled Small Modular Advanced High-Temperature Reactor (SmaHTR)," Oak Ridge National Laboratory, Report ORNL/TM-2010/199, Oak Ridge, TN, 2010.
- [7] P. M. Bardet, E. D. Blandford, M. Fratoni, A. Niquelle, E. Greenspan and P. F. Peterson, "Design, Analysis and Development of the Modular PB-AHTR," in *International Conference on Advances in Nuclear Power Plants (ICAPP)*, Anaheim, CA, 2008.
- [8] C. Andreades, A. T. Cisneros, J. K. Choi, A. Y. K. Chong, M. Fratoni, S. Hong, L. R. Huddar, K. D. Huff, D. L. Krumwiede, M. R. Laufer, M. Munk, R. O. Scarlat, N. Zweibaum, E. Greenspan and P. F. Peterson, "Technical Description of the "Mark 1" Pebble-Bed Fluoride-Salt-Cooled High-Temperature Reactor (PB-FHR) Power Plant," University of California Berkeley, Report UCBTH-14-002, Berkeley, CA, 2014.
- [9] E. Blandford, "Commercializing the FHR: From IRP to Kairos Power," in *Enabling the Next Nuclear: FHRs to Megawatts Workshop*, West Falls Church, VA, 2018.

- [10] S. G. Nagley, C. M. Barnes, D. L. Husser, M. L. Nowlin and W. C. Richardson, "Fabrication of Uranium Oxycarbide Kernels for HTR Fuel," Idaho National Laboratory, Report INL/CON-10-18858, Idaho Falls, Idaho, 2010.
- [11] J. Leppänen et al., "The Serpent Monte Carlo code: Status, development and applications in 2013.," *Annals of Nuclear Energy*, vol. 82, pp. 142-150, 2015.
- [12] K. Terrani, "Accelerating the Deployment of Advanced Nuclear Energy Systems," *Nuclear News*, pp. 34-37, April 2020.
- [13] M. P. Trammell, B. C. Jolly, M. D. Richardson, A. T. Schumacher and K. A. Terrani, "Advanced Nuclear Fuel Fabrication: Particle Fuel Concept for TCR," Oak Ridge National Laboratory, Report ORNL/SPR-2019/1216, Oak Ridge, Tennessee, 2019.
- [14] F. Rahnema, B. Petrovic, C. Edgar, D. Zhang, P. Avigni, M. Huang and S. Terlizzi, "The Current Status of the Tools for Modeling and Simulation of Advanced High Temperature Reactor Neutronics Analysis," Georgia Institute of Technology, Report CRMP-2015-12-001, Atlanta, Georgia, 2015.
- [15] A. T. Cisneros and D. Ilas, "Neutronics and Depletion Methods for Multibatch Fluoride Salt-Cooled High-Temperature Reactors with Slab Fuel Geometry," *Nuclear Technology*, vol. 183, pp. 331-340, 2013.
- [16] R. Kelly and D. Ilas, "Verification of a Depletion Method in SCALE for the Advanced High Temperature Reactor," in *PHYSOR 2012 - Advances in Reactor Physics - Linking Research, Industry, and Education*, Knoxville, Tennessee, USA, April 15-20, 2012.
- [17] B. Petrovic and G. Maldonado, "Fuel and Core Design Options to Overcome the Heavy Metal Loading Limit and Improve Performance and Safety of Liquid Salt Cooled Reactors," Nuclear Energy University Programs, Report 12-3870, 2016.
- [18] C. Gentry, G. I. Maldonado, O. Chvala and B. Petrovic, "Neutronic Evaluation of a Liquid Salt-Cooled Reactor Assembly," *Nuclear Science and Engineering*, 2017.
- [19] C. Gentry, K. S. Kim and G. I. Maldonado, "Two-Step Procedure for Liquid-Salt-Cooled- Reactor Analysis," *Nuclear Technology*, vol. 204, no. 3, pp. 299-317, 2018.
- [20] L. M. Huang, "Neutronic Analysis and Optimization of the Advanced High Temperature Reactor Fuel Design Using Machine Learning," Georgia Institute of Technology, Atlanta, GA, 2017.
- [21] RELAP5-3D Code Development Team, "RELAP5-3D Code Manual Volume I: Code Structure, System Models and Solution Methods," Idaho National Laboratory, Idaho Falls, Idaho, 2012.

- [22] P. Avigni and B. Petrovic, "Fuel Element and Full Core Thermal-Hydraulic Analysis of the AHTR for the Evaluation of the LOFC Transient," *Annals of Nuclear Energy*, vol. 64, pp. 499-510, 2014.
- [23] P. Avigni and B. Petrovic, "On-line Refueling for the Advanced High Temperature Reactor (AHTR)," *Nuclear Engineering and Design*, vol. 340, pp. 166-182, 2018.
- [24] K. M. Ramey and B. Petrovic, "Monte Carlo modeling and simulations of AHTR fuel assembly to support V&V of FHR core physics methods," *Annals of Nuclear Energy Vol. 118*, pp. 272-282, 2018.
- [25] P. Chiotti, R. W. C and K. M, "Thermodynamic properties of uranium oxycarbides," *Journal of the Less-Common Metals vol. 10*, pp. 273-289, 1966.
- [26] N. N. Ngoepe and J. P. R. de Villiers, "The thermal expansion of 3C-SiC in TRISO particles by high temperature X-ray diffraction," *Journal of Nuclear Materials*, no. 438, pp. 88-93, 2013.
- [27] S. Stecura and W. J. Campbell, "Thermal expansion and phase inversion of rare-earth oxides," United States Department of the Interior, Bureau of Mines, Washington DC, 1961.
- [28] P. Hidnert and G. W. B, "Thermal expansion of molybdenum," *Scientific Papers of the Bureau of Standards vol. 19*, pp. 429-444, 1924.
- [29] P. J. Hacker, G. B. Neighbour and M. Brian, "The coefficient of thermal expansion of nuclear graphite with increasing thermal oxidation," *Journal of Physics D: Applied Physics vol. 33*, pp. 991-998, 2000.
- [30] G. V. Tsagareishvili, T. G. Nakashidze, J. S. Jobava, G. P. Lomidze, D. E. Khulelidze, D. S. Tsahareighvili and O. A. Tsagareishvili, "Thermal Expansion of Boron and Boron Carbide," *Journal of the Less-Common Metals*, no. 117, pp. 159-161, 1986.
- [31] J. Maxwell, *A Treatise on Electricity and Magnetism Volume I*, London: MacMillan and Co., 1873.
- [32] B. Petrovic, K. M. Ramey and I. Hill, "Benchmark Specifications for the Fluoride-salt High-temperature Reactor (FHR) Reactor Physics Calculations," Organization for Economic Cooperation and Development -Nuclear Energy Agency - Nuclear Science Committee (Report NEA/NSC/R(2020)5), Paris, France, 2020.
- [33] R. Stainsby, M. Worsley, A. Grief, F. Dawson, M. Davies, P. Coddington, J. Baker and A. Dennier, "Development of Local Heat Transfer Models for Safety Assessment of High Temperature Gas-Cooled Reactor Cores - Part I: Pebble Bed Reactors," *Journal of Engineering for Gas Turbines and Power*, vol. 132, 2010.
- [34] J. J. Duderstadt and L. J. Hamilton, *Nuclear Reactor Analysis*, Ann Arbor, Michigan: John Wiley & Sons, Inc., 1976.

- [35] D. M. McEligot, W. D. Swank, D. L. Cottle and F. I. Valentin, "Thermal Properties of G-348 Graphite," Idaho National Laboratory, Report INL/EXT-16-38241 1330693, Idaho Falls, ID, 2016.
- [36] A. A. Campbell, Y. Katoh, M. A. Snead and K. Takizawa, "Property Changes of G347A Graphite Due to Neutron Irradiation," *Carbon*, vol. 109, pp. 860-873, 2016.
- [37] J. Leppänen and M. Aufiero, "Development of an Unstructured Mesh Based Geometry Model in the Serpent 2 Monte Carlo Code," in *PHYSOR 2014*, Kyoto, Japan, 2014.
- [38] C. J. Werner and e. al., "MCNP6.2 Release Notes," Los Alamos National Laboratory, report LA-UR-18-20808, Los Alamos, New Mexico, 2018.
- [39] D. van Veen, "Efficient improvement of local power estimation in the general purpose Monte Carlo code MCNP5," Delft University of Technology, Delft, Netherlands, 2009.
- [40] D. Van Veen and J. E. Hoogenboom, "Efficiency Improvement of Local Power Estimation in the General Purpose Monte Carlo Code MCNP," *Progress in Nuclear Science and Technology*, vol. 2, pp. 866-871, 2011.
- [41] N. Khandera and e. al., "Properties of Equilateral Triangles," Brilliant, 2020. [Online]. Available: <https://brilliant.org/wiki/properties-of-equilateral-triangles/>. [Accessed 18 August 2020].
- [42] B. Petrovic, K. Ramey, I. Hill, E. Losa, M. Elsayi, Z. Wu, C. Lu, J. Gonzales, D. Novog, G. Chee, K. Huff, M. Margulis, N. Read and E. Shwageraus, "Preliminary Results of the NEA FHR Benchmark Phase I-A and I-B (Fuel Element 2D Benchmark)," in *The International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering*, Raleigh, NC, 2021.
- [43] K. Shibata, O. Iwamoto, T. Nakagawa, N. Iwamoto, A. Ichihara, S. Kunieda, S. Chiba, K. Furutaka, N. Otuka, T. Ohsawa, T. Murata, H. Matsunobu, A. Zukeran, S. Kamada and J. Katakura, "JENDL-4.0: A New Library for Nuclear Science and Engineering," *Journal of Nuclear Science and Technology*, vol. 48, no. 1, pp. 1-30, 2011.
- [44] A. Johnson, D. Kotlyar, S. Terlizzi and G. Ridley, "serpentTools: A Python Package for Expediting Analysis with Serpent," *Nuclear Science and Engineering*, 2020.
- [45] A. E. Isotalo, J. Leppanen and J. Dufek, "Preventing xenon oscillations in Monte Carlo burnup calculations by enforcing equilibrium xenon distribution," *Annals of Nuclear Energy*, vol. 60, pp. 78-85, 2013.
- [46] J. Leppanen, "Serpent - a Continuous-energy Monte Carlo Reactor Physics Burnup Calculation Code," VTT Technical Research Centre, Finland, 2013.

- [47] B. Petrovic, T. Flaspöehler and K. Ramey, "Benchmarking FHR Core Physics Simulations: 2D Fuel Assembly Model," in *International Conference of the Croatian Nuclear Society: Nuclear Option for CO₂ Free Energy Generation*, Zadar, Croatia June 3-6, 2018.
- [48] T. Maruyama and M. Harayama, "Neutron irradiation effect on the thermal conductivity and dimensional change of graphite materials," *Journal of Nuclear Materials*, no. 195, pp. 44-50, 1992.
- [49] J. Dufek and J. E. Hoogenboom, "Numerical stability of existing Monte Carlo burnup codes in cycle calculations of critical reactors," *Nuclear Science and Engineering*, vol. 162, no. 3, pp. 307-311, 2009.
- [50] D. Kotlyar, Y. Shaposhnik, E. Fridman and E. Shwageraus, "Coupled neutronic thermo-hydraulic analysis of full PWR core with Monte-Carlo based BGCore system," *Nuclear Engineering and Design*, no. 241, pp. 3777-3786, 2011.
- [51] J. Dufek, D. Kotlyar, E. Shwageraus and J. Leppänen, "Numerical stability of the predictor-corrector method in Monte Carlo burnup calculations of critical reactors," *Annals of Nuclear Energy*, no. 56, pp. 34-38, 2013.
- [52] D. Kotlyar and S. E., "On the use of predictor-corrector method for coupled Monte Carlo burnup codes," *Annals of Nuclear Energy*, no. 58, pp. 228-237, 2013.
- [53] J. R. Lamarsh and A. J. Baratta, *Introduction to Nuclear Engineering*, Third Edition, Upper Saddle River, New Jersey: Prentice-Hall, 2001.
- [54] J. Dufek, D. Kotlyar and E. Shwageraus, "The stochastic implicit Euler method - A stable coupling scheme for Monte Carlo burnup calculations," *Annals of Nuclear Energy*, no. 60, pp. 295-300, 2013.
- [55] D. Kotlyar and E. Shwageraus, "Numerically stable Monte Carlo-burnup-thermal hydraulic coupling schemes," *Annals of Nuclear Energy*, no. 63, pp. 371-381, 2014.
- [56] D. Kotlyar and E. Shwageraus, "Monitoring and preventing numerical oscillations in 3D simulations with coupled Monte Carlo codes," *Annals of Nuclear Energy*, vol. 71, pp. 198-205, 2014.
- [57] V. Valtavirta and J. Leppänen, "New stochastic substep based burnup scheme for Serpent 2," in *PHYSOR 2018*, Cancun, Mexico, 2018.
- [58] D. Kotlyar and E. Shwageraus, "Stochastic semi-implicit substep method for coupled depletion Monte-Carlo codes," *Annals of Nuclear Energy*, vol. 92, pp. 52-60, 2016.
- [59] Serpent Development Team, "Serpent Wiki," VTT Technical Research Centre of Finland, 2021. [Online]. Available: https://serpent.vtt.fi/mediawiki/index.php?title=Input_syntax_manual&oldid=4980.

- [60] F. B. Brown, "'K-effective of the World' and Other Concerns for Monte Carlo Eigenvalue Calculations," *Progress in Nuclear Science and Technology*, vol. 2, pp. 738-742, 2011.
- [61] T. Ikonen, "FINIX - Fuel behavior model and interface for multiphysics applications - Code documentation for version 0.13.9," VTT, Espoo, Finland, 2013.
- [62] S. 2. User, "Question regarding hex-y detector lattice," 2019. [Online]. Available: URL requires access to the SERPENT 2 Discussion Forum (<https://ttuki.vtt.fi/serpent/viewtopic.php?f=25&t=3093&p=10489&hilit=Hexagonal+mesh+detector#p10489>). [Accessed 18 August 2020].