

Neutronic-Thermal Simulation of Micro Reactor Designs for the Purpose of Analyzing the Impact of Thermal Expansion and Hydrogen Migration in Metal Hydride Moderators

by

W. Reed Kendrick

B.S. Mechanical Engineering, University of Texas at Austin, 2018

Submitted to the Department of Nuclear Science & Engineering
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY IN NUCLEAR ENGINEERING

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2024

© 2024 W. Reed Kendrick. This work is licensed under a [CC BY-NC-ND 4.0](#) license.

The author hereby grants to MIT a nonexclusive, worldwide, irrevocable, royalty-free license to exercise any and all rights under copyright, including to reproduce, preserve, distribute and publicly display copies of the thesis, or release the thesis under an open-access license.

Authored by: W. Reed Kendrick

Department of Nuclear Science & Engineering

August 13, 2024

Certified by: Benoit Forget

Korea Electric Power Professor of Nuclear Engineering, Thesis Supervisor

Accepted by: Ju Li

Tokyo Electric Power Company Professor in Nuclear Engineering

Neutronic-Thermal Simulation of Micro Reactor Designs for the Purpose of Analyzing the Impact of Thermal Expansion and Hydrogen Migration in Metal Hydride Moderators

by

W. Reed Kendrick

Submitted to the Department of Nuclear Science & Engineering
on August 13, 2024 in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY IN NUCLEAR ENGINEERING

ABSTRACT

The recent increased interest in microreactor designs has presented the opportunity to take advantage of the smaller core dimensions to perform steady state neutronic-thermal coupled simulation with the inclusion of an additional physics system. This work accomplishes this, by adding thermal expansion and zirconium hydride-based hydrogen diffusion to the neutronic-thermal simulation of multiple heat pipe microreactor designs. Microreactors' smaller cores are inherently characterized by more leakage than gigawatt-scale reactor cores. The inclusion of thermal expansion's representation in the coupling system may reveal neutronic or thermal impacts of geometric expansion that have yet to be noted for these smaller scale geometries. This is the impetus for the work on thermal expansion. The work on hydrogen diffusion is inspired by the common use of zirconium hydride in microreactor designs as a moderator. This material provides high density of hydrogen with high melting point, but features a well documented increase in mobility of hydrogen within the zirconium lattice at high temperatures. Coupling this migration of hydrogen within the neutronic-thermal simulation is performed in order to identify and analyze neutronic and thermal impacts due to the movement of hydrogen within the moderator. Additionally, a heat pipe failure case is simulated for each microreactor geometry studied, aimed to analyze the impacts of multi-pipe failure on both thermal expansion and hydrogen diffusion, as well as their downstream neutronic-thermal effects.

Thesis supervisor: Benoit Forget

Title: Korea Electric Power Professor of Nuclear Engineering

Acknowledgments

I could not have completed the work in this thesis, nor the process of finishing my PhD without so many people who have helped along the way. The following list of people does not do this fact justice, but it is my best effort.

Benoit Forget, my thesis advisor; Thank you for being an inspiration for how one can be brilliant and yet incredibly understanding and personable.

Guillaume Giudicelli, my friend and thesis reader; Thank you for hours spent helping me debug, teaching me theory, and for being someone always interested in whatever I have to talk about.

Amelia Trainer, my friend and office-mate; Thank you for being an attentive friend, for being a shoulder to lean on during this process, and for showing me what it means to be truly a good friend to someone.

Ben Dacus, my friend and classmate; Thank you for being a source of laughter and happiness, and for being someone I can be myself around.

My family, my brother, and my grandparents; Thank you for being a constant source of love and support, both financially and emotionally. I will forever appreciate the privilege it was to have you all in my life during this process.

My dog, Moose; Thank you for being a cuddly friend to come home to, and for surviving your many episodes of eating things you absolutely should not eat.

My parter, Kelly O’Laughlin; Thank you. Thank you for supporting me during this process. Thank you for listening to me complain. Thank you for taking more than your fair-share of chores when I was working. Thank you for loving me. This is as much my accomplishment as it is yours. I love you.

Contents

Title page	1
Abstract	3
Acknowledgments	5
List of Figures	9
List of Tables	15
1 Introduction	17
1.1 Background	17
1.2 Objectives and Scope of Thesis	18
2 General Background	21
2.1 Heat Pipes	21
2.2 Neutronic-Thermal Coupling	22
2.3 Core Geometries	23
2.3.1 MRaD Core	23
2.3.2 Simba Core	25
2.3.3 SR Core	26
3 General Methods	27
3.1 MOOSE	27
3.1.1 Non-conservation	27
3.2 OpenMC	28
3.2.1 Heating Tally	28
3.3 MITHP	29
3.3.1 MOOSE Implementation	31
3.3.2 Comparisons	31
4 Thermal Expansion	33
4.1 Background	33
4.1.1 Thermal Expansion	33
4.1.2 Geometric Deformation in Nuclear Reactors	35
4.1.3 DAGMC	36

4.1.4	Literature Review	38
4.2	Methods	39
4.2.1	Coupling Scheme	39
4.2.2	Thermal Expansion in MOOSE	40
4.2.3	DAG-OpenMC and Heat Pipes	45
4.3	Results	47
4.3.1	MRaD Core	51
4.3.2	SR Core	58
4.4	Conclusion	64
5	Hydrogen Diffusion	65
5.1	Background	65
5.1.1	Why Metal-Hydride Moderators	65
5.1.2	Zirconium Hydride	67
5.1.3	Diffusion of H in ZrH	68
5.1.4	Functional Expansions	71
5.1.5	Literature Review	77
5.2	Methods	81
5.2.1	Coupling Scheme	81
5.2.2	Diffusion in MOOSE	83
5.2.3	Continuously Varying Material Tracking	85
5.2.4	Function Expansion Use in Meshes	88
5.2.5	Functional Expansion Use with Annular Cylinders	92
5.3	Test Cases	96
5.3.1	1D Case	96
5.3.2	2D Case	97
5.4	Results	101
5.4.1	MRaD Core	103
5.4.2	Simba Core	116
5.4.3	SR Core	128
5.5	Conclusion	140
6	Conclusions	143
	References	147

List of Figures

1.1	Depiction of BWXT Terrestrial Micro RX Microreactor on semi-truck, retrieved from [5].	18
1.2	General coupling scheme used for coupled simulation in this thesis.	19
2.1	General schematic of a heat pipe, featuring liquid flow down wick and vapor flow centrally. Figure retrieved from [18]	21
2.2	MRaD Core Geometry	23
2.3	Simba Core Geometry	25
2.4	SR Core Geometry	26
3.1	Cumulative count of "frozen" heat pipes versus nonlinear iteration.	30
3.2	Graphical representation of MITHP heat pipe process.	31
3.3	Comparison of HTPIPE and MITHP temperature profile using the same inputs.	32
3.4	Comparison of SOCKEYE and MITHP temperature profile using similar inputs. Dashed lines indicate transitions from condenser to adiabatic region to evaporator.	32
4.1	Lennard-Jones Potential for a paired atom, showing that an increase in temperature results in an increased atomic spacing.	34
4.2	MRaD core geometry comparison of CAD mesh versus facet geometry, clipped in the axial mid-plane.	37
4.3	Detailed thermal expansion coupled simulation scheme. All processes not in bold are performed via Python script.	40
4.4	Graphical representation of the structure of the MOOSE thermal expansion solve, with some aspects omitted. Letters correspond to more detailed explanations that can be found in Section 4.2.2.	41
4.5	Example of mesh collapse where thermal expansion caused the air gap between control drum and reflector to collapse in the MRaD geometry. The center, darker mesh surface indicates where parts of the mesh element have inverted.	43
4.6	SR mesh focused on one of the inner "dummy" assemblies. Highlights how tightly packed the inner TRI mesh elements are, which may causes overlaps during displacement and simulation failure.	44
4.7	MRaD geometry with thermal expansion incorrectly represented in air mesh blocks.	45

4.8	Depiction of the CSG heat pipe inserted into the CAD-tesselated heat pipe void. Note the gaps and overlaps caused by the polygonized geometry.	46
4.9	Comparison of the original MRaD geometry (left) with the thermal expansion form (right). Note the lack of air and transformation of control drums into contiguous reflector.	47
4.10	Comparison of the original SR geometry (below) with the thermal expansion form (above). Note the removal of control drums and shutdown rod space, as well as removal of gas from fuel annulus.	48
4.11	k_{eff} and relative change of k_{eff} versus iteration count for MRaD and SR cores during thermal expansion coupled simulation.	49
4.12	L2 norm of cell volume relative residuals versus iteration count for MRaD and SR cores during thermal expansion coupled simulation.	50
4.13	MRaD temperatures after final coupled thermal expansion iteration.	51
4.14	MRaD displacement (cm) after final thermal expansion iteration. Displacement is relative to the axial mid-plane.	52
4.15	MRaD displacement in centimeters in alternative case with bottom plane constrained. Results of this constraint showed no neutronic or thermal difference from a case with no axial constraint.	53
4.16	Cell percent volume change due to thermal expansion.	54
4.17	Percent change in pin powers of MRaD core after thermal expansion, including standard deviation of changes.	55
4.18	Cell-averaged temperature change after thermal expansion of MRaD geometry.	55
4.19	Selected heat pipe failure locations and the resulting mid-plane temperature effect on the MRaD core.	56
4.20	Comparison of mesh displacement between the heat pipe failure case and the no-failure case for the MRaD geometry. Heat pipe temperature impact is constrained by surrounding functioning heat pipes and the overall displacement effect is minor. Displacement is relative to the axial mid-plane.	57
4.21	SR temperatures after final coupled thermal expansion iteration.	58
4.22	SR displacement (cm) after final coupled thermal expansion iteration. Displacement is relative to the axial mid-plane.	59
4.23	SR Cell percent volume change due to thermal expansion.	60
4.24	Percent change in pin powers of SR core after thermal expansion, including standard deviation of changes.	61
4.25	Cell-averaged temperature changes due to thermal expansion of SR geometry.	61
4.26	Selected heat pipe failure locations and the resulting mid-plane temperature effect on the SR core.	62
4.27	Comparison of mesh displacement between the heat pipe failure case and the non-failure case for the SR geometry. Heat pipe temperature impact is constrained by the surrounding functioning heat pipes and the overall displacement is negligible.	63
5.1	H-Zr phase diagram with isobars of equilibrium H ₂ pressures indicated by colored lines. Figure retrieved from Olander et. al. [16]	68

5.2	Hydrogen atom density as a function of temperature for various metal hydrides in equilibrium with 1 atm of hydrogen gas. Figure retrieved from [102]. . . .	69
5.3	Legendre polynomials plotted up to 5 th order.	73
5.4	Zernike polynomials plotted on a unit disk up to 5 th order.	74
5.5	Axial stoichiometric ratio of hydrogen in YH _x sampled in varying radial regions. Innermost radial regions experiences highest temperature gradients. Figure retrieved from [131].	78
5.6	Griffin SN(3,3) calculated axial power distributions with and without YH-hydrogen diffusion via SWIFT. Figure retrieved from [33].	79
5.7	Axial hydrogen stoichiometric ratio and radially-averaged moderator temperature, sourced from [37].	80
5.8	Simplified hydrogen diffusion coupled simulation scheme. CAD here referred to CAD-based mesh geometries, as MOOSE interacts with meshes rather than strictly CAD geometries.	81
5.9	Detailed hydrogen diffusion coupled simulation scheme. All processes not in bold are performed via Python script.	82
5.10	Graphical representation of the structure of the MOOSE solve, with some aspects omitted. Letters correspond to more detailed explanations that can be found in Section 5.2.2.	84
5.11	Arbitrary functional expansion heating distribution used in conjunction with Figure 5.12. No axial shape is included.	89
5.12	Figures showing error when normalization radius does not account for polygonal mesh geometry. The cylindrical radius is 1cm.	90
5.13	Figures showing normalization radius convergence and error minimization for functional expansion input. The resulted in an optimal normalization radius of 1.024cm when used with the originally 1cm radius fuel pin.	91
5.14	Figures showing normalization radius convergence and error minimization for functional expansion output. The resulted in an optimal normalization radius of 0.9755cm when used with the originally 0.975cm radius moderator rod. . .	92
5.15	Heating mesh tally results for annular and cylindrical geometries. Units are eV per source particle.	93
5.16	L2 norm of the percent error from using Zernike FETs versus a fine cylindrical heating mesh tally.	94
5.17	L2 norm of the percent error from using Zernike FETs with normalization versus a fine cylindrical mesh tally.	95
5.18	1D test case graphical summary.	96
5.19	1D test case results comparing transmission and reflection probabilities. . . .	97
5.20	Simplified graphic of the workflow for the 2D test case.	98
5.21	Hydrogen migration result due to different ΔT values and their respective impact on axial thermal neutron current. Asymmetry is due to stochastic uncertainty.	99
5.22	Contour plot of total thermal neutron current versus temperature gradient and moderator width.	100

5.23	(a) shows "True Width" versus thermal gradient. This value is the equivalent width of moderator with no temperature gradient that would result in the same total thermal neutron current leaving the moderator. (b) shows "Necessary Width" versus thermal gradient. This value is the equivalent width of moderator needed that would conserve the original $\Delta T = 0$ thermal neutron current leaving the moderator.	100
5.24	L2 norm of axial hydrogen residuals for hottest moderator rod, for each core design during hydrogen diffusion coupled simulation.	102
5.25	MRaD temperatures after final coupled iteration.	103
5.26	Final hydrogen concentration values in MRaD moderator rods, returned by MOOSE.	104
5.27	Temperature and H/Zr stoichiometric ratio of hottest moderator and fuel, sampled via axial slice at the core mid-plane.	106
5.28	Temperature and H/Zr stoichiometric ratio of hottest moderator and fuel, sampled via axial slice at the core top.	106
5.29	Axially-averaged hydrogen profile for MRaD hottest moderator rod.	107
5.30	Percent change in pin powers of MRaD core due to hydrogen migration, and their respective standard deviations.	108
5.31	Axial segment-averaged volumetric heating rate of MRaD hottest fuel pin from initial and final iteration, and the percent change between the two.	109
5.32	Axial segment-averaged volumetric heating rate of MRaD hottest moderator rod from initial and final iteration, and the percent change between the two.	110
5.33	Change in MRaD fuel and moderator mid-plane temperature due to hydrogen migration.	111
5.34	Comparisons of initial and final iteration axial temperature of MRaD hottest fuel pin and moderator rod.	112
5.35	Selected heat pipe failure locations and the resulting mid-plane temperature effect on the MRaD core.	113
5.36	Axially-averaged hydrogen profile for MRaD hottest moderator rod after heat pipe failure.	114
5.37	Axial segment-averaged volumetric heating rate of MRaD hottest fuel pin from initial and final iteration with failed heat pipes. Includes the percent change between the two cases.	115
5.38	Change in MRaD hottest fuel pin caused by hydrogen migration with heat pipe failure.	115
5.39	Simba temperature distributions after final coupled iteration.	116
5.40	Final hydrogen concentration values in Simba moderator rods, returned by MOOSE.	117
5.41	Temperature and H/Zr stoichiometric ratio of hottest moderator and fuel, sampled via axial slice at the core mid-plane.	118
5.42	Temperature and H/Zr stoichiometric ratio of hottest moderator and fuel, sampled via axial slice at the core top.	119
5.43	Axially-averaged hydrogen profile for Simba hottest moderator rod.	119
5.44	Percent change in pin powers of Simba core due to hydrogen migration, and their respective standard deviations.	121

5.45	Axial segment-averaged volumetric heating rate of Simba hottest fuel pin from initial and final iteration, and the percent change between the two.	121
5.46	Axial segment-averaged volumetric heating rate of MRaD hottest moderator rod from initial and final iteration, and the percent change between the two.	122
5.47	Change in fuel and moderator mid-plane temperature due to hydrogen migration in the Simba core.	123
5.48	Comparison of initial and final iteration axial nodal-average temperatures of Simba core hottest fuel pin and moderator rod.	124
5.49	Selected heat pipe failure locations and the resulting mid-plane temperature effect on the Simba core.	125
5.50	Axially-averaged hydrogen profile for Simba hottest moderator rod after heat pipe failures	126
5.51	Axial segment-averaged volumetric heating rate of Simba hottest fuel pin from initial and final iteration with failed heat pipes. Includes the percent change between the two cases.	126
5.52	Change in Simba hottest fuel pin temperature caused by hydrogen migration with heat pipe failure.	127
5.53	SR core temperatures after final coupled iteration.	128
5.54	Final hydrogen concentration values in SR moderator rods, returned by MOOSE.	130
5.55	Temperature and H/Zr stoichiometric ratio of hottest moderator and fuel, sampled via axial slice at the core mid-plane. Moderator rods are the smaller pins on the left side.	131
5.56	Temperature and H/Zr stoichiometric ratio of hottest moderator and fuel, sampled via axial slice at the core top. Moderator rods are the smaller pins on the left side.	132
5.57	Axially-averaged hydrogen profile for SR hottest moderator rod.	132
5.58	Percent change in pin powers of SR core due to hydrogen migration and their respective standard deviations.	133
5.59	Axial segment-averaged volumetric heating rate of SR hottest fuel pin from initial and final iteration, including the percent change between the two.	134
5.60	Change in fuel and moderator mid-plane temperature due to hydrogen migration in the SR core.	135
5.61	Comparison of initial and final iteration axial nodal-average temperatures of SR core hottest fuel pin and moderator rod.	136
5.62	Selected heat pip failure locations and the resulting mid-plane temperature effect on the SR core.	137
5.63	Axially-averaged hydrogen profile for SR hottest moderator rod after heat pipe failures.	138
5.64	Axial segment-averaged volumetric heating rate of SR hottest moderator rod from initial and final iteration with heat pipes. Includes the percent change between the two cases.	138
5.65	Change in SR hottest fuel pin temperature caused by hydrogen migration with heat pipe failure.	139

List of Tables

2.1	Material summary for the MRaD core geometry.	24
2.2	Material summary for the Simba core geometry.	25
2.3	Material summary for the SR core geometry.	26
4.1	OpenMC particle count settings thermal expansion studies.	49
4.2	Global neutronic results and differences from non-expansion MOOSE run for MRaD core. "AR" is Absorption Rate, "FF" is fission factor. Absorption rates are in units of "per source particle".	53
4.3	Global neutronic results and differences from non-expansion MOOSE run for SR core. "AR" is Absorption Rate, "FF" is fission factor. Absorption rates are in units of "per source particle".	60
5.1	Some common moderating materials and their moderating properties. ξ is the average logarithmic energy decrement, MSDP is the macroscopic slowing down power, MR is the moderating ratio. Values from [84]–[86]	66
5.2	OpenMC particle count settings for hydrogen diffusion studies.	101
5.3	Final axial values for x in ZrH_x , sampled from all MRaD moderator rods. The layer under each result are the percent change from original $ZrH_{1.67}$	105
5.4	Global neutronic results and differences from initial flat-hydrogen case for MRaD core. "AR" is Absorption Rate, "Thermal FF" is the ratio of fissions caused by thermal neutrons compared to total fissions. Absorption rates are in units of "per source particle". \pm values are one standard deviation.	108
5.5	Final axial values for x in ZrH_x , sampled from all Simba moderator rods. The layer under each result is the percent change from original $ZrH_{1.67}$	118
5.6	Global neutronic results and differences from initial flat-hydrogen case for Simba core. "AR" is Absorption Rate, "Thermal FF" is the ratio of fission caused by thermal neutrons compared to total fission. Absorption rates are in units of "per source particle".	120
5.7	Final axial values for x in ZrH_x , sampled from all SR moderator rods. The layer under each result is the percent change from original $ZrH_{1.67}$	129
5.8	Global neutronic results and differences from initial flat-hydrogen case for SR core. "AR" is Absorption Rate, "Thermal FF" is the ratio of fissions caused by thermal neutrons compared to total fissions. Absorption rates are in units of "per source particle".	133

6.1 Comparison of OpenMC particle simulation rates vesus method and geometry.
Rates are particles per second. 144

Chapter 1

Introduction

1.1 Background

The nuclear community is currently evaluating the role of smaller reactors, especially those in the "micro reactor" category [1]. The "micro reactor" label is generally applied to nuclear reactors with less than 50 megawatts of electrical production capability. These reactors are characterized by smaller geometries, often with steam generators and other components designed to be included as an integrated feature of the reactor. This gives the reactor a "plug-and-play" capability, being able to be transported to a site and immediately activated to generate power for the customer [2].

These designs aim to leverage the smaller geometry of the reactor to gain economic advantages in multiple ways. Firstly, vendors are hoping to claim a reduced need for the emergency planning zone and other safety considerations due to the smaller footprint and lower fuel inventory of these cores [3]. This will require agreement by the US Nuclear Regulator Commission and a new licensing structure for prospective microreactors. This may be a difficult task given the short deployment timeline that microreactor vendors are hoping for [2], [4].

Secondly, the previously mentioned mobility means an opportunity to enter more economically viable markets that have limited options for power. While economic analysis of microreactors are still hypothetical at this point, the general consensus of estimates is that the levelized cost of electricity (LCOE) of these reactors is likely to be higher than that of traditional gigawatt reactors [6]. While this may make microreactors a difficult sell in traditional settings, for remote regions like rural Alaska that rely primarily on diesel for energy, these microreactors may provide a cheaper alternative [7]–[9]. Many microreactor designs envision a capability to transport the reactor via truck, for example Figure 1.1, allowing them to penetrate regions yet untapped by nuclear power.

Finally, microreactors are extremely attractive for space applications as both sources of power and propulsion. The high density of power and long lifetimes before refueling make microreactors appealing for space applications where mass is a highly restricted quantity. The KRUSTY reactor experiments on the use of a microreactor with Stirling engine were highly successful in showing the potential for microreactor use for space missions [10].

This is not an exhaustive list of the arguments in favor of microreactor development,



Figure 1.1: Depiction of BWXT Terrestrial Micro RX Microreactor on semi-truck, retrieved from [5].

but merely two examples of the reasons why current interest in microreactor development is so high. The work of this thesis is inspired by a combination of this interest and the conduciveness of the smaller core geometry for multiphysics simulation using high-fidelity methods.

1.2 Objectives and Scope of Thesis

The microreactor concepts that exist are highly varied in design, featuring many different proposed fuel materials, moderating materials, and assembly layouts. This presents a unique opportunity for reactor physics analysis. The small scale of these microreactors allow for relative ease in multiphysics analysis, needing fractions of the computational resources required for the same analysis of a gigawatt-scale reactor. This work takes advantage of this feature, performing neutronic-thermal coupled simulation of multiple heat pipe microreactor cores with an additional physics system included in simulation, specifically thermal expansion and hydrogen diffusion. The general layout of this coupling scheme can be seen in Figure 1.2.

The first physics system added is thermal expansion. Small reactor geometries have been shown to exhibit stronger feedback responses to geometric deformation [10], due to the higher leakage that comes with smaller dimensions. This increased importance of geometric deformation indicates that there may be utility in simulating thermal expansion along with neutronic-thermal simulation. The creation of Direct Accelerated Geometry Monte Carlo (DAGMC) [11] allows mesh-based geometries to be used in particle transport simulation with OpenMC [12], [13]. Combining this with the ability of the MOOSE [14] Solid Mechanics module to perform thermal expansion-based mesh deformation, a thermally deformed geometry can be represented in OpenMC particle transport. This thesis utilizes this process to simulate thermal expansion and quantify the effect that this thermal expansion has on steady-state neutronic and thermal behavior of the core.

The second physical phenomenon studied in this thesis is hydrogen diffusion in the zirco-

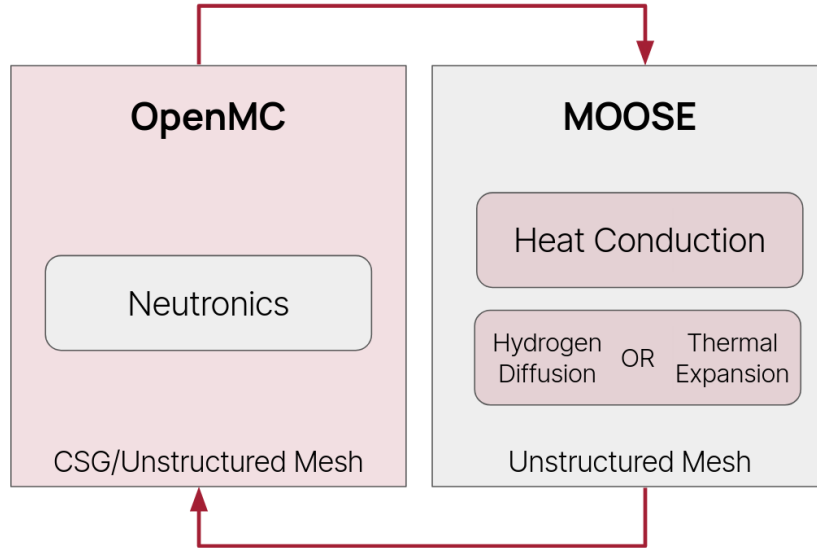


Figure 1.2: General coupling scheme used for coupled simulation in this thesis.

niium hydride moderator. Zirconium hydride is a common moderating material suggested for microreactors due to its high hydrogen density and high melting point [15], [16]. One of the caveats to zirconium hydride is that the hydrogen dispersed throughout its zirconium lattice is quite mobile at normal core operating temperatures [15]. Given the importance of the moderator on neutronic response, the movement of hydrogen within the zirconium lattice will undoubtedly cause ripple effects in neutronic behavior, and in the thermal response following. Quantifying the magnitudes and locations of these responses will help inform future reactor designs for reactors that use zirconium hydride as a moderator.

Unlike the case of thermal expansion, where MOOSE and OpenMC are coupled by utilizing DAGMC to perform simulations on essentially the same deformed geometry, the hydrogen diffusion process described in this work opts to transfer spatial information between the neutronic and thermal simulation via functional expansion representation. Functional expansions are also used to represent the volumetric heating rates for fuel and moderator in the core. This allows for continuous representation of power and hydrogen concentration, and allows analysis of hydrogen migration’s impact on power distribution in the fuel to be continuous as well.

Also of note in this thesis are the use of heat pipes as a method of heat removal in all cores studied. Heat pipes are included in microreactor concepts because of their reliability as a heat sink method that lacks moving parts. They are also known to fail due to operational limits or manufacturing defects [17]. Because heat pipe failure causes strongly increased temperatures in the core, which in turn increases thermal expansion and hydrogen diffusion, each microreactor studied in this thesis includes a heat pipe failure study.

Chapter 2

General Background

2.1 Heat Pipes

Heat pipes are typically metal tubes filled with highly heat-conducting material and feature a central wick to enable capillary action. A generalized heat pipe geometry can be seen in Figure 2.1, showing the movement of vapor and liquid between evaporator and condenser sections. Heat pipes feature some advantageous characteristics for reactor use, such as lack of moving components. A downside is that failure is not uncommon for heat pipes due to manufacturing defects and the proximity of several limits during operation[17]. Heat pipes can also fail as part of a "cascade", where other heat pipe failures increase the thermal load on neighboring heat pipes, exceeding an operational limit, and thus causing secondary failures. This risk is particularly high when the heat pipes regularly operate near these limits. This makes their presence an interesting note for the reactor design discussed in this work. Therefore, this work will include a case study by "failing" four selected heat pipes in the center of the core, aiming to simulate a minor cascade failure, and analyze the impact on hydrogen migration and its effects.

For the purposes of this work, heat pipe behavior is simulated in MOOSE as a custom object, named MITHP, detailed in Section 3.3. The code within was based on Los Alamos Laboratory's HTPIPE [19], and features a simple 1D iterative pressure and temperature solver. The codes response has been compared against HTPIPE and SOCKEYE [20] results




Figure 2.1: General schematic of a heat pipe, featuring liquid flow down wick and vapor flow centrally. Figure retrieved from [18]

in similar environments, and has been determined to be sufficiently accurate to be used as a hypothetical potassium heat pipe response object.

2.2 Neutronic-Thermal Coupling

Neutronic-thermal coupling is the general category for any simulation process that couples neutron transport simulation with heat transfer simulation. This coupling provides resolution of the most significant form of feedback in nuclear reactor core; neutron flux determines energy deposition in the core, which determines heating rates and thermal profile in the core, which cause cross-section broadening via Doppler feedback mostly in the fuel and density changes in the moderator, which in turn affects neutron flux.

Representing this coupled system has been a task that various groups in the nuclear industry have accomplished in different forms. Some early examples of this include the USNRC codes RELAP5/PARCS [21], [22], and ATHLET/DYN3D [23], [24], just to name a few. These coupled simulations typically consist of a few-group diffusion or transport solver for neutronics, loosely coupled to a subchannel or larger systems codes to simulate thermal hydraulics and thermal kinetics. The strength of these codes are often their computational efficiency and high accuracy when solving within the confines of the problem they were developed to solve and on which they were validated. These codes tend to struggle or outright fail when presented with novel reactor designs featuring non-Cartesian patterns or non-typical fuel and moderator choices [25].

Recent neutronic-thermal coupling techniques tend to leverage the use of unstructured meshes and the ability to generalize the reactor geometry. MOOSE is a framework that allows for the addition of other physics systems, including subchannel codes [26] and deterministic neutronics solvers [14], [27]. This lets an engineer go from mesh generation of a core design all the way to neutronic-thermal hydraulic results without leaving the MOOSE environment.

Deterministic codes will always have a dependence on the validity of their cross-sectional data, however, and lack of understanding in how the used group cross-sections are generated can drastically affect results. Instead, continuous energy neutronics codes such as OpenMC can be used for the neutronic simulation. This is an increasingly common technique in the field of multiphysics coupling [28]–[31], and is the technique this thesis utilizes. One additional part of the appeal for the MOOSE-OpenMC coupled choice is that both codes are open source with large international user bases, which has been used for extensive code validation and developed numerous workflows. The trade-offs of this method are a significant penalty in computational time required, and the introduction of stochastic uncertainty to the coupled system.




Figure 2.2: MRaD Core Geometry

2.3 Core Geometries

The different heat pipe microreactor geometries studied in this thesis are listed in this Section, detailing their material composition and operation parameters. Note that the MRaD and SR geometries are slightly altered for their use in Chapter 4, and these changes are highlighted in section 4.3.

2.3.1 MRaD Core

The MRaD core designs originates from the work of Stauff to describe a microreactor concept under analysis by the NEAMS Micro-Reactor Application Drivers [32], [33]. This core has a rated thermal power of 2 MW, with the slice geometry simulated prescribed for 345.6 kW thermal. There are some material changes from Stauff’s work made in order to align with the purposes of this work. Firstly, the original UCO-form TRISO fuel is homogenized in the OpenMC simulation in order to reduce computational complexity. Homogenization cannot exactly replicate the neutronic behavior that TRISO fuel has [34], but mass and mass ratios are preserved in order to maintain general neutronic characteristics.

Secondly, the moderator is changed from YH_2 to $\text{ZrH}_{1.67}$. Given that half of this thesis’ work is on the topic of hydrogen migration in zirconium hydride, this is a requirement to utilize this geometry. Additionally, YH and ZrH are often mentioned in the same contexts for their moderating capability and hydrogen migration aspects [35]. In a rough simulation comparison of the core with YH versus ZrH as a moderator, minor differences were observed in neutronic response. An estimation of hydrogen number density in YH_2 resulted in 0.0627 atoms per cm-barn, and the value for $\text{ZrH}_{1.67}$ is 0.0613 atoms per cm-barn, explaining why the materials should behave similarly as moderators.

Table 2.1: Material summary for the MRaD core geometry.

Feature	Material	Notes
Fuel	UCO	TRISO fuel with 40% packing fraction at 14% ^{235}U enrichment.
Moderator	$\text{ZrH}_{1.67}$	-
Heat pipe	Potassium	Neutronically represented homogeneously with the stainless steel envelope
Monolith	Graphite	-
Reflector	Boron Iron	-
Absorber	B_4C	-
Air	Natural Air [36]	-

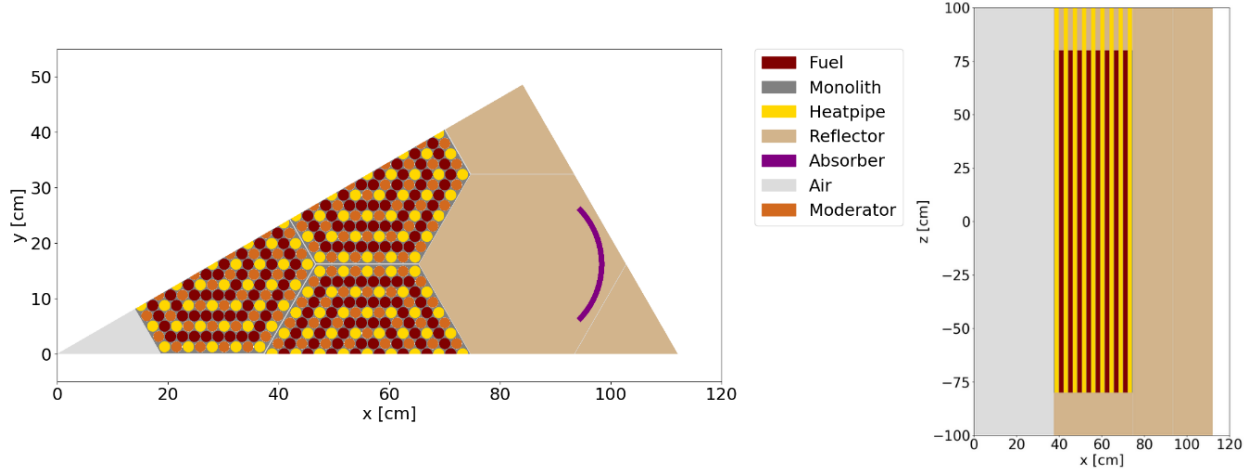


Figure 2.3: Simba Core Geometry

Table 2.2: Material summary for the Simba core geometry.

Feature	Material	Notes
Fuel	UN	19.75% enrichment.
Moderator	ZrH _{1.67}	-
Heat pipe	Potassium	Neutronically represented homogeneously with the stainless steel envelope
Monolith	Graphite	-
Reflector	Beryllium Metal	-
Absorber	B ₄ C	-
Air	N ₂	-

2.3.2 Simba Core

The Simba reactor geometry originates from the work of Terlizzi and Labourè [37] on the Simplified Microreactor Benchmark Assessment problem. Like the MRaD geometry, some changes to materials have been made. Once again, the YH₂ moderator has been replaced with ZrH_{1.67} for the purposes of Chapter 5. In addition, the heat pipes used in the work of Terlizzi are sodium based, while the heat pipe simulation via MITHP (described in Section 3.3) is potassium-based.

Unique features of this core versus the next introduced core and the MRaD cores include air gaps between each assembly and a gap around each moderator rod. The gap surrounding the moderator is thermally represented with radiative heat transfer in MOOSE. These gaps influence the thermal profiles of the core, with the moderator gap being particularly influential on the results for the Simba core seen in Chapter 5.

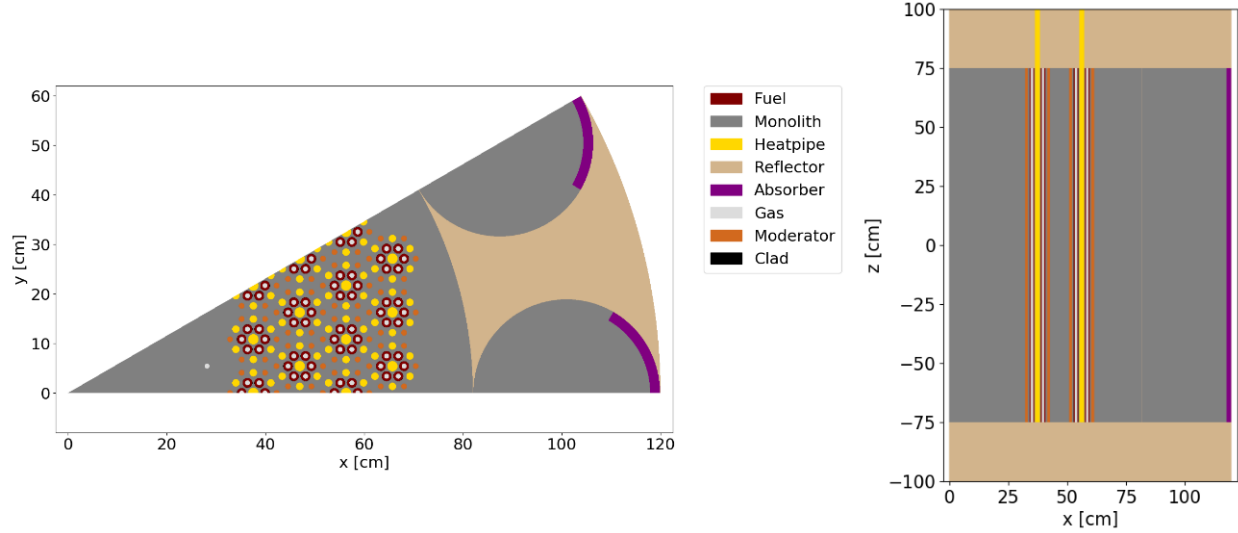


Figure 2.4: SR Core Geometry

Table 2.3: Material summary for the SR core geometry.

Feature	Material	Notes
Fuel	UO ₂	5% enrichment.
Moderator	ZrH _{1.67}	-
Heat pipe	Potassium	Neutronically represented homogeneously with the stainless steel envelope
Monolith	Graphite	-
Reflector	Magnesium Oxide	-
Absorber	B ₄ C	-
Cladding	FeCrAl	-
Air	N ₂	-

2.3.3 SR Core

This reactor geometry was developed by Professor Koroush Shirvan at MIT and has been referred to as "Shirvan's reactor" in previous works [38]. As such, this geometry will be called the SR core design. The specific OpenMC design of the core was created based on the work of my colleague Carmen Crawford [39].

This core is the only core studied with annular fuel, and also the only core represented with cladding on the fuel. The influence of these design choices are remarked on in the results sections of each chapter, in particular for hydrogen diffusion in Chapter 5. This core is originally rated for 15 MW thermal power, but for this work that power was decreased to 2 MW thermal in order to keep the moderator temperature below limits mentioned in Section 5.1.3, pertaining to ensuring the zirconium hydride remains in its δ -phase.

Chapter 3

General Methods

3.1 MOOSE

MOOSE is a finite element physics framework developed by several U.S. National Laboratories [14]. It features extremely parallelizable code for high efficiency on remote computing clusters, along with well documented and user-friendly methods for implementing one's own physics kernels. MOOSE is used for all thermal conduction, thermo-mechanics, and hydrogen migration simulation in this work, relying on modules that are readily available as part of the open-source MOOSE environment [40], [41].

All simulations using MOOSE utilize the finite element method with Lagrangian variable representation. The Newton scheme of non-linear solving was applied, with the default MOOSE Jacobian preconditioning.

3.1.1 Non-conservation

One thing to note about finite element (FE) solves is that there is an element of non-conservation that can potentially occur, depending on the problem layout and the specific finite element method chosen. Continuous finite element methods are commonly assumed to be globally conservative, but not necessarily locally conservative, while the finite volume (FV) method is conservative in both. This is due to finite volume equation typically originating from an integral form of the conservation equation in use, which forces local conservation on the volume that is solved [42]. Locally conservative methods are by definition also globally conservative.

Instances where local conservation is necessary are often fluid-related simulations that deal with shocks or steep discontinuities [43]. The work of this thesis is more reliant on global conservation, as this reflects a preservation of energy and total hydrogen in the system. Unfortunately, the use of Dirichlet boundaries in all MOOSE simulations in this work cause some global non-conservation. The strong enforcement of Dirichlet boundaries overwrites the original finite element equations at these boundaries without regard to the flux required to achieve these nodal values, thus the solution may not satisfy conservation [44], [45]. A way to avoid this is to use Neumann boundary conditions, which pertain to the solution derivative at the boundary. MITHP, the utilized heat pipe simulation code described in

Section 3.3, functions by reading a heat flux profile and enforcing a Dirichlet temperature profile on the heat pipe boundary, so some amount of error will be present in each MOOSE simulation. In order to avoid this, MITHP would instead need to return either a heat flux or heat transfer coefficient with wall temperature. Both of these are potential solutions that may be explored in future works.

The error caused by this non-conservation is reduced by refinement of the mesh or the use of higher order meshes. For every model studied in Chapter 5, this error accounted for a roughly 1% loss in total energy. For the MRaD model studied in Chapter 4, however, the loss was roughly 10% because a first-order mesh had to be used for reasons that will be discussed in Section 4.2.2. Again, it is worth noting that the finite volume method would eliminate the error here. MOOSE is capable of utilizing this method, but only with constant, monomial variables and thermal expansion is not available for finite volume solves. It's for these reasons that the work of this thesis only uses the finite element method.

3.2 OpenMC

OpenMC is an open-source Monte Carlo neutron and photon transport simulation code initially developed by the Massachusetts Institute of Technology. OpenMC can perform steady state k-eigenvalue calculations on complicated geometries with both continuous energy and multigroup with tallying capabilities for heat deposition[46]. OpenMC was developed with a strong emphasis on parallelism[12], leaning on the fact that the Monte Carlo particle transport method has some inherent advantages to parallelism[47]. This parallelization allows the software to scale well with increased usage of processors, which this work will make use of during the simulations.

One of OpenMC's strengths is its Python API, allowing the program to flexibly couple to a range of other Python-based packages. This allows taking advantage of high user-accessibility and access to data structures provided by packages like pandas [48]. Additionally, Coreform Cubit[49] has a Python API that the iteration scheme takes advantage of, allowing for mid-iteration geometry modification that is absolutely integral to the coupling methods described later in this thesis.

Cross-section temperature dependence was set via interpolation, with the default temperatures for all OpenMC models being 800 Kelvin. In the case of thermal expansion, volume-average temperatures are retrieved from MOOSE results. For the hydrogen diffusion study, OpenMC temperatures remained constant at 800 Kelvin in order to isolate the effect of hydrogen's spatial variation on neutronic response.

3.2.1 Heating Tally

Because OpenMC is significantly used in this thesis as a method of generating spatial power profiles in the reactor cores studied, the process of going from eigenvalue calculation to volumetric heating rates should be discussed. This process begins by tallying a "heating" score in every cell requested [46]. This heating score is calculated by

$$H(E) = \phi(E) \sum_i \rho_i \sum_r \kappa_{i,r}(E) \quad (3.1)$$

where $\phi(E)$ is the particle flux in units of particles per cm per source particle, ρ_r is the weighting factor for reaction r , and $\kappa_{i,r}(E)$ is the KERMA (Kinetic Energy Release in Materials [50]) of reaction r and isotope i . This KERMA coefficient is in units of eV-barn, functionally similar to a cross-section. These KERMA values are generated in NJOY [51] by evaluating the Q-value of the reaction and subtracting the energy component of particles that will be transported (e.g. prompt/delayed neutrons & photons) as well as neutrinos.

Note the units for the heating term $H(E)$ are eV per source particle. This is a reminder that in an eigenvalue simulation flux magnitude is unknown, similarly to how reactor power can be any level (within realistic bounds) at a specific eigenvalue. To convert these tally results to a MOOSE-readable input of volumetric heating rates, the score is altered via

$$q_n''' = \frac{H_n}{H_{\text{total}}} \frac{P}{V_n} \quad (3.2)$$

where n is the cell sampled, H_{total} is the sum of the geometry heating tally, P is the input power for the system, and V_n is the volume of cell n . This is essentially using the heating tally results to spatially distribute the input power for the geometry. The resultant volumetric heating rates q_n''' can be used as an input for MOOSE finite element solves.

3.3 MITHP

MITHP (MIT Heat Pipe) is a C++ translated version of the Los Alamos National Laboratory heat pipe code HTPIPE [19] with alterations made to integrate into the MOOSE finite element framework.

The code solves in an iterative manner, first sampling the input heat flux across the heat pipe surface. This heat input translates into a pressure differential along the heat pipe, which is then used to look up a saturation temperature in a table of potassium fluid properties for a discretized length of the heat pipe. The heat removal from the fictitious condenser section of the heat pipe is calculated, and the difference between input and output heat is minimized iteratively.

The original HTPIPE code is not designed for use in another iterative solver like MOOSE, which leads to two additional changes to the code being required.

Firstly, during the MOOSE thermal conduction solve, each iteration changes nodal temperatures as the problem converges on a solution. A result of this is that nodal temperatures around the heat pipe may decrease from one iteration to the next, which results in a negative total heat flux recorded on the heat pipe. HTPIPE's original code is not designed to deal with this situation, and a negative heat flux causes temperatures to spike to the arbitrary upper temperature limit of 1800 Kelvin. Clearly this is not conducive for accurate heat pipe simulation. Instead, a check for negative heat flux was added in MITHP. If this occurs, the simulation switches to a simple 1D conductive heat transfer model tailored to return a small decrease in temperature. I found that this would keep some modicum of realism in the

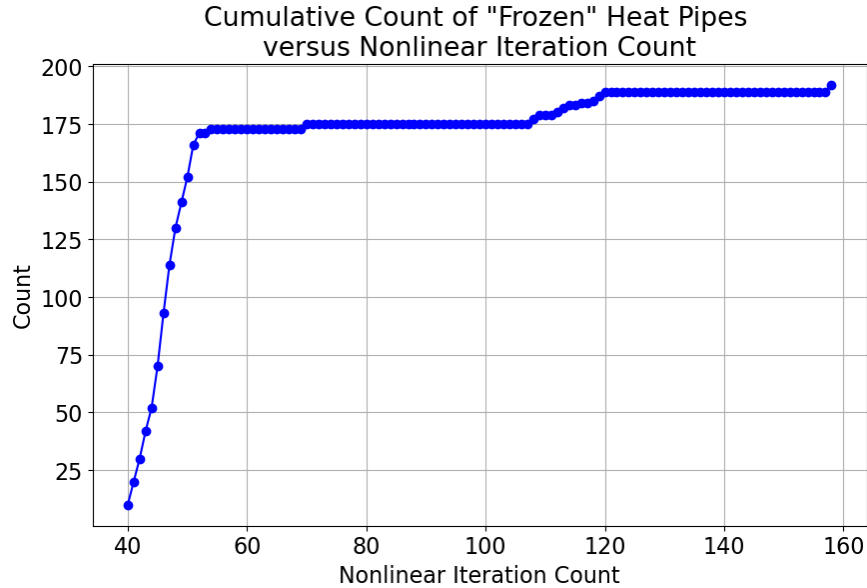


Figure 3.1: Cumulative count of "frozen" heat pipes versus nonlinear iteration.

physics of the heat pipe, while keeping the response to negative heat fluxes small enough to avoid instability in the overall simulation.

The second change necessary has to do with my mention of instability. MOOSE convergence is based on the residual from the weak form of the working equations for the problems solved in MOOSE. For the thermal conduction simulation, this residual can be seen as a value measuring the iterative change in temperature across the mesh. MITHP is solved in a Picard scheme in MOOSE because of its removed nature without the implementation of its contribution to the Jacobian of the heat conduction solve. This causes an undamped MITHP to spike residuals upon calculation, and results in instability that cannot be converged. To combat this, strict damping of the temperature change on the surface of the heat pipe is enforced, along with a "freeze" criterion. A short history of the heat pipe's temperatures are tracked, up to the last five iterations, and if the largest change in temperature amongst the five iterations is less than a set criterion, the heat pipe will "freeze" and lock its temperatures in place. This allows the overall problem to converge without needing adjustments to problem convergence parameters. That criterion is set to 0.1 Kelvin for the simulations included in this work, in an attempt to limit the error caused by treating the heat pipes in this manner. Figure 3.1 shows a cumulative count of the number of "frozen" heat pipes versus nonlinear MOOSE iteration.

These changes, combined with the MOOSE implementation method described in Section 3.3.1, result in a heat pipe representation that can be used for thermal MOOSE solves. The computational cost of MITHP is relatively low; in a test case with a single heat pipe, the MITHP solve represented 7% of the time spent simulating. This does not include the increase of simulation time due to the problem becoming non-linear, however. When comparing between a simple heat conduction problem with the MITHP boundary versus a Dirichlet boundary condition, the transition from linear to non-linear caused a 1400% increase in simulation time.

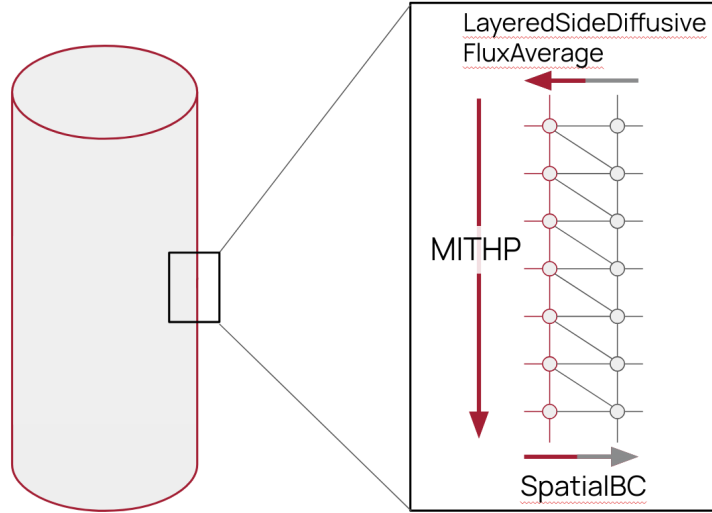


Figure 3.2: Graphical representation of MITHP heat pipe process.

3.3.1 MOOSE Implementation

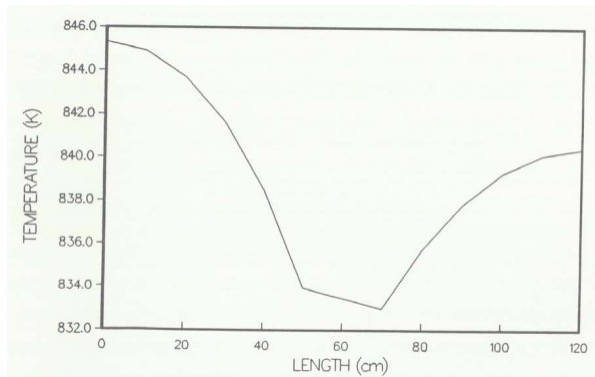
The MITHP heat pipe representation in MOOSE relies firstly on sampling the heat flux along the heat pipe boundary. This is accomplished with the existing MOOSE UserObject `LayeredSideDiffusiveFluxAverage`, which returns flux averages along an axially segmented surface. These heat fluxes are used as an input for MITHP’s code, which outputs a temperature profile in return. This temperature profile is sampled by a custom BC `SpatialBC` which sets a Dirichlet BC to the heat pipe boundary. A graphical representation of this flow can be seen in Figure 3.2

3.3.2 Comparisons

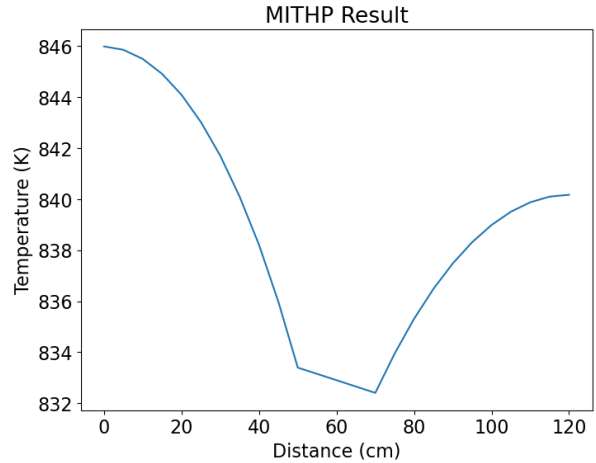
In order to use MITHP with confidence in the larger coupled simulations that will be described in the following chapters, there must be agreement established between the operation of MITHP and the current heat pipe models that exist. The first of these models to compare MITHP against is its originator, HTPIPE. Being essentially a translation of HTPIPE into C++ with some changes in order to fit the MOOSE environment, MITHP should be able to replicate HTPIPE results without problem.

The HTPIPE documentation includes an example problem featuring a flat power input into the condenser summing to 15 kilowatts, which can be replicated with MITHP. The results of each heat pipe’s simulation is shown in Figure 3.3, and features clear agreement between MITHP and HTPIPE. Again, this should be the expected outcome as MITHP is derived from the HTPIPE code.

A modern heat pipe code used in contexts similar to the work of this thesis is SOCKEYE [20]. This heat pipe simulation tool fits in the MOOSE framework, and is particularly strong at simulating heat pipe response in transient settings. Because I did not have access to SOCKEYE due to licensing disagreements, no direct comparison can be made. However, the results of the MRaD core provided in the VTB Repository [52] include the input and




(a) Plot of HTPIPE results from [19].



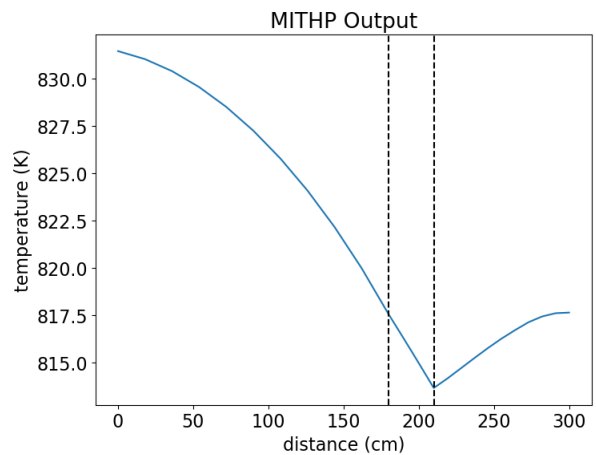
(b) MITHP results using HTPIPE input parameters.

Figure 3.3: Comparison of HTPIPE and MITHP temperature profile using the same inputs.

output of a SOCKEYE run. This input problem, a flat input of 1.8 kilowatts to the heat pipe, can be translated to MITHP and a comparison can be made, as seen in Figure 3.4. This comparison shows some differences, namely an overall higher temperature recorded by MITHP, however, the general shape is conserved. The slightly higher temperature is likely due to slight differences in parameters within the codes. Note that there is a clear difference in the condenser region, which I believe has more to do with the underlying HTPIPE code's simplicity.



(a) SOCKEYE heat pipe results from [52] for the MRaD geometry.



(b) MITHP results using similar inputs.

Figure 3.4: Comparison of SOCKEYE and MITHP temperature profile using similar inputs. Dashed lines indicate transitions from condenser to adiabatic region to evaporator.

Chapter 4

Thermal Expansion

4.1 Background

The first of the two additional physics included with neutronic-thermal coupling is thermal expansion. Thermal expansion is of interest in the context of microreactor design and neutronic-thermal coupling because of the uniqueness that microreactor designs have given their smaller geometry.

Reactor geometries are not stationary when increasing in temperature, they expand and deform in different ways depending on their environment and the specific design. Microreactors exist in a novel context, where size of the core exists as a selling point [1], and changes in that size may be important to note. Additionally, while conventional reactors leverage their size (along with other design choices) in order to flatten power shape, microreactors will inherently live with a more peaked power shape due to their size and unavoidable leakage. Changes in dimension may be more impactful on such a system with a strongly peaked flux profile. In particular, one can expect stronger changes in leakage due to geometric deformation in a smaller reactor compared to a large reactor.

For these reasons, this chapter details the work undertaken and results of coupling thermal expansion with neutronic-thermal simulation of multiple heat pipe microreactor designs.

4.1.1 Thermal Expansion

Thermal expansion is a well defined topic, being a phenomenon that humans have recorded for hundreds of years and have experienced for far more. Thermal expansion is the reason for expansion joints seen in bridges, the reason for the development of continuous welded rail for railway travel, and the reason power lines sag on hot days [53]. Thermal expansion occurs on the molecular level as increased temperatures increase the vibrational frequency of each atom. This vibration leads to an increase in distance from the surrounding atoms because of the anharmonicity between the attractive and repulsive forces that govern the lattice. A visual representation of this effect can be seen in Figure 4.1, a plot of the empirically derived Lennard-Jones 12-6 potential [54]. Here r is the distance between atoms, σ is the distance at which there is zero potential, V is the potential, and ϵ is the depth of the potential well. Increasing temperature acts as "filling" the potential well.

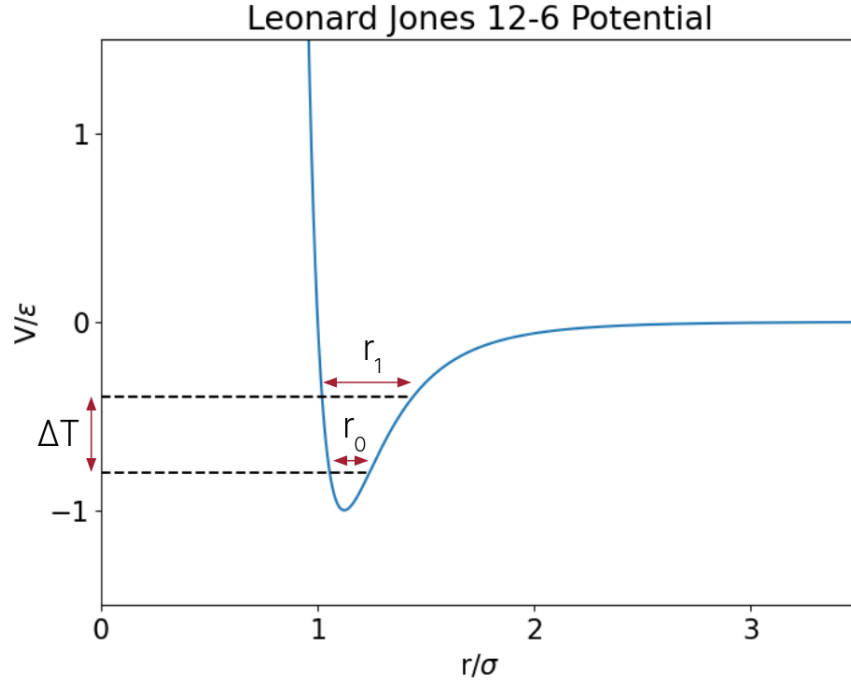


Figure 4.1: Lennard-Jones Potential for a paired atom, showing that an increase in temperature results in an increased atomic spacing.

Similarly to the Lennard-Jones Potential, thermal expansion is a material trait described with empirically-derived formulas. Thermal expansion properties are typically described with a linear expansion coefficient,

$$\alpha_L = \frac{1}{L} \frac{dL}{dT} \quad (4.1)$$

where L is the sample length and T is the temperature. Because the thermal expansion coefficient is most often dependant on temperature and the change in temperature is usually given in a problem, one can use Equation 4.2 to determine the relative change in length, most commonly referred to as (thermal) strain or ϵ .

$$\frac{\Delta L}{L} = \exp\left(\int_{T_i}^{T_f} \alpha_L(T) dT\right) - 1 = \epsilon \quad (4.2)$$

For an isotropic material where expansion is equal in all three orthogonal directions, the volumetric expansion coefficient is simply $\alpha_V = 3\alpha_L$. However, some materials are anisotropic in their expansion, particularly crystal lattices with noncubic symmetry [55]. These materials require a linear expansion coefficient defined for each axis of symmetry, which themselves may change orientation depending on the temperature.

The core of a nuclear reactor is an environment with extreme temperatures and sees significant thermal expansion of its materials, isotropically and anisotropically. Quantifying the per-degree effect of temperature variation on neutron balance via a reactivity coefficient has been a longstanding practice for reactor engineers. Depending on the study, thermal

expansion is included in this measurement. In traditional light-water reactors at steady state, the effect of thermal expansion on eigenvalue is fairly minor, with the largest effect being a change in the leakage of fast spectrum neutrons [56].

In spite of its relatively low impact on neutronic behavior, thermal expansion is still very much a topic of analysis for reactor engineers. This is because of the somewhat more mundane worry of material failure due to thermal stress, either within a material itself or at a contact point with another material. Some of the more notable concerns of this nature that exist for traditional nuclear reactors are:

- During fast transients like a control rod ejection, the differing thermal expansion coefficients between fuel and its cladding has been shown to lead to fracturing and clad "ballooning" [57].
- During regular operation of core, heat conduction through fuel pellets is largely radial, leading to non-uniform thermal expansion and radial cracking which degrades fuel thermal conductivity and increases fission gas release [58], [59].
- When using graphite at high temperature, the release of Wigner energy exacerbates expansion and requires designers to balance designed space for expansion and the need for heat removal. [60], [61].

Graphite in particular is an interesting problem when tackling thermal expansion, as graphite is notably asymmetric in its expansion vectors. The coefficients of expansion are dependent on the manufacturing process, but one study gives the two differing crystallographic expansion coefficients as $\alpha_a \approx 26.5 \times 10^{-6} \text{ K}^{-1}$ and $\alpha_c \approx -1.5 \times 10^{-6} \text{ K}^{-1}$ [62]. Given that graphite is typically used as a large-scale component of the core, like a monolith, this conflicting expansion and contraction of graphite requires that crystal orientation be considered during design.

For the purposes of this chapter, neither anisotropic expansion nor potential mechanical failure are taken into account, although anisotropy may be possible to model by defining material lattice planes. Instead, the focus of this work is on the impact of non-uniform, isotropic-in-cell, elastic thermal expansion on core neutronic behavior and simulating the thermal-neutronic feedback loop. It was previously mentioned that thermal expansion has a minor impact on eigenvalue for traditional light water reactors, derived from a change in fast-spectrum leakage. In a microreactor setting, like the reactors studied in this thesis, the smaller and more leakage-dominated systems are likely to experience more significant responses to thermal expansion. This is a primary impetus for the work of this chapter.

4.1.2 Geometric Deformation in Nuclear Reactors

It is important to note that although thermal expansion is the only deformation mechanism included in this study, it is not the only method of geometric deformation that nuclear reactors experience during operation. A notable one is fission-induced swelling of the fuel material. This process occurs as fission products collide and decelerate in the fuel material, creating vacancies and displacements as well as depositing themselves throughout. These fission products are a mix of solid and gaseous in form, and result in a swelling effect for the fuel [63].

The gaseous fission products are especially accelerative for the swelling process, as the fission products will migrate to grain boundaries and cavities. These gasses then experience larger thermal expansion than their solid-state counterparts, increasing expansion and encouraging material fracture [64]. Because of the inherent risk of highly radioactive, gaseous fission products escaping the fuel, the U.S. Nuclear Regulator Commission includes fission gas release and swelling modeling as a part of the fuel qualification process [65]. This has led to fuel swelling being a well documented and modeled phenomenon [66].

While nuclear fuel is unique in that fission products drive material expansion, the process of voids or cavities occurring due to irradiation is not isolated to the fuel material. Steel in particular has been characterized to undergo swelling due to void formation from neutron irradiation [67]. Certain absorbers also experience this swelling, due to their nature of having inherently high collision-rates [68]. These mechanisms are mostly found in high-fluence environments, however, such as fast reactors.

An additional source of deformation in the reactor core comes from volume expansion or contraction due to metamorphic phase change. Alloys such as zirconium hydride, the primary topic of Chapter 5, have detailed phase diagrams that describe transition points between different phases depending on temperature and their specific stoichiometric ratio [69]. Because these phases changes imply a reordering of lattice structure, the overall volume of the material changes as well. ZrH in particular sees a 14% decrease in density when transitioning from β to δ -phase [70].


While this study neglects these previously mentioned methods of geometric deformation that do regularly occur in reactor cores, the sole study of thermal expansion and its effects is useful as it provides a baseline for similar deformation cases. Being able to define that x centimeters of axial displacement causes y change in eigenvalue allows the next study that includes some of the previously mentioned methods to estimate neutronic impact even before simulation. Additionally, the method of coupling described in this chapter are applicable for all the above deformation modes, given a model exists that describes the deformation.

4.1.3 DAGMC

As has already been mentioned in the content of Section 2.2, the largest hurdle involved for neutronic-thermal coupling is reconciling the difference in native environment that each type of simulation exists in. Thermal coupling is typically performed on an unstructured mesh geometry via the finite element method where calculations are performed iteratively over quadrature points located within mesh elements. Neutronic calculations, when simulated via Monte Carlo particle transport, are typically performed on CSG geometries where surfaces are defined with analytical geometric expressions and cell spatial regions are defined by boolean half-spaces. In order to couple the two together, something must be altered. Chapter 5 describes a process of altering the information transferred into a format readable by both MOOSE and OpenMC, while this chapter will utilise Direct Accelerated Geometry for Monte Carlo (DAGMC) to instead alter the geometry that particle transport occurs on.

DAGMC is a software library originally intended for neutronics modeling of fusion reactor geometries, allowing the user to translate CAD geometries into Monte Carlo-solvable geometry inputs[11]. The software has since seen use in both fission and fusion applications, and features an ability to couple with the Sandia National Laboratories-based Coreform

Cubit software to automate CAD-to-CSG transfers[49]. By leveraging the built-in graphical tessellation feature inherent to the Cubit software, a coarsened facet-based geometry can be exported. Recent collaboration between Coreform and the DAGMC team has enabled exporting DAGMC geometries directly from coarse meshes, a distinct upgrade over the graphical-faceting method [71]. An example of the facet geometry versus CAD geometry can be seen in Figure 4.2, highlighting the surface-only nature of the facet geometry. This facet geometry replaces the geometry block of information traditionally used for OpenMC problem description. In order to do so, temperatures and material assignments are made in Cubit via groups and are stored within the facet geometry.



(a) CAD mesh representation.

(b) Facet representation.

Figure 4.2: MRaD core geometry comparison of CAD mesh versus facet geometry, clipped in the axial mid-plane.

The DAGMC geometry is essentially a geometry made solely of planar bodies, which OpenMC is well-equipped to perform transport on. DAGMC’s integration with OpenMC is well documented and there exists a robust literature documenting its use [13]. While the planar bodies work well for representing tessellated surfaces, this does still mean that surfaces that were originally second-order or higher, such as the outer surface of a cylinder, will incur severe computational penalties. When running a simple reflective-boundary pin-cell problem, I saw a roughly 92% slowdown in active and inactive particle simulation rates. This decrease is dependent on the geometry of the problem as well as tallies and transport settings, and several compilation options have been suggested as a method of decreasing this slowdown [72]. The details of the specific use of DAG-OpenMC for this chapter’s purposes are expounded on in Section 4.2.3.

4.1.4 Literature Review

The concept of coupling thermal expansion to neutronic-thermal simulation is not a novel one, in fact the last five years have seen a number of papers published on this topic featuring a wide range of coupling methods and reactor geometries [73]–[82]. Of particular note is the work by Stauff et. al in [33], who used MOOSE coupled with Griffin [27] to simulate thermal stress and strain on the MRaD geometry. Stauff notes a maximum axial displacement of 1.5 cm and a maximum stress located between neighboring heat pipes, but doesn't note any impact to neutronic or thermal behavior due to thermal expansion.

In a study on a fast-spectrum reactor, Walker et. al. reported an impact of roughly 7 pcm per Kelvin due to combined axial expansion of fuel and radial expansion of reflector and cladding [80]. In experimental testing of the KRUSTY heat pipe space reactor, Poston et. al. reported that roughly 90% of reactivity feedback was related to thermal expansion of the core, likely due to its small size [10]. The cumulative sum of these works suggest that not only does thermal expansion have an impact on core behavior, but also the smaller the core the larger this impact may be.

The work described in this chapter aims to quantify the thermal expansion experienced in two different heat pipe microreactor geometries, as well as that thermal expansions' downstream effect on neutronic and thermal behavior of each respective core. In addition, the impact of heat pipe failure on these quantities is analyzed.

4.2 Methods

The sections in this chapter detail the methods used for solving thermal expansion coupled with particle transport and heat conduction. The general method of coupling between OpenMC and MOOSE is the first topic, detailing data transfers and other intermediary steps required. Following that is an explanation of the general MOOSE input layout, covering what physics kernels and input blocks are used to simulate heat conduction and thermal expansion. Also included in this section are some notes on difficulties inherent with this form of modeling in these reactor geometries and their solutions. In the final section, the DAG-OpenMC simulation phase is explained with more detailed, including an implementation detail in the simulation setup that arises because of heat pipe representation in the core geometry.

4.2.1 Coupling Scheme

A flowchart diagram showing the coupling scheme is included in Figure 4.3, showing the flow from mesh generation to final results compilation. The initial work done before the main iteration is primarily focused on mesh generation and parsing. Mesh generation is performed using the MOOSE Reactor module [40]. This module allows regularized mesh generation for reactor core geometries, in particular geometries with repeated features such as Cartesian and hexagonal lattices.

The process of parsing the mesh occurs following this generation. Because the DAG-OpenMC geometry is derived from the mesh, all cell references now must be references to mesh volume IDs. Building the arrays of IDs that establish what materials are found in each volume as well as calculating and tabulating volumetric data for each mesh volume is handled during the parsing step. This is accomplished via Coreform Cubit’s Python API, without which a manual inspection of the mesh would be required. A dictionary cataloguing each core material and its properties, along with mesh volume IDs that feature that material, is instantiated and populated at this time as well.

The beginning of the main iteration loop and start of the neutronic side of the process involves first generating the DAGMC facet geometry. This step leverages the previously mentioned Cubit Python API to perform automatically, and material and temperature assignments are attached to their respective volume IDs by references the arrays built during the parsing step. Following this, the OpenMC XML input file is generated, featuring two changes from typical OpenMC input layout. Firstly, because volumes will expand and contract during iteration and mass must be conserved, material density will be updated every iteration step. An unfortunate side-effect of this is that a unique material is needed to be generated for every volume in the mesh, as a unique density means a unique material definition. More recent DAGMC updates allow density to be applied during the facet generation step which would alleviate this, but for this work that was not yet available. Additionally, if depletion were utilized with this scheme, a unique material would be necessary regardless. For the MRaD geometry this means there are 466 unique material definition in the OpenMC input XML. This causes a minor slowdown during cross-section lookup, but is negligible when compared to the slowdown inherent to DAGMC-OpenMC. The second change from

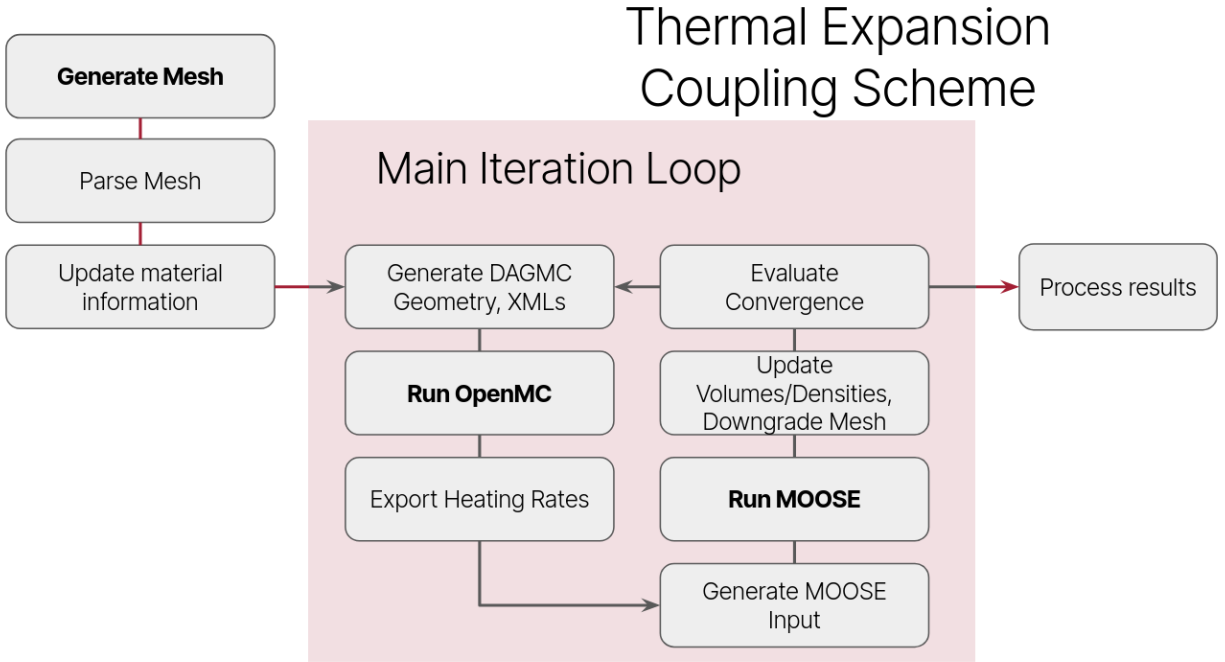


Figure 4.3: Detailed thermal expansion coupled simulation scheme. All processes not in bold are performed via Python script.

typical OpenMC input layout is described in more details in Section 4.2.3.

After OpenMC has run and output a result, the cell heating tallies are converted to volumetric heating rates. The MOOSE input file is generated via a template, applying the OpenMC-generated heating rates to each of its respective block in the mesh. The specific blocks used in the MOOSE input are described in Section 4.2.2. After MOOSE has iterated and converged on a solution, a parse of the output deformed geometry occurs. This parsing operation samples and tabulates the new cell volumes, and exports a "downgraded" mesh. This downgrade is a conversion from second-order mesh to first-order, necessary to simplify the next DAGMC geometry that will be generated. During my testing I found that generating a DAGMC facet geometry based on a second order mesh significantly impacted computation time without substantially affecting the neutronic result.

Convergence is then evaluated by inspecting the relative residual of the eigenvalue and the L2 norm of all relative residuals of cell volume changes. In this manner, convergence can be assessed for both the OpenMC solve and the MOOSE solve. If convergence is not achieved, the iteration continues by generating a facet geometry based on the MOOSE-output deformed mesh. If the eigenvalue and volumetric changes are assessed to be converged, results are compiled.

4.2.2 Thermal Expansion in MOOSE

One of MOOSE's strongest selling points for this sort of study is its library of existing modules that cover a wide range of physics kernels. Included amongst those is the Solid Mechanics module, which contains the specific material and physics blocks used to simulate

thermal expansion in the mesh. A representation of the MOOSE input file is included in Figure 4.4. This figure does not include post-processors and the material blocks used to apply thermal and mechanical material properties. An explanation for each subset of the input, corresponding to the connected letters seen in Figure 4.4, is as follows:

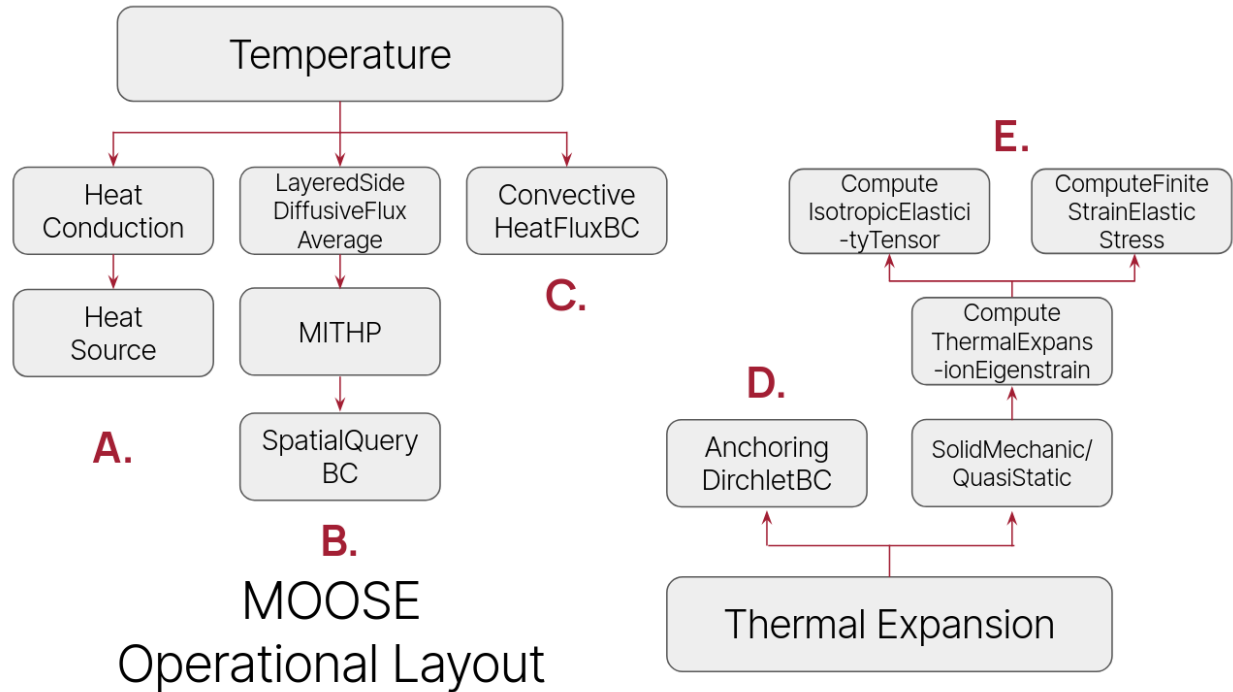


Figure 4.4: Graphical representation of the structure of the MOOSE thermal expansion solve, with some aspects omitted. Letters correspond to more detailed explanations that can be found in Section 4.2.2.

- A **Kernel activating heat conduction and the input volumetric heating rates for every mesh volume.** Each volume has the OpenMC-generated volumetric heating rate applied via the HeatSource block.
- B **Heat pipe simulation.** First an axially-discretized surface heat flux calculation is performed on each heat pipe surface. Then a unique custom UserObject MITHP is called, linked to these fluxes as an input, in order to calculate a resulting axial temperature profile. Finally the custom boundary condition SpatialQueryBC samples the temperature profile generated by MITHP and establishes a DirichletBC. Again, also unique to each heat pipe.
- C **Outer, top, and bottom boundary heat removal.** This is an arbitrary natural convection heat loss representation, mirroring the inputs used in [33] found on the Virtual Test Bed Github Repository [52]. This inclusion allows for a more realistic temperature profile for the outer air materials, given their otherwise isolated condition.

D Anchoring Boundary Conditions. Because the geometries studied in this work are slices of larger cores, displacement at these boundaries needs to be anchored with a `DirichletBC` in order to preserve symmetry. Note that this displacement should be still allowed in the direction tangential to the reflected surface, and only anchored in the direction orthogonal to the surface. I was not able to define this level of directionality in my MOOSE input, so all displacement besides axial is frozen at each reflective border, which leads to some inaccuracy in radial displacement. This over-constraint of the sides restricts the total radial expansion modeled. In a side case not shown in this thesis, after managing to orthogonally pin one of the sides, the radial expansion increased by roughly 50%. As will be seen in the results of 4.3, the overall magnitude of radial expansion is not dramatic, however. Modeling the entire core rather than a slice would remove this issue and accurately represent radial expansion in the geometry. A correct orthogonal pinning of the both the sliced geometry's sides would do the same.

E Instantiation and solving of thermal expansion-based eigenstrains, strains, and stress. The first block, `SolidMechanics/QuasiStatic` is a MOOSE Action object that instantiates all the needed data for a Solid Mechanics problem.

`ComputeThermalExpansionEigenstrain` calculates the eigenstrain tensor resulting from isotropic thermal expansion, given the temperature at the quadrature point versus the input stress-free temperature. These eigenstrains lead to stresses calculated by `ComputeFiniteStrainElasticStress` given the material properties computed by a `ComputeIsotropicElasticityTensor`.

The result of these inputs is a solution that thermally solves temperature in the system and displaces mesh node points based on thermal expansion eigenstrain tensors. The output mesh is now deformed, and can be converted via DAGMC for an OpenMC particle transport solve.

While the thermal expansion portion of the MOOSE solve does not involve any custom-made code, it does introduce two significant complications necessary to be navigated in order to solve the core geometry.

Mesh Collapse

The first of these complications involves when displacement of mesh node points leads to one node passing a neighboring node, inverting the mesh cell and causing a negative Jacobian and rapid simulation failure. I refer to this as "mesh collapse" because it most often occurs in areas where thermal expansion causes a narrow mesh to be "crushed" by the expansion of a neighboring part of the mesh.

A very common area this might be seen is in the air gap between a fuel pin and its cladding. During real life operation, that gap will close as the fuel heats up and expands from the burnup process, which is generally a positive occurrence as it increases heat removal from the fuel. If this example takes place in a mesh environment, the process of closing the gap becomes a problem. Meshes cannot have their elements or node points removed during simulation, only moved, and if the fuel makes contact with the cladding, then the air mesh volumes between are inverted. This causes negative Jacobians upon the next calculation sweep and immediate simulation failure. A similar example of this, where the air

gap surrounding the control drum in the MRaD geometry collapses during thermal expansion, can be seen in Figure 4.5.




Figure 4.5: Example of mesh collapse where thermal expansion caused the air gap between control drum and reflector to collapse in the MRaD geometry. The center, darker mesh surface indicates where parts of the mesh element have inverted.

Another instance where I believe a similar error occurs is when mesh quadrature points are very close in position but not a part of the same mesh element. For example, during testing with thermal expansion, simulation failure would often occur due to negative Jacobians being reported in TRI mesh elements, typically when those mesh elements were small. Increasing the size of the element or downgrading from second order to first order would sometimes resolve this problem. My suspicion is that at the center of the mesh block, all element edges radiate outwards from a single point and the quadrature points that are near each other but not on the same element may cross during deformation. This could also be a more mundane error related to TRI meshes and Libmesh [83]. Regardless, this provided a pervasive constraint on the mesh geometry and was one of the reasons why the meshes simulated in the results of this chapter are slightly altered from their forms presented in Section 2.3.

Both the aforementioned types of error were troubling for the microreactor geometries studied in this thesis, particularly because each reactor's mesh is generated from the MOOSE Reactor module, which by default uses the TRI meshing scheme for the center of cylindrical features. There is an option to instead use the QUAD scheme for some features, but others, such as when defining control drums, have no alternative. To navigate this issue, as well as the previously mentioned more conventional "gap collapse" case, each core geometry had some mesh simplifications applied, and the MRaD core mesh is simulated in first-order for the results of this chapter. The Simba core, on the other hand, was unable to function without simulation failure and is not included in the results of this chapter because of this.




Figure 4.6: SR mesh focused on one of the inner "dummy" assemblies. Highlights how tightly packed the inner TRI mesh elements are, which may cause overlaps during displacement and simulation failure.

Air Representation

The second complication to the MOOSE solve introduced by the inclusion of thermal expansion relates to the inclusion of air in the reactor geometries. For this description, "air" includes any gas representation. Air is featured in each of the core geometries studied in this thesis: The MRaD design has air replace the center assembly as well as outside the core, the Simba design features air between each assembly, and the SR design features air in the annulus in the center of each fuel pin.

These air gap representations complicate thermal expansion simulations, as the Solid Mechanics module is meant for solids, which is clear from the name. Mistakenly including air with thermal expansion and its relative Elastic modulus will result in Figure 4.7 and rapid simulation failure.

Instead, one would prefer to exclude these air blocks from the solid mechanics simulation. Unfortunately, I did not find it possible to perform thermal expansion in these geometries while excluding specific blocks of the mesh from deformation-related kernels and materials, though this should be possible in future MOOSE updates. The most obvious solution to this is to include the air mesh blocks in deformation without thermal expansion, and set their blocks' Elastic Modulus to a very low value, such that the air mesh will "give way" to the expansion of the other materials, like it would in real life. Unfortunately this causes the errors described in the previous subsection, where some mesh elements collapse due to this unimpeded expansion.

Instead, the best method I found to avoid simulation failure was in fact to remove the air blocks altogether. This removes the threat of mesh collapsing due to extreme expansion,




Figure 4.7: MRaD geometry with thermal expansion incorrectly represented in air mesh blocks.

and avoids needing to represent a non-solid object with physics meant for solids. The impact of this choice, beyond allowing simulation to continue, is twofold:

1. Firstly, the air has value as a thermal isolation. Although its conductance is low, it does serve to correctly represent the reactor thermal profile as it would be in the real system. While the air in the VTB geometry is generally negligible for this element, being located on the innermost and outermost portions of the core, the air in the SR geometry is found in the fuel pin annulus. Again, while the gas in the annulus isn't there for its thermal conducting properties, the lack of it likely does impact the shape or magnitude of the temperature profile the fuel will see.
2. Secondly, because the mesh is what becomes the geometry that particle transport is simulated upon, losing these mesh blocks means losing the cells for particle transport. Given that the gas is inherently low density, this impact on neutronics is likely to be insignificant.

4.2.3 DAG-OpenMC and Heat Pipes

Noted in Section 4.2.2, heat pipe representation during the thermal/mechanical solve in MOOSE is essentially a surface representation. The heat flux along a surface is used as an input to define what that surface's temperature profile will be. At no point is the interior of the heat pipe used in any valuable way. In fact, an inclusion of a material inside the heat




Figure 4.8: Depiction of the CSG heat pipe inserted into the CAD-tesselated heat pipe void. Note the gaps and overlaps caused by the polygonized geometry.

pipe negatively impacts the operation of the MITHP User Object, as it alters the heat flux recorded at the heat pipe boundary.

The traditional way of handling this is to remove the mesh blocks that describe the heat pipes, as only the surface of the blocks are needed. This is a process that similar works using SOCKEYE [20] for heat pipe representation have employed [33], [37]. For a coupling method that performs neutronics on the DAGMC facet geometry based on the mesh output, this would mean removing the heat pipes from the neutronic solve of the system. Unlike the previous case where air is removed from the system, heat pipes are neutronically an important feature. In the MRaD design, the heat pipes account for 10% of all absorptions in the geometry, due to the stainless steel casing and potassium fluid. Neglecting this would make the neutronic results significantly less realistic.

Instead, the solution I found was to insert heat pipes into the geometry on the OpenMC side by importing the DAGMC geometry not as the root universe, but instead as a cell. This is a somewhat complicated process of geometry definition, where the region that the DAGMC cell inhabits is defined by a set of core bounding surfaces and by excluding every heat pipe region. This is possible by taking advantage of the OpenMC-defined geometry that exists for the work in Chapter 5.

A side effect of this arises from the fact that the cylindrical regions for the heat pipes do not exactly fit the empty regions in the mesh geometry, due to the tessellation from meshing. This causes slight gaps at the boundary between the heat pipe and the monolith/reflector. An example of this can be seen in Figure 4.8. These gaps result in a minor increase of streaming from the core that wouldn't exist in real life. However, that effect is minor compared to the results of completely excluding heat pipes.

4.3 Results

The previous two sections note changes to microreactor core geometry that were made in order to avoid simulation failure and to enable an accurate coupling process. Namely, mesh simplification and air exclusion to avoid mesh-related simulation failure in MOOSE, and heat pipe representation in OpenMC to preserve heat pipe influence on the neutronic solve. Remembering that the Simba core is not included in the results of this chapter, Figures 4.9 and 4.10 show the change in geometry for all results included in this chapter. These changes result in a roughly 400 pcm difference in eigenvalue for both cores studied.

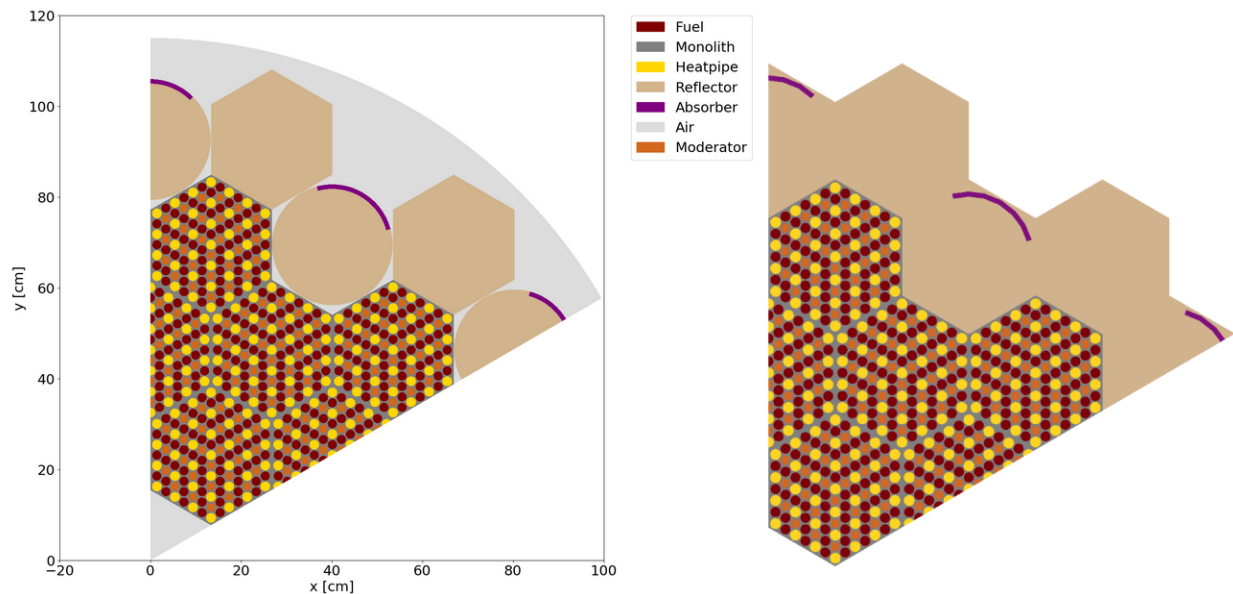


Figure 4.9: Comparison of the original MRaD geometry (left) with the thermal expansion form (right). Note the lack of air and transformation of control drums into contiguous reflector.

The results covered in this section cover the location and magnitude of displacement observed in the final converged iteration result, as well as that displacement's effect on neutronic and thermal behavior. This includes thermal expansion's impact on global statistics such as eigenvalue and leakage, as well as fuel pin powers and reactor temperatures. In order to isolate the impact of thermal expansion, the final converged solution will be compared with a solve that only transferred temperatures to OpenMC and did not deform the mesh. This way the impact of Doppler broadening can be excluded and the impact seen can be attributed wholly to geometric deformation. OpenMC particle count parameters can be seen in Table 4.1 and is consistent for each core study.

As previously mentioned, convergence of the coupled simulation was based on the relative change of the eigenvalue from OpenMC and the L2 norm of the relative residual of all volume changes from MOOSE. The eigenvalue details can be seen in Figure 4.11. Note that the original convergence criterion was set to a relative residual of $1.0e-4$, however the MRaD case fails to meet this. This geometry was still presumed to be converged, given the results seen in Figure 4.12, depicting the L2 norm of cell volume relative residuals. Both cores

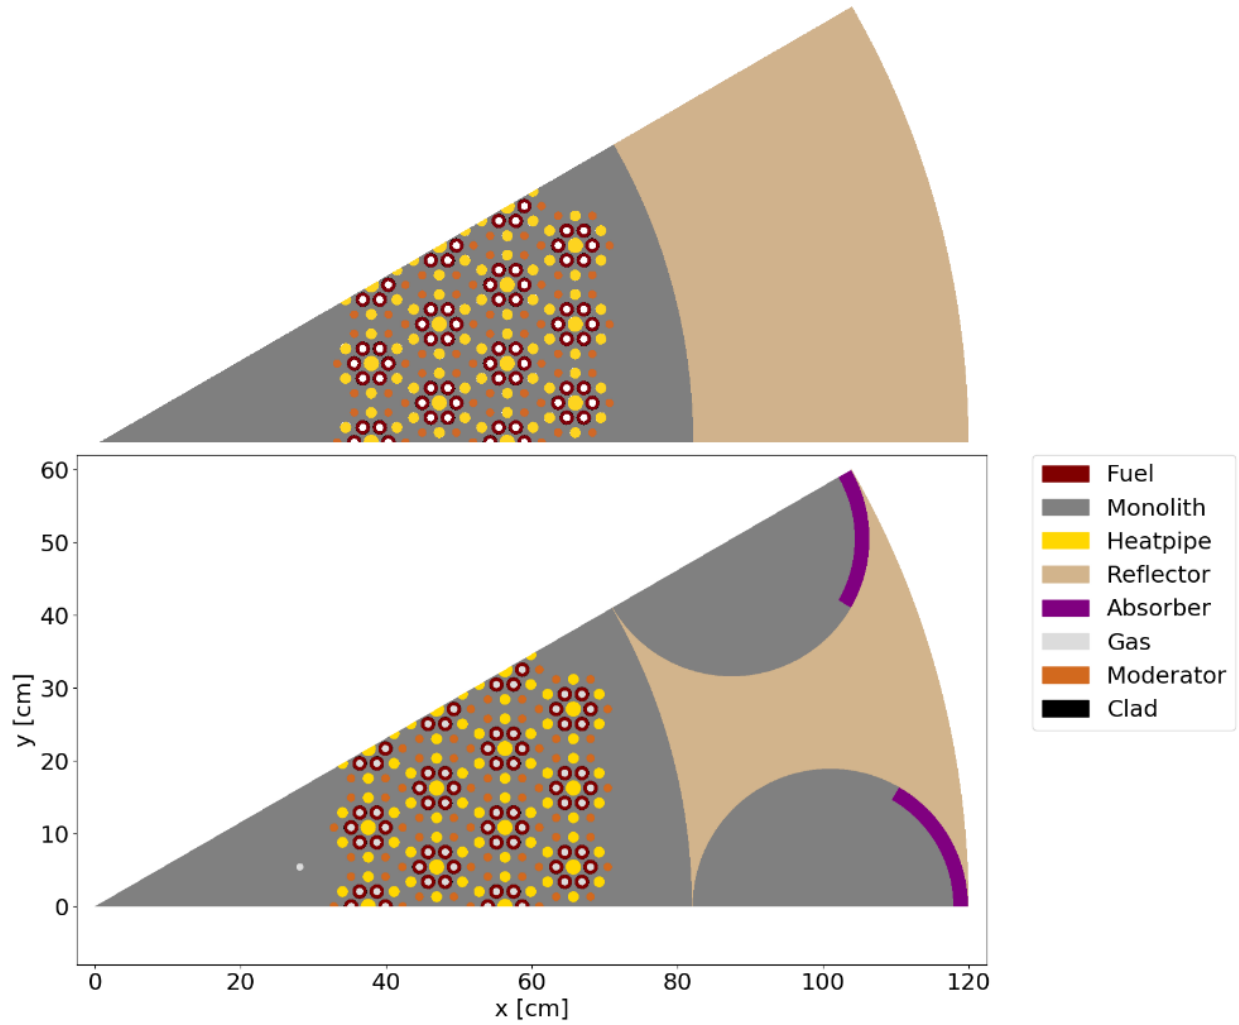



Figure 4.10: Comparison of the original SR geometry (below) with the thermal expansion form (above). Note the removal of control drums and shutdown rod space, as well as removal of gas from fuel annulus.

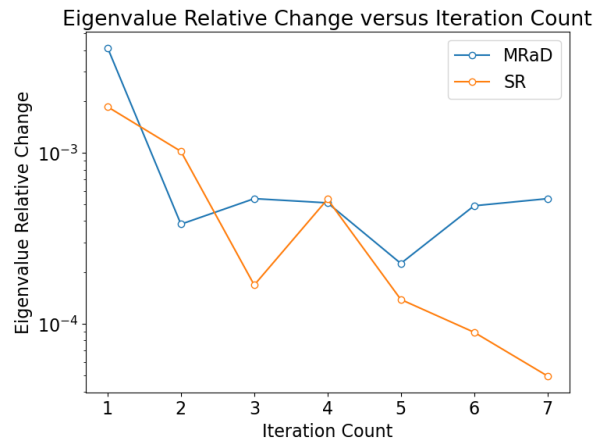
reached a value beneath $1.0e-4$ relating to the change in volume per iteration, with the SR case being larger due to more relative expansion in the fuel. It should also be noted that convergence here has an assumption of convergence in other parameters that have been introduced, such as mesh refinement. The correctness of these parameters are assumed in this work.

Table 4.1: OpenMC particle count settings thermal expansion studies.

Parameter	Count
Active Batches	150
Inactive Batches	50
Particles	50000



(a) Eigenvalue per iteration for MRaD and SR cores.



(b) Relative changes in eigenvalue per iteration for MRaD and SR cores.

Figure 4.11: k_{eff} and relative change of k_{eff} versus iteration count for MRaD and SR cores during thermal expansion coupled simulation.

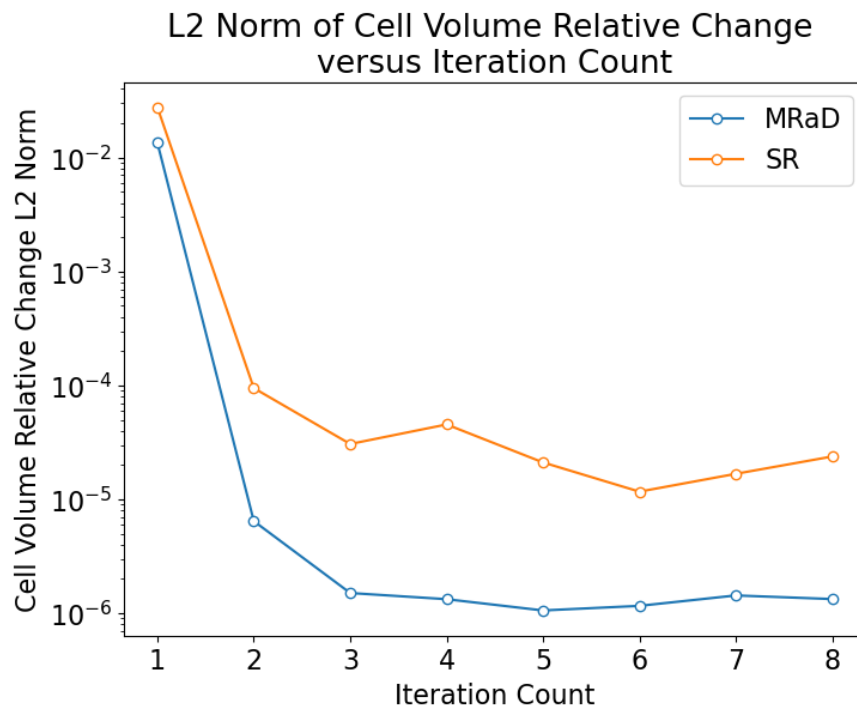
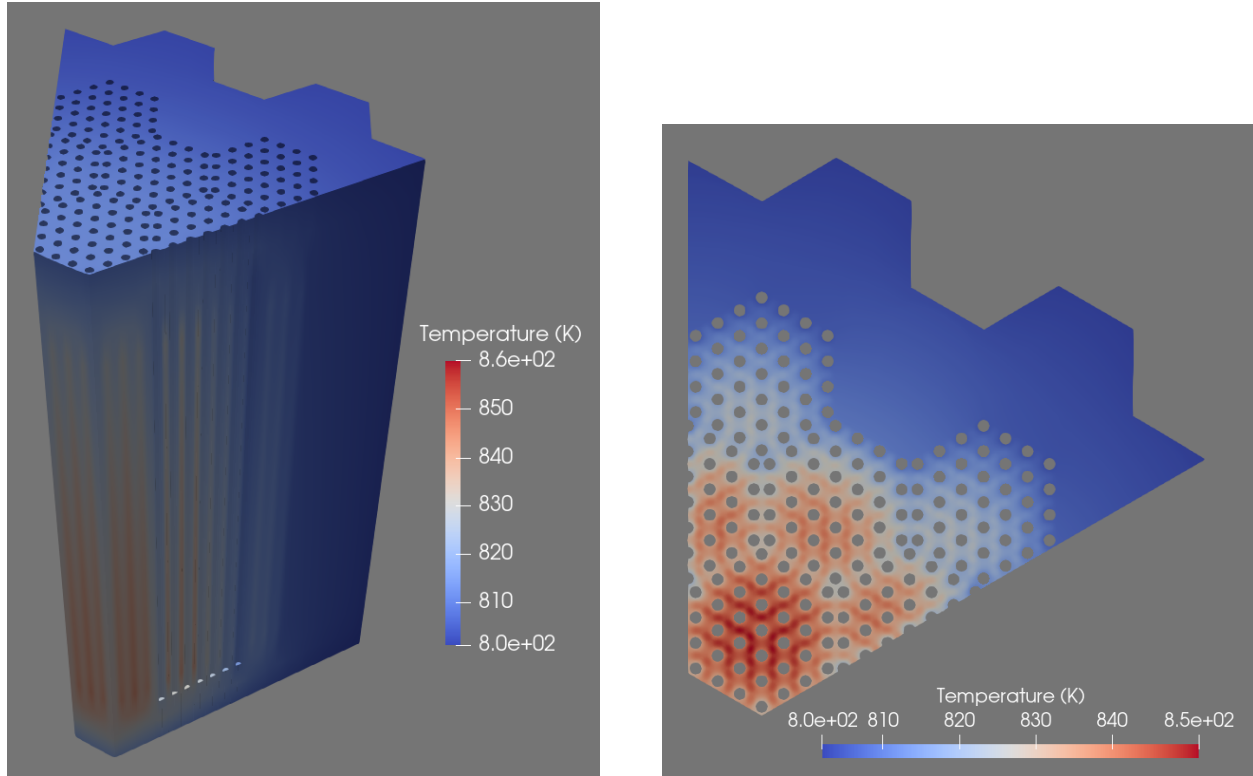


Figure 4.12: L2 norm of cell volume relative residuals versus iteration count for MRaD and SR cores during thermal expansion coupled simulation.

4.3.1 MRaD Core



(a) Isometric view of MRaD Core temperatures after final coupled iteration.

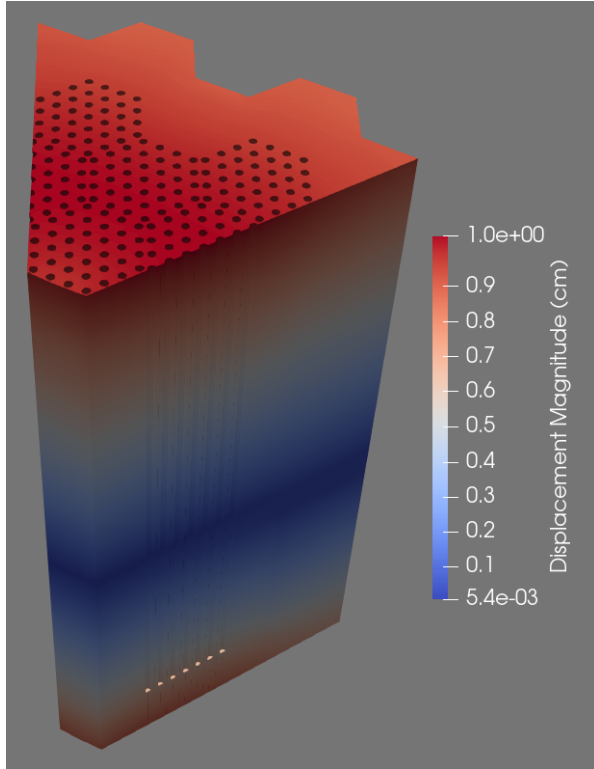
(b) Axial mid-plane slice view of MRaD Core temperatures after final coupled iteration.

Figure 4.13: MRaD temperatures after final coupled thermal expansion iteration.


The MOOSE-output temperature profile of the MRaD core can be seen in Figure 4.13. Note that temperatures are slightly lower than seen in the work of Stauff [33] and in the results of Chapter 5. This is due to the mesh elements being first-order, which increases the global non-conservation error from imposing Dirichet boundary conditions on the finite element solve. This topic is discussed in Section 3.1.1. By calculating the heat removed from the heat pipes and convective surface transfer and comparing this value to the input power, roughly 10% of power input is not accounted for. This causes a roughly 10 Kelvin difference in peak temperature, which is a low-enough value to consider negligible.

The calculated displacement due to thermal expansion can be seen in Figure 4.14. Clearly axial displacement is the dominating mode in the core, with a roughly 1 centimeter difference at each end. Radially, the reflector sees roughly 0.3 cm of outwards expansion. A mid-plane slice showing percent change in volume for each cell can be seen in Figure 4.16. The moderator rods experience the largest expansion of about 2%, and the reflector slightly below that.

Note that for these results, displacement is measured with the mid-plane as a reference. An alternative case was run with the bottom plane axially constrained, which resulted in Figure 4.15. This alternative displacement did not result in any difference in results com-



(a) MRaD displacement in centimeters after convergence.



(b) Mid-plane slice of MRaD showing displacement in centimeters after convergence.

Figure 4.14: MRaD displacement (cm) after final thermal expansion iteration. Displacement is relative to the axial mid-plane.

pared to a non-axially constrained simulation. Attempts were made to correctly constrain the sides of the geometry, as constraining them in x and y does restrict radial expansion. These attempts were only partially successful, but showed that the true magnitude of radial displacement may be on the order of 0.5-0.6 centimeters, which would not significantly deviate the results shown in this chapter.

Selected global neutronic statistics and the impact from including thermal expansion on these statistics are included in Table 4.2. The largest impact visible is a decrease of 5.1% in leakage due to capturing thermal expansion in the coupled simulation. The eigenvalue's 315 pcm increase and minor increase to thermal fission factor suggest that the decreased leakage led to an increase in moderation and thermal fission in the core. The inclusion of the deformed geometry also caused a minor softening in the core flux spectrum, with a 0.11% increase in thermal (<10 eV) neutrons and 0.40% decrease in fast (>10 eV) neutrons.

The effect on core pin powers can be seen in Figure 4.17. The peak pin power remained at the same level before thermal expansion, and no clear pattern of effect on the pin power distribution is noticeable. Overall, the fuel saw a 0.3% increase in energy deposited per source particle. Because thermal power in the core is assumed to be constant, this equates to an increase in power in the fuel relative to power deposited elsewhere in the core. The effect of these neutronic changes on the thermal outcome from MOOSE can be seen in Figure




Figure 4.15: MRaD displacement in centimeters in alternative case with bottom plane constrained. Results of this constraint showed no neutronic or thermal difference from a case with no axial constraint.

Table 4.2: Global neutronic results and differences from non-expansion MOOSE run for MRaD core. "AR" is Absorption Rate, "FF" is fission factor. Absorption rates are in units of "per source particle".

Quantity	Result Value	Difference	Pct. Dif. (%)
Eigenvalue	1.01685 ± 0.00035	0.00315 ± 0.00049	-
Leakage	0.05798 ± 0.0001	-0.00312 ± 0.00013	-5.11 ± 0.22
Thermal FF	$0.79814 \pm 1.2e-4$	$0.00121 \pm 1.7e-4$	0.15 ± 0.02
Fuel AR	$0.65359 \pm 2.2e-4$	$-0.00103 \pm 3.0e-4$	0.16 ± 0.05
Heat Pipe AR	$0.10609 \pm 7.0e-5$	$0.00009 \pm 9.0e-5$	0.08 ± 0.09
Moderator AR	$0.04286 \pm 2.0e-5$	$-0.00048 \pm 2.8e-5$	1.12 ± 0.07
Reflector AR	$0.14111 \pm 1.6e-4$	$0.00024 \pm 2.2e-4$	0.17 ± 0.16





Figure 4.16: Cell percent volume change due to thermal expansion.

4.18, a slice showing cell-average temperature change cause by expansion. These values are mostly without a clear pattern, though perhaps there is a slight increase in temperature towards the center of the core and a decrease in the outer assemblies. It is more likely, in my opinion, that this temperature difference is influenced by stochastic noise.



(a) Change in pin power after thermal expansion.

(b) Standard deviation of pin powers after thermal expansion.

Figure 4.17: Percent change in pin powers of MRaD core after thermal expansion, including standard deviation of changes.

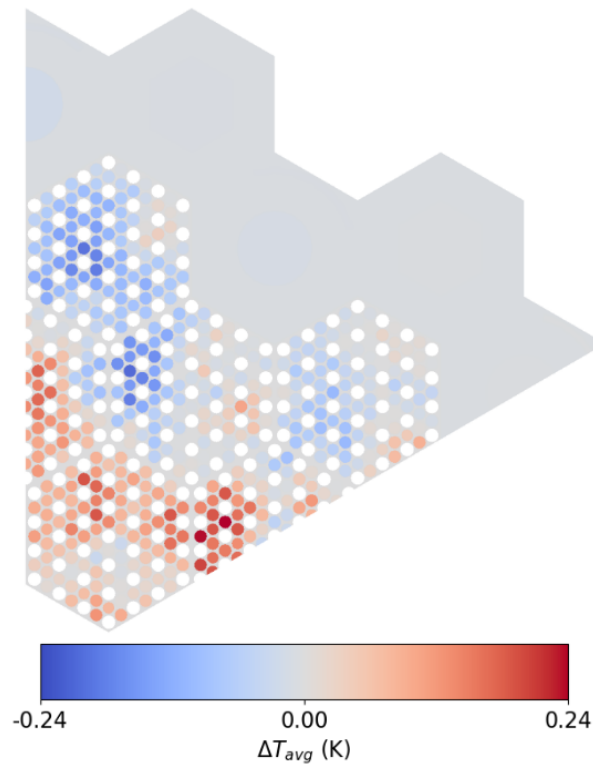


Figure 4.18: Cell-averaged temperature change after thermal expansion of MRaD geometry.

Heat Pipe Failure

Heat pipe failure causes a highly localized impact on the reactor core, with a single heat pipe failure not being overly impactful to the core. Instead, this heat pipe failure study looked at failing four heat pipes in the center hottest portion of the core, to attempt to simulate a severe failure scenario. The selected heat pipes failed can be seen in Figure 4.19, along with the resulting mid-plane temperature profile.

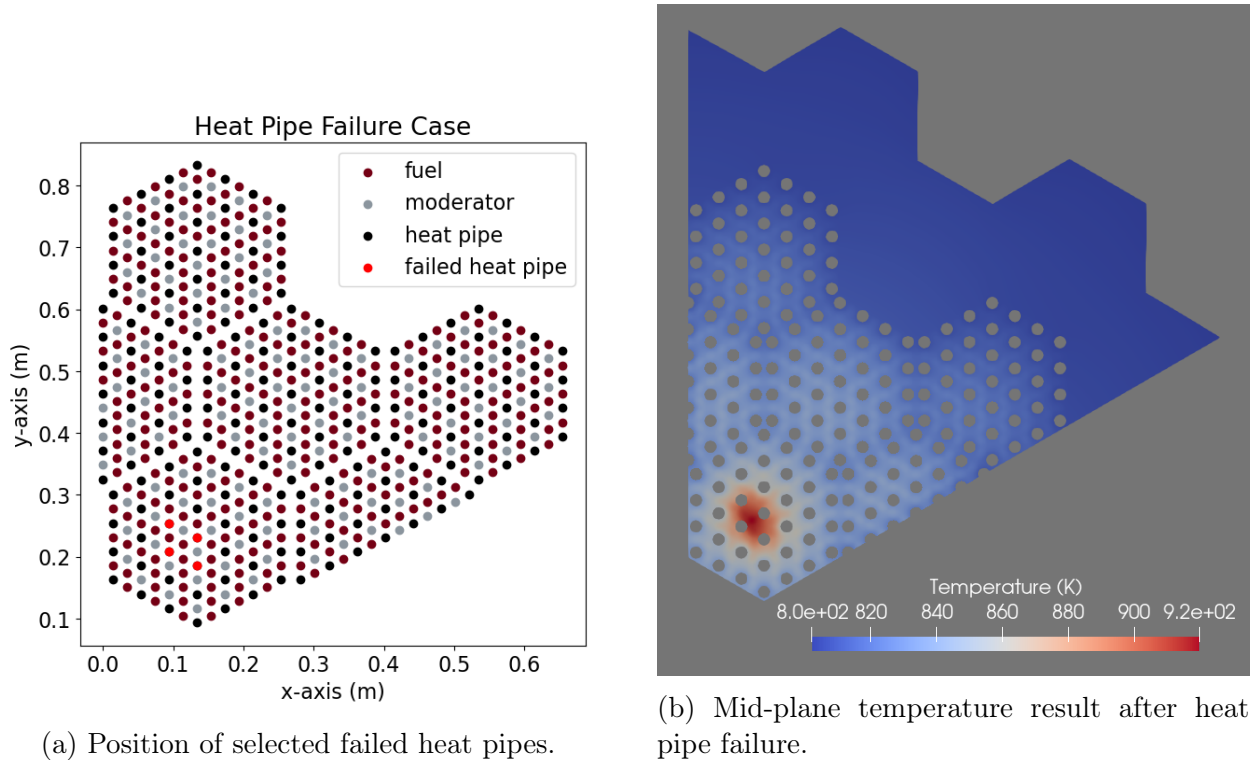


Figure 4.19: Selected heat pipe failure locations and the resulting mid-plane temperature effect on the MRaD core.

Even with four failed heat pipes, the surrounding heat pipes keep the temperature change relegated to such a small area in the core that displacement due to the failure is incredibly minor. This can be seen in Figure 4.20, where the change to expansion seen in the mesh due to the heat pipe failures is less than a millimeter. This results in the difference between including expansion and not including expansion being the same results seen from the non-heat pipe failure case.


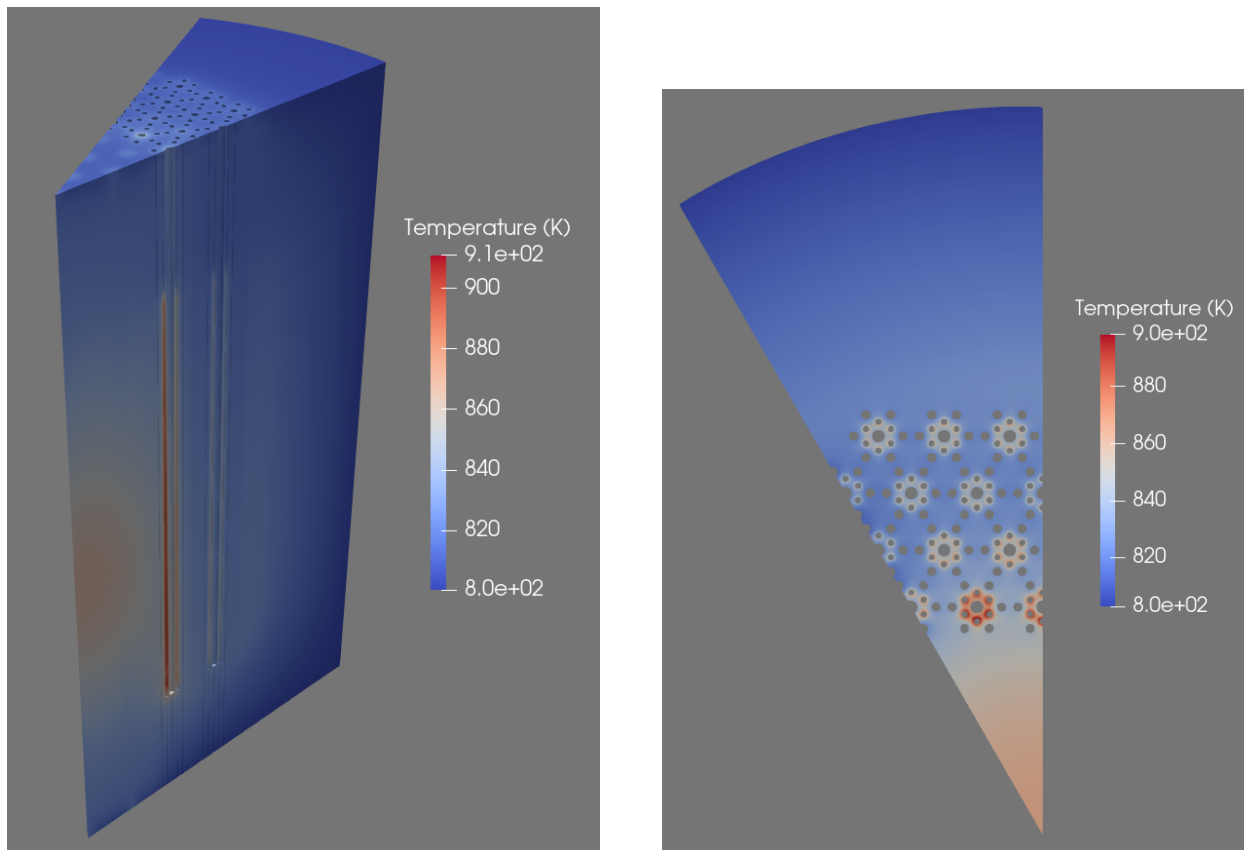


Figure 4.20: Comparison of mesh displacement between the heat pipe failure case and the no-failure case for the MRaD geometry. Heat pipe temperature impact is constrained by surrounding functioning heat pipes and the overall displacement effect is minor. Displacement is relative to the axial mid-plane.

4.3.2 SR Core

The SR core thermal profile post-convergence can be seen in Figure 4.21, and the displacement profile of the core can be seen in Figure 4.22. The cladding of the fuel pins causes higher fuel temperatures than seen in the MRaD geometry, though both cores see a similar axial displacement on the order of 1 centimeter. Figure 4.23 shows the percent change




(a) Isometric view of SR Core temperatures after final coupled iteration.


(b) Axial mid-plane slice view of SR Core temperatures after final coupled iteration.

Figure 4.21: SR temperatures after final coupled thermal expansion iteration.

in volume that each cell experienced due to thermal expansion. The fuel experienced the largest expansion in the SR geometry by a large margin, on average experiencing a nearly 4.5% increase in volume. Next highest was the radial reflector. The effect of this thermal expansion on neutronic statistics can be seen in Table 4.3. Similarly to the MRaD core, thermal expansion increased eigenvalue in the core and decreased leakage. While the leakage decrease was less significant than in the case of the MRaD core, the eigenvalue increase was larger, at 407 pcm. This is due to the increased fuel volumes previously mentioned. Overall, the increased fuel volumes and decreased leakage led to a slight increase in fission rate in the fuel and decreased absorption in other materials. The increased fuel volumes result in very minor changes to pin powers, seen in Figure 4.24. There is no distinguishable pattern of change for the pin powers, and the average pin power change is below the uncertainty. When



(a) SR displacement in centimeters after iteration.



(b) Mid-plane slice of SR showing displacement in centimeters after iteration.

Figure 4.22: SR displacement (cm) after final coupled thermal expansion iteration. Displacement is relative to the axial mid-plane.

comparing energy deposition tallies for the non-expanded case versus the expanded case, the fuel saw a 0.5% increase in energy deposition. When parsing the cell-average temperatures of the SR core in Figure 4.25, it seems that there may be a net-increase in temperature of the fuel due to this increased energy deposition. This increase in temperature is small, however, and may be influenced by stochastic uncertainty.





Figure 4.23: SR Cell percent volume change due to thermal expansion.

Table 4.3: Global neutronic results and differences from non-expansion MOOSE run for SR core. "AR" is Absorption Rate, "FF" is fission factor. Absorption rates are in units of "per source particle".

Quantity	Result Value	Difference	Pct. Dif. (%)
Eigenvalue	1.00702 ± 0.00034	0.00407 ± 0.00047	-
Leakage	0.09964 ± 0.00014	-0.00281 ± 0.00021	-2.74 ± 0.20
Thermal FF	$0.76651 \pm 1.2e-4$	$-0.00367 \pm 2.0e-4$	-0.48 ± 0.03
Fuel AR	$0.73425 \pm 2.5e-4$	$0.00510 \pm 4.0e-4$	0.69 ± 0.05
Heat Pipe AR	$0.07840 \pm 6.0e-5$	$-0.00240 \pm 8.0e-5$	-2.97 ± 0.10
Moderator AR	$0.02388 \pm 1.2e-5$	$-0.00017 \pm 1.7e-5$	-0.69 ± 0.07
Reflector AR	$0.06500 \pm 8.0e-5$	$-0.00005 \pm 1.1e-4$	-0.08 ± 0.17



(a) Change in pin power after thermal expansion.

(b) Standard deviation of pin power changes after thermal expansion.

Figure 4.24: Percent change in pin powers of SR core after thermal expansion, including standard deviation of changes.




Figure 4.25: Cell-averaged temperature changes due to thermal expansion of SR geometry.

Heat Pipe Failure

Similarly to the MRaD case, three heat pipes were chosen to be simulated failures in the SR core, specifically the heat pipes located in the highest temperature region. The selected cores and temperature outcome of these failures can be seen in Figure 4.26. Again, the temperature impact of these heat pipe failures are highly concentrated around the failed heat pipes themselves, with the majority of the core remaining at comparable temperatures to before failure. Because the thermal change is so limited, the difference in thermal expansion is also limited, which results in Figure 4.27. The difference in displacement between the heat pipe failure case and the non-failure case is far less than would register as a neutronic response.

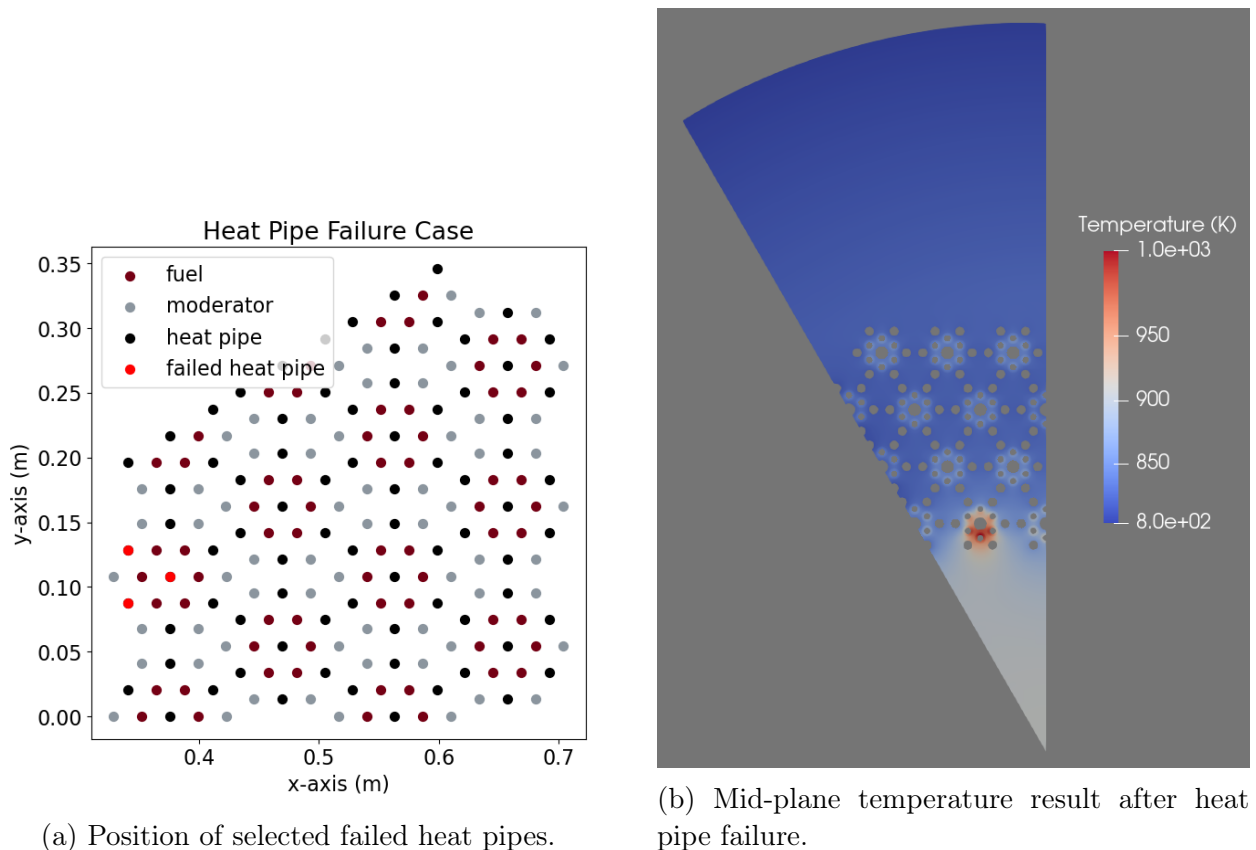


Figure 4.26: Selected heat pipe failure locations and the resulting mid-plane temperature effect on the SR core.




Figure 4.27: Comparison of mesh displacement between the heat pipe failure case and the non-failure case for the SR geometry. Heat pipe temperature impact is constrained by the surrounding functioning heat pipes and the overall displacement is negligible.

4.4 Conclusion

The results of this chapter of work show the process and output of coupling thermal expansion with neutronic-thermal coupling via OpenMC and MOOSE. A focus is the neutronic and thermal impact of the inclusion of thermal expansion-based deformation, highlighting what changes in results occur because of the expanded geometry. Also included are studies of the impact of heat pipe failure on thermal expansion.

For both the MRaD and SR cores studied, the inclusion of thermal expansion resulted in a decrease in leakage in the system with an accompanying increase to eigenvalue, on the order of 300 to 400 pcm. The downstream effect of this depended on the core design, and in particular depended on what features of the core experienced the largest thermal expansion.

For the MRaD core, the moderator rods expanded the most, followed by the reflector. This caused an increase in moderation of the core, slightly softening the flux spectrum. For the SR core, the fuel pins by far thermally expanded the most, which led to an increase in fission rate in the core and a minor 0.5% increase in power located in the fuel. For both cores, these changes did not noticeably affect pin power distributions nor the thermal profile of the core.

Both cores also aligned in their results from the heat pipe failure case. The thermal impact of heat pipe failures were highly relegated to small sections of the cores, causing net increases in thermal expansion to be below margins that one would normally care to measure. This resulted in no perceivable difference in thermal expansion-based effects between the failure case and the non-failure case.

For future work on this topic, there are plenty of areas of improvement for the methods used. Probably the largest improvement that could be made is refinement of the meshes used to represent the reactor cores. Because both the MOOSE and OpenMC simulation live on these mesh geometries, a highly refined mesh with axial segmentation of features such as fuel pins would provide a more accurate result to analyze. Accomplishing this while avoiding simulation failure due to the hurdles described in Section 4.2.2 would likely be challenging. Additionally, the mechanical side of the coupled simulation could be significantly improved by more accurately representing non-isometric expansion for materials such as graphite, higher accuracy modeling of fuel pins to analyze flowering, correctly representing radial expansion, and by including more detailed stress calculations. Finally, a correct constraintment of the sides of the sliced geometries would result in more accurate radial expansion representation, as explained in Section 4.2.2. This might reveal further interesting results.

Chapter 5

Hydrogen Diffusion

5.1 Background

The second of the two main topics presented in this thesis is the inclusion of hydrogen diffusion in reactor neutronic-thermal simulation, and its resulting impact on steady-state core behavior. Also of study are the effects of heat pipe failure on this feedback system.

As a general summary, some reactor designs call for the use of metal hydride moderators, as the increased hydrogen density provides higher moderating capabilities that a smaller reactor may need. These materials are known to feature hydrogen migration at the temperature ranges prescribed by reactor operation. Movement of the hydrogen within the moderator will invariably cause a resulting response in neutron flux and all other downstream responses.

This section of work details the efforts to calculate that spatial redistribution of the hydrogen in the moderator, and then calculate the impact of this redistribution on neutronic behavior and the resulting spatial power and thermal profile.

5.1.1 Why Metal-Hydride Moderators

When designing a thermal reactor, a dominating need across all reactor designs is the need for moderation. For example, all three types of reactors analyzed in this portion of the thesis use uranium-235 as their fissioning isotope, which favors thermalized neutrons in order to undergo fission. The term "thermal" in these contexts represents neutrons below 10 electronvolts in energy. To slow neutrons down to these optimal energies, moderating material is used as a medium to encourage neutron scattering collisions and loss of energy through kinetic energy transfer.

Qualifying a material as a "good" moderator [84] requires that: 1) the energy-lost-per-collision be relatively large, 2) the probability of scattering be relatively large, and 3) the probability of scattering in the material far outweighs the probability of being absorbed. This first component is typically quantified for elastic scattering via the average logarithmic energy decrement, ξ , where $\xi = 1 + \frac{\alpha}{1-\alpha} \ln(\alpha)$ and $\alpha = \left(\frac{A-1}{A+1}\right)^2$ where A is the atomic mass number. Looking at the inputs, it is clear that lighter elements have a larger value of ξ . Kinematically, this makes sense, as the energy transfer from a projectile to a target is likely to be larger when the target matches the mass of the projectile. Think a baseball hitting

Table 5.1: Some common moderating materials and their moderating properties. ξ is the average logarithmic energy decrement, MSDP is the macroscopic slowing down power, MR is the moderating ratio. Values from [84]–[86]

Material	ξ	MSDP	MR
H ₂ O	0.920	1.35	71
D ₂ O	0.509	0.176	5670
Be	0.207	0.158	143
C	0.158	0.060	192
ZrH ₂	0.030	1.45	55
YH ₂	0.048	1.2	25

another baseball, versus a baseball hitting the earth.

The second component is typically quantified via a Macroscopic Slowing Down Power (MSDP), calculated by multiplying the logarithmic decrement by the macroscopic scattering cross-section ($\text{MSDP} = \xi \cdot \Sigma_s$). This value is meant to take into account the fact that energy is only lost if the particle actually collides with the target, so a higher chance of that happening implies a "better" moderator.

The final component of being a "good" moderator is typically quantified via the Moderating Ratio (MR), calculated as $\text{MR} = \xi \cdot \Sigma_s / \Sigma_a$, where Σ_s and Σ_a are the macroscopic scattering and absorption cross-sections, respectively. This ratio weights the energy decrement value to account for materials that have a high energy decrease per collision but also are likely to absorb neutrons, like lithium-6. Table 5.1 lists some common moderating materials used in reactors, along with their respective values for ξ , MSDP, and MR.

What are not reflected in Table 5.1, however, are how the specific requirements of the reactor design can weigh each factor. Because the microreactor geometries included in this work are severely limited in size, materials with high MR but low MSDP must be excluded. These materials preserve neutron economy at the cost of requiring larger dimensions to allow for the additional collisions needed for neutron thermalization.

This filter narrows down material choice to either a metal hydride or H₂O. Water has strict thermal constraints, however. Its low boiling point (relative to liquid metals, for example) often becomes the dominant temperature limit for the system. Although reactor designs with supercritical water have been explored, the requirement of pressurizers and the risk of steam flashing make this moderator a poor choice for a microreactor that must be highly mobile. Zirconium hydride and yttrium hydride, in contrast, are suited for use in environments up to 600°C and 750°C, respectively [15], [87]. In order to produce power at the levels designed for, H₂O as a moderator is not an option for the microreactors studied in this work.

Left with ZrH_x or YH_x to choose from, the final choice is often based on a more mundane reason; Zirconium is cheaper to procure than yttrium [85]. Yttrium hydride features some advantages over zirconium hydride that will be mentioned in Section 5.1.3, but yttrium's relative scarcity as a rare earth element means a higher cost for reactor use. The economics of the reactors included in this work are not a topic of study, but it bears mentioning that microreactors already are expected to face a relatively high Levelized Cost of Electricity [6].

It is fair to assume, then, that using a moderating material that has significantly higher costs yet similar performance would be an unlikely choice. In addition, and of largest importance to this work, zirconium hydride’s long history of use means its material properties are better defined than the rarer yttrium hydride.

It is for these reasons that zirconium hydride will be the moderator of choice for all reactor designs studied in this work.

5.1.2 Zirconium Hydride

Zirconium is a material with a long history in the nuclear power industry. Zirconium was originally designated for use in naval nuclear reactors by Admiral Rickover, due to its corrosion-resistance and neutron transparency [88]. This would eventually lead to the Zircaloy claddings that are still in use today. In fact, the large body of work detailing the behavior of these cladding zirconium in presence of hydrogen is the source of a majority of the physics used in this work [89]. One of the first instances of zirconium hydride being proposed as a moderator comes from a 1960 boiling water reactor design [90]. Arguably its most famous occurrence comes from its use as the combination fuel-moderator UZrH featured in SNAP and TRIGA reactor designs [16].

The method of hydriding zirconium metal is not particularly complicated. First, highly pure zirconium metal must be acquired for the process. Zirconium as found in nature tends to contain hafnium, which has a large thermal neutron cross-section and therefore would negatively impact moderating performance [88]. This pure zirconium metal is then placed in a furnace, pulled to vacuum, flushed with a gas like argon, and then raised in temperature to roughly 900°C in order to degas. The temperature is then lowered to a specified fabrication temperature and hydrogen gas is introduced to the system. The quantity of hydrogen introduced will dictate the hydrogen concentration that the ZrH reaches. Once the concentration is constant throughout the material, the temperature is lowered, following the temperature-pressure isochor. These steps are explained in more detail in the work of Hu et al. [85]. Recent studies have looked at the viability of using Zircaloy as the initial material, due to its smaller grain size and therefore faster hydriding rate [91].

Because zirconium hydride is an alloy, its metamorphic behavior can be described with a phase diagram, as seen in Figure 5.1. Zircaloy and zirconium-based cladding materials are characterized as α -phase ZrH [69], given that their hydridization is a result of the surrounding environment (typically the H₂O coolant). When used in a moderating context, the zirconium hydride is typically in its δ -phase [15]. For the studies included in this thesis, all moderator material will begin with a default composition of ZrH_{1.67}.

The phase diagram for zirconium hydride is surprisingly controversial. A literature review on the topic shows distinctly different conclusions about what that phase diagram looks like, particularly where the $\delta \rightarrow \epsilon$ transition happens, the bounds of the narrow two-phase $\delta + \epsilon$ region, or if a stable γ -phase exists. [92]–[98]. Because representing these complex phase-change dynamics goes beyond the purposes of this work, an approximation is made that the moderator will remain in the δ -phase regardless of the changes in hydrogen concentration or temperature calculated. The work by Trainer [99] suggests that the phase of ZrH_{*x*} is less impactful in comparison to hydrogen concentration. The steady-state powers of all reactors analyzed are limited in order to constrain peak moderator temperatures to below 600°C, as




Figure 5.1: H-Zr phase diagram with isobars of equilibrium H_2 pressures indicated by colored lines. Figure retrieved from Olander et. al. [16]

a way of trying to guarantee the δ -phase assumption's validity. Heat pipe failure studies that feature moderator temperatures beyond this limit will contain the warning that real-life behavior may deviate from the results shown, as ϵ -phase ZrH has been shown to potentially have significantly higher rates of hydrogen diffusion [70], [100].

5.1.3 Diffusion of H in ZrH

Mentioned several times now throughout the previous sections, including the end of Section 5.1.1, the quality of zirconium hydride that is of interest to this thesis is the mobility of hydrogen within the zirconium crystal lattice. The δ -phase zirconium exists in a face-centered-cubic lattice, with hydrogen atoms residing in the tetrahedral interstitial sites between zirconium atoms [16]. This hydrogen mobility is what allows the hydriding process described in [85], and is a feature shared by other crystal lattices that have been proposed for hydride moderator use, such as titanium and yttrium [85], [101]. Figure 5.2 depicts a variety of these materials along with their hydrogen atom density versus temperature. Clearly depicted is how responsive to temperature these hydride materials are. Additionally, looking at yttrium's behavior, the increased retention of hydrogen up to temperatures beyond 1000°C are why YH is often proposed as a moderating material in place of ZrH.




Figure 5.2: Hydrogen atom density as a function of temperature for various metal hydrides in equilibrium with 1 atm of hydrogen gas. Figure retrieved from [102].

Simulation Equations

As previously noted, interest in the α -phase of zirconium hydride for cladding purposes has been a significant source of characterizing hydrogen migration in zirconium. A good example of this is Courty's work [89], modeling hydrogen behavior in Zircaloy-4, which is the source for the following equations that describes hydrogen migration in zirconium. The migration of hydrogen through these lattices is driven by three factors: 1) the hydrogen concentration itself, 2) the temperature gradient, and 3) any mechanical stress on the material.

Diffusion by concentration is most commonly referred to as Fick's Law, the equation of which is included in Equation 5.1. J_{Fick} is the diffusive flux of hydrogen, D is the diffusion constant, and C is the hydrogen concentration.

$$J_{Fick} = -D \cdot \nabla C \quad (5.1)$$

Diffusion by temperature gradient is referred to as the Soret Effect, or Thermodiffusion. This effect comes from the natural repulsion of particles from areas of higher thermal motion due to kinetic transfer, creating a "thermophobic" effect as particles concentrate in colder regions [103]. This effect is represented as an equation in Equation 5.2, where Q^* is the heat of transport, R is the ideal gas constant, and T is the temperature in Kelvin. This heat of transport is assumed to be a value of 25.1 kJ/mol, in agreement with [89]. The diffusion constant is calculated via the Arrhenius equation, expressed in Equation 5.3, where A_D is the pre-exponential factor, and Q_D is activation energy. For this work, a diffusion coefficient of $D = 1.53 \times 10^{-7} \cdot \exp(-14000/RT)$ cm²/s, defined in [104] and utilized in [105].

$$J_{Soret} = -\frac{DCQ^*}{RT^2} \cdot \nabla T \quad (5.2)$$

$$D = A_D \cdot \exp\left(-\frac{Q_D}{RT}\right) \quad (5.3)$$

Hydrogen diffusion via stress occurs when stress applied to the material leads to strain-induced hydrogen ordering phenomena. Work by Steuwer et. al. has shown transformation of the δ -hydride to a γ -hydride in the subset of grains which are suitably aligned with the loading direction [106]. Because the magnitude of this effect is less than that of Fick's Law and Soret Effect's influence, the impact of stress on hydrogen diffusion is typically neglected [107]–[109]. This is also the case for work described in this chapter of the thesis, leaving us with a final equation describing hydrogen diffusion in zirconium, Equation 5.4.

$$J = J_{Fick} + J_{Soret} = -D \cdot \nabla C - \frac{DCQ^*}{RT^2} \cdot \nabla T \quad (5.4)$$

Given that the hydrogen in the zirconium lattice is the true moderator for the reactor's neutrons, it is obvious why movement of hydrogen would be of concern. Movement of the moderator directly translates to a spatial response from the bulk neutron population. If the magnitude of that response is large enough, it could translate to a change in power shape during reactor operation. A back-of-the-envelope estimation of the magnitude of hydrogen migration isn't quite feasible in this case, as the feedback between hydrogen migration, neutronics, and temperatures make this a nonlinear coupled problem. Verifying the magnitude of both the hydrogen migration and neutronic-thermal response is the primary goal of this chapter of the thesis.

It is easy enough to calculate an estimated diffusion coefficient from Equation 5.3, done in Equation 5.5. Using this coefficient, we can perform a simple 1D diffusion time estimation, Equation 5.3, to estimate the time it would take for a hydrogen particle to migrate 1 cm, a hypothetical radius for a moderator rod. The result is around 260 days, not a particularly fast rate, but given that many microreactor designs are expected to operate for upwards of five years [110], a steady state should be expected to be reached during operation.

$$D = 1.53 \times 10^{-7} \cdot \exp\left(\frac{-14000}{8.3145 \cdot 870}\right) \approx 2.2 \times 10^{-8} \text{ cm}^2/\text{s} \quad (5.5)$$

$$t \approx \frac{x^2}{2D} = \frac{1^2}{2 \cdot 2.2 \times 10^{-8}} \approx 263 \text{ days} \quad (5.6)$$

Hydrogen Loss

While redistribution of hydrogen in the moderator is of interest to this work for all the reasons listed earlier, leakage of hydrogen out of the material has not been mentioned so far. This is potentially much more impactful of a phenomenon than simple redistribution of hydrogen, as H_2 is highly combustible as learned from past reactor failures [111]. Even without this threat, loss of hydrogen from the moderator will reduce overall moderation,

which in turn reduces the thermal fission rate in the core. This hydrogen loss was the main performance constraint for SNAP reactor operations [16], [112].

Characterizing the desorption of hydrogen from the zirconium lattice is a complicated topic, with loss rates strongly dependent upon the surface state of the ZrH. The material with clean surface conditions saw rapid hydrogen escape at 600°C temperature ranges, while surfaces with an oxide layer saw a slower loss rate [113], [114]. Due to these factors, cladding is required for zirconium hydride moderators [15]. The specific cladding used will likely vary based on specific design, but Alumina has been suggested as one material option.

For the work in this thesis, all moderators are presumed to have a "perfect cladding", meaning no leakage of hydrogen occurs.

5.1.4 Functional Expansions

One distinct challenge in coupling Monte Carlo neutronics to a finite element-based solver is the difference in the physical domain representation. Monte Carlo neutronics, like OpenMC, generally use Constructive Solid Geometry (CSG), where surfaces are analytically defined with expressions and volumes are created through boolean operation with those surfaces. Tallies are then accumulated during simulation by volume integration of flux times a score multiplier[12], typically with some filter applied. This essentially ends up being a discrete histogram representation of the quantity of interest. Finite element solvers, like MOOSE, define quantities over the element volumes, calculate quantities locally, often at each node or each quadrature point, and volume averages are a tertiary operation that functions by integrating over point-wise data [14].

This implies that traditional coupling of the two simulations would mean passing volume-averaged data from one to inform the behavior of the other. Refinement of the quality of information being passed requires refining the problem geometry itself, discretizing the CSG model, making the already costly neutronic solve even costlier [30].

The use of functional expansions to represent spatial heterogeneity has been suggested as a solution to this challenge [115], [116]. By calculating expansion coefficients that fit orthonormal basis sets of functions to an input distribution, a series of easy-to-calculate, continuous functions can replace the often complicated fine-grid data.

Representing otherwise complicated information with a series of simple expressions (often with mathematically useful traits, like orthogonality) is one of the core lessons of calculus. Nearly every student in calculus will learn that functions with difficult to describe local limit behaviors can be expressed instead as an n^{th} order Taylor series that is readily differentiable [117]. Via Equation 5.7, $f(x)$ is the hypothetical troubling function, and the right-hand side is a potentially simpler representation of this function about $x = a$.

$$f(x) = \sum_{n=0}^{\infty} \frac{f^n(a)}{n!} = f(a) + f'(a)(x - a) + \frac{f''(a)}{2!}(x - a)^2 + \dots \quad (5.7)$$

The analogous situation for the work of this thesis is that we have randomly sampled distribution of discrete particle reactions in a volume (for an analog Monte Carlo simulation) as our $f(x)$, and the right-hand side of Equation 5.7 is an arbitrary series of basis functions. This sampled data has no analytical or closed form, which make the basis functions a handy

replacement. Equation 5.8 represents this, where a_n is the n^{th} expansion coefficient, k_n is the normalization factor, and ψ_n is the n^{th} function in the basis set chosen.

$$f(x) = \sum_{n=0}^{\infty} a_n k_n \psi_n(x) \quad (5.8)$$

In the work described in this chapter, the goal is to use these functional expansions to represent the spatial distribution of hydrogen in the moderator rods and the volumetric heating rates in both moderator and fuel.

History of Functional Expansions

An early example of using a series expansion to represent the shape of a statistical distribution is Rice's work in 1945, expressing random noise of an electrical signal with a Fourier series expansion [118]. The application of these expansion-based representations to nuclear purposes is first published in 1975 with a spherical harmonics expansion representing the angular distribution of x-ray photo-emission [119]. The first application for Monte Carlo uses is in 1984 when Noel and Wio used Legendre polynomials to describe spatial and angular distributions of neutron flux in a 1D shielding application [120].

For the next 20 years there were various publications that utilize functional expansions in similar ways, but in a mostly isolated manner without truly generalizing the method or providing formal proofs of statistical claims [115]. This changes with the works of Brown, Martin, and Griesheimer in the period of 2003-2005 [115], [121]–[123], who in-detail describe the methods of functional expansion tallying in Monte Carlo along with introducing a method of sampling varying material properties during transport that will be highlighted in Section 5.2.3.

These works directly enabled implementation of functional expansion tallying in Monte Carlo codes OpenMC [30] and Serpent2 [124], [125], which makes the work described in this chapter possible.

Choice of Basis Set

The historical examples of functional expansion utilization presented previously each include the use of different basis sets that suit their purposes. The choice of function series used for this process strongly impacts the accuracy and efficiency with which the true spatial information can be reproduced. The series used must be selected so that their domain of definition matches the domain of interest.

In the work of this chapter, the domains represented with series expansion are finite cylindrical domains, specifically representing fuel pins and moderator rods. Because no single basis set fits for this geometry, to the best of my knowledge, the dimensions are discretized into an axial basis set and an azimuthal/radial basis set. A common choice, and the choice that this work uses, is rebuilding the cylindrical distributions using Zernike polynomials for the radial and azimuthal dimensions, and Legendre polynomials for the axial dimension [126].

Legendre polynomials originate from Legendre's 1782 work as coefficients for the expansion of the Newtonian potential, which applies to gravitational potential or Coulomb




Figure 5.3: Legendre polynomials plotted up to 5th order.

potential [127]. These polynomials have been used widely since then for a varying spread of purposes, and there exist multiple definitions for the generation these polynomials L_n . One common expression is Rodrigues' formula, Equation 5.9. A plot of $L_n(x)$ up to 5th order is included in Figure 5.3.

$$L_n(x) = \frac{1}{2^n n!} \frac{d^n}{dx^n} (x^2 - 1)^n \quad (5.9)$$

The Legendre polynomial series is a well-suited fit for representing axial hydrogen concentration and heating profile because:

- Legendre polynomials are so ubiquitous there exist an implementation of the series in most software and packages. This includes OpenMC, MOOSE, and Python via Numpy. This means little-to-no required work on our side to utilize Legendre polynomials for functional expansion.
- Legendre polynomials are well-fit to work with cosine inputs, which is important as the heating rates for core fuels and moderators are expected to have a cosine shape due to axial leakage.

Zernike polynomials, named after Nobel laureate Frits Zernike, are a series of polynomials orthogonal over the unit disk, and were originally intended for characterizing aberrations in optical systems [126]. Similarly to Legendre polynomials, their use has appeared in topics far beyond optics since. Zernike polynomials can be calculated via a combined radial and azimuthal expression, and a visual of $Z_n^m(r, \theta)$ can be found in Figure 5.4.

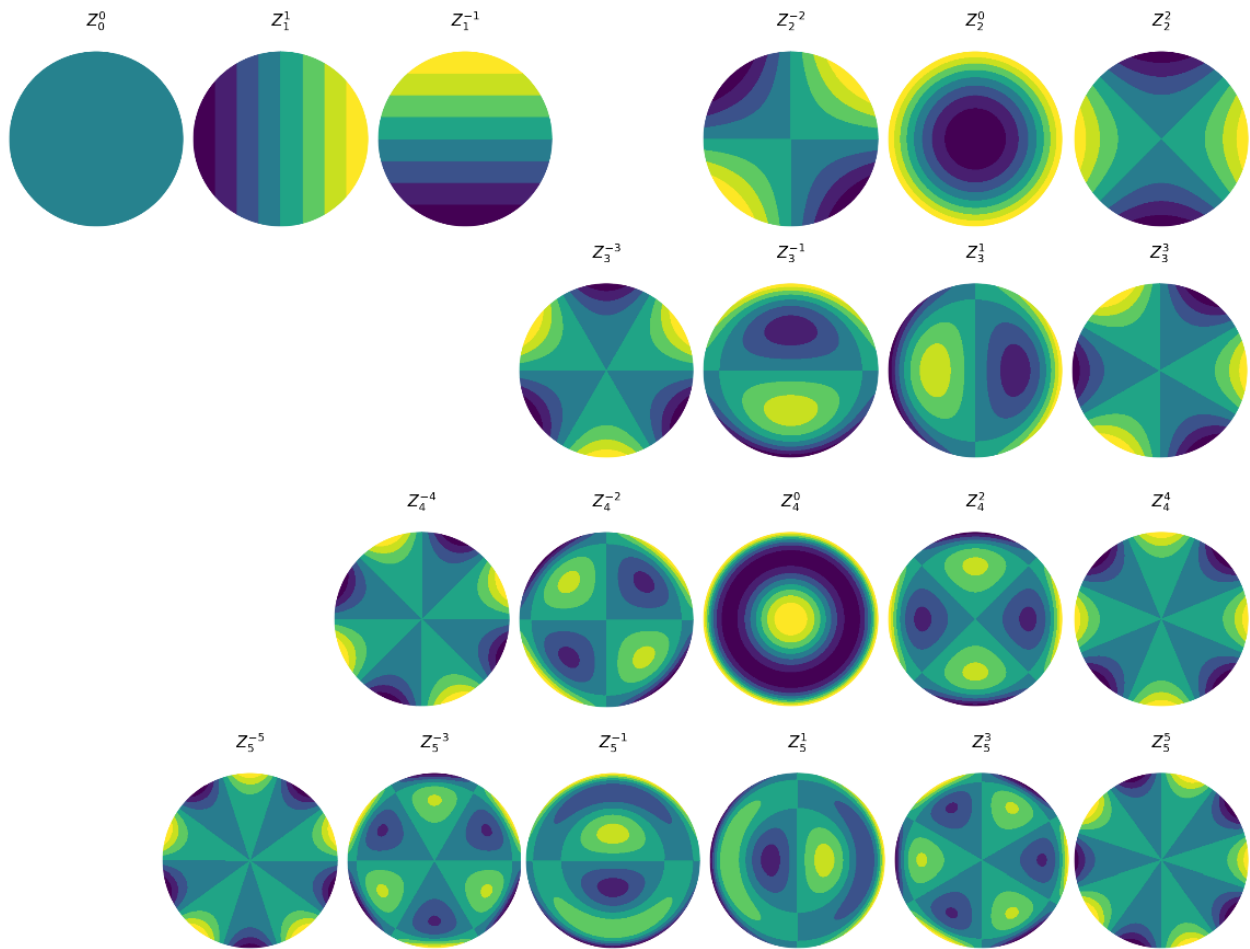


Figure 5.4: Zernike polynomials plotted on a unit disk up to 5th order.

One benefit of using Zernike polynomials is that the individual radial and azimuthal components of each n^{th} basis can be factorized, allowing efficient computational implementation. For example, rather than storing each function uniquely, equations for even Zernike polynomials (Eq. 5.10) and odd Zernike polynomials (Eq. 5.11) are retrievable by solving the radial equation (Eq. 5.12) and multiplying by a sine or cosine.

$$Z_n^m(r, \theta) = R_n^m(r) \cos(m \cdot \theta) \quad (5.10)$$

$$Z_n^m(r, \theta) = R_n^m(r) \sin(m \cdot \theta) \quad (5.11)$$

$$R_n^m(r) = \sum_{k=0}^{\frac{n-m}{2}} \frac{(-1)^k (n-k)!}{k! \left(\frac{n+m}{2} - k\right)! \left(\frac{n-m}{2} - k\right)!} r^{n-2k} \quad (5.12)$$

A disadvantage of using Zernike polynomials is that the unequal distribution of nodal lines over the unit disk introduces ring effects near the perimeter of the unit disk [128]. Additionally, arbitrarily increasing the order of Zernike polynomials used can result in excessive errors if the sampling is not refined enough to have confidence in the accuracy of calculated higher-order moments [115].

Combining the axial and radial/azimuthal functional representations, a necessary step to truly represent the cylindrical domain, means convoluting the two basis sets into a 3D one [126]. This cross-multiplication means if both expansions are of 5th order, the cylindrical basis set will have 126 terms ($n_{cyl} = n_l \times n_z = 6 \times 21$) and 126 expansion coefficients. Tallying the highest order expansion coefficients (such as the combined 5th order Zernike 5th order Legendre terms) accurately would require immense particle counts for OpenMC and numerous quadrature points for MOOSE. Instead, this work truncates higher order cross-terms beyond a *combined* 5th order to minimize error. Additionally, for all work in this chapter, both 5th order Legendre and Zernike polynomials are used.

Tallying and Enforcement

Tallying the moments of the functional expansion across the cylindrical geometry is straightforward for both OpenMC and MOOSE and exists in an easily implementable form for users. The term "tallying" here means calculating expansion coefficients. The expansion coefficients can be used later with their corresponding polynomial basis set to rebuild a fit of the original data that was sampled against. In a Monte Carlo simulation, these expansion coefficients are estimated by

$$\hat{a}_n = \frac{1}{N} \sum_{i=1}^N \psi_n(z_i, r_i, \theta_i) \rho(z_i, r_i, \theta_i) \quad (5.13)$$

where \hat{a}_n is the estimated n^{th} expansion coefficient, N is the number of particles tallied, ψ_n is the n^{th} term in the basis set, and ρ is the weighting function [129]. This ρ would be the flux multiplied by a scoring value for the case of tallying in an OpenMC simulation, i represents the individual particle, and ψ in this work is the previously-mentioned cylindrical basis set.

Reconstructing the original information via the expansion coefficients means sampling the cylindrical basis set weighted by the previously-calculated moments at a requested point and using the result as an input. We have two different use cases for these expansions in this work, for each code. For MOOSE, that input is the volumetric heating rate for the unique fuel pin or moderator rod. For OpenMC, that input is the number density of hydrogen in the moderator rod. Explicitly stated, sampling is accomplished via

$$V(z, r, \theta) = \sum_{i=0}^I \sum_{n=0}^N \sum_{m=-n}^M \hat{a}_{l,n}^m L_i(z) Z_n^m(r, \theta) \quad (5.14)$$

where V is the value rebuilt from the expansion representation, i , n , and m are the indexes describing order for Legendre and Zernike polynomials, L_i is the Legendre polynomial of order i , and Z_n^m is the Zernike polynomial of orders (n, m) .

MOOSE is capable of taking expansion coefficients as inputs to define a `FunctionSeries` Function, that then can be enforced multiple different ways [130]. For this work, where these expansion coefficients rebuild volumetric heating shape distributions, this Function is used as a parameter of each respective `HeatSource` Kernel. MOOSE has a `CylindricalDuo` `FunctionSeries` definition, using the same convoluted basis set described previously, also used in Wendt [126]. This allows users to enforce and tally the needed expansion coefficients without any source code modification.

In the case of OpenMC, only the ability to tally expansion coefficients is included with publicly-accessible source code. What is tallied is an un-normalized set of expansion coefficients via Equation 5.13. In order to enforce some behavior using expansion coefficients, for this work's purpose being hydrogen number density, an additional complication is required, as seen in Section 5.2.3.

Orthonormalization

One last important caveat about functional expansions' use in this work, all polynomials and therefore expansion coefficients must be orthonormalized after tallying. Many basis sets are inherently orthogonal or have defined methods of orthogonalization, as orthogonality allows for recurrence relations and standardization. Recurrence relations allow computational optimizations, a well-defined need for any multiphysics simulations like the one described in this chapter, and standardization allows equivalent mapping between the computed domain and the "standard" domain the functional expansion exists on. Legendre and Zernike polynomials are readily orthogonalized.

Orthonormality can be seen as a sort of subset of orthogonality, where the product of non-equal basis set functions ψ_i and ψ_j equal the Kronecker function δ_{ij} rather than some coefficient. Orthonormalization is required in order to correctly equate the original data set sampled with the rebuilt data set calculated via expansion coefficients. Because each polynomial basis set is orthogonal to begin with, orthonormalization is accomplished by normalizing the expansion coefficients. This requires multiplying each term by a orthonormalization constant c_{ψ_i} unique to its basis function and order. The explanation here relies heavily on the work of Wendt [126], whose explanation is the best I have seen on this topic.

The orthonormalization constants for Legendre polynomials are

$$b_n = \frac{2n+1}{2} a_n \quad (5.15)$$

while the constants for Zernike polynomials, taking into account both the radial and azimuthal component, are

$$b_n^m = \begin{cases} \frac{n+1}{\pi} a_n^m & \text{if } m = 0 \\ \frac{2(n+1)}{\pi} a_n^m & \text{otherwise} \end{cases} \quad (5.16)$$

With orthonormalization, the flow for how (generalized) tallying and enforcement occurs is:

- Expansion coefficients a_n are calculated from a spatial distribution $F(\vec{r})$ in domain Γ , \vec{r} being the location in phase space (z,r,θ) .

$$a_n = \int_{\Gamma} \psi_n(\vec{r}) \rho_{\psi}(\vec{r}) F(\vec{r}) \quad (5.17)$$

- Apply normalization constant c_{ψ_n} to the expansion coefficient a_n .

$$b_n = c_{\psi_n} \cdot a_n \quad (5.18)$$

- After exchanging coefficients with the next solve, the equivalent distribution $f(\vec{r})$ is rebuilt and sampled as an input.

$$f(\vec{r}) = \sum_{n=0}^N b_n \psi_n(\vec{r}) \approx F(\vec{r}) \quad (5.19)$$

Noted before this orthonormalization subsection, OpenMC's functional expansion tallies are un-normalized, while MOOSE can orthonormalize coefficients before output.

5.1.5 Literature Review

Efforts to understand hydrogen migration's impact on zirconium hydride-moderated core behavior are only a recent trend in the reactor physics community. The majority of the reason for this is because of the computational cost of simulating this sort of coupling. Without the advancements in computer architecture and software like OpenMC and MOOSE that focus on paralellizability, performing coupled 3D heterogeneous neutronic-thermal simulation would not be possible. On top of that, microreactor geometries have become a topic of wider interest. With all the advancements today, this would still likely be too-costly of a study to perform conventionally on a gigawatt-level reactor geometry, and would require the introduction of additional methods.

With a newly-focused interest in microreactors, along with MOOSE being developed for multiphysics coupling, studies on the impact of hydrogen migration on reactor operation are beginning to appear in the field. I will briefly mention a couple of these in order to point out similarities and differences in methodology and results.




Figure 5.5: Axial stoichiometric ratio of hydrogen in YH_x sampled in varying radial regions. Innermost radial regions experiences highest temperature gradients. Figure retrieved from [131].

The work of Mehta et. al. analyzing the impact of hydrogen diffusion on a hypothetical U-Mo fuel, $YH_{1.8}$ moderator heat pipe reactor is, I believe, the first of its kind for this topic [131]. Hydrogen migration and thermal conduction are solved via Abaqus [132], and neutronics are solved via MCNP [133]. Results showed migration of hydrogen away from the central highest-temperature areas, seen in Figure 5.5, and a resultant central axial dampening of power and temperature. The geometry studied isn't representative of a true core, but the trends observed would be expected to be reproduced in the cores studied for this work.

It should be noted that the scheme for performing thermal conduction and hydrogen diffusion in MOOSE described in Section 5.2.2 has been used before in the same context of hydrogen migration. In fact, two of the geometries studied in this thesis were created for purposes similar to the contents of this chapter. Firstly, the "MRaD" core geometry, sourced from the Idaho National Laboratory National Reactor Innovation Center Virtual Test Bed (VTB) [134]. Stauff et. al. created and utilized this geometry in coupled simulations with MOOSE, Griffin, and SWIFT [32], [33]. SWIFT is a Los Alamos National Laboratories code, standing for Stoichiometry With Internal Fluctuating Temperature, was used to solve the hydrogen redistribution in the geometry's moderator [135]. Griffin is a deterministic reactor physics code used for calculating neutronics and power distribution [27]. Both codes were designed to function inside of the MOOSE framework and are used as such. The coupled model with SWIFT is not available on the VTB.

The general results of the hydrogen-focused portion of Stauff's work showed diffusion of hydrogen to the axial extremes of the moderator rods and a minor impact to eigenvalue and power shape, specifically a minor depression of power in the axial-center of the fuel, seen in Figure 5.6. These results are consistent with this thesis' results (coming in Section 5.4.1) for the MRaD geometry, though there is not a one-to-one comparison available due to the




Figure 5.6: Griffin SN(3,3) calculated axial power distributions with and without YH-hydrogen diffusion via SWIFT. Figure retrieved from [33].

differences in Stauff’s work and this thesis. Particularly, the change in moderator material to ZrH, the lack of hydrogen leakage, and the use of OpenMC with functional expansion applications make the results not directly comparable. They are similar-enough to note agreement in magnitude of eigenvalue change and axial power dampening.

A similar study is performed by Terlizzi et. al. [37], utilizing the Simba geometry. As with the MRaD case, this study’s alterations of moderator material choice and differences in neutronics solver make a direct comparison tenuous, but some comparisons can certainly be made. Again, a shift in hydrogen concentration to the top and bottom of the rods can be seen, with a magnitude of roughly 4% difference between maximum and minimum concentration in the rod. This difference can be seen in Figure 5.7. This hydrogen shift’s impact to eigenvalue is reported to be roughly 14 pcm, a value not far off from the results of 5.4.2.

To the best of my knowledge, this work is the first to solve this problem in a general approach that can continuously capture the migration in axial and radial directions, and to do so on multiple reactor geometries. The results shown later in this chapter align well with the previous trend of results shown, and include detail and analysis of the impact of hydrogen-migration that I have not seen on the topic yet.




Figure 5.7: Axial hydrogen stoichiometric ratio and radially-averaged moderator temperature, sourced from [37].




Figure 5.8: Simplified hydrogen diffusion coupled simulation scheme. CAD here referred to CAD-based mesh geometries, as MOOSE interacts with meshes rather than strictly CAD geometries.

5.2 Methods

5.2.1 Coupling Scheme

Similarly to the coupling scheme for thermal expansion, hydrogen diffusion coupling is done via an operator-split with Picard fixed point iterations. No information is shared in-memory between the neutronics solve or the finite element solve. Instead, results are processed after each respective run and necessary information is included in the input file for the next simulation's run. The information transferred is dependent on the solve, with the OpenMC solve providing volumetric heating rates (in the form of functional expansion moments for fuel and moderator) and the MOOSE solve providing hydrogen concentrations (also in the form of functional expansion moments). A simplified form of this coupling can be seen in Figure 5.8.

Note that temperature is not exchanged between MOOSE and OpenMC. This is on purpose, intended to help isolate the impact that hydrogen migration has on all results. In the case that hydrogen migration caused temperature changes large enough to noticeably broaden cross-sections, this would be an oversight. As will be seen in the results in Section 5.4, this is not a concern.

A more detailed form of the coupling is viewable in Figure 5.9, and the following subsections explain some of these steps. This includes the pre-iteration processes, and the neutronic and finite element solve processes that are a part of the main iteration loop.

Pre-Iteration Steps

To begin with, mesh generation is done via MOOSE itself, utilizing the Reactor module [40]. This module is designed for standardized mesh generation of Cartesian and hexagonal lattice core geometries. I have found it to be significantly more useful than the traditional method of building a CAD geometry and meshing afterwards. It is particularly useful to have a mesh generation tool that allows the user to explicitly define the number of mesh rings and azimuthal sectors. The mesh generation input for the MRaD geometry and Simba geometry

Hydrogen Diffusion Coupling Scheme

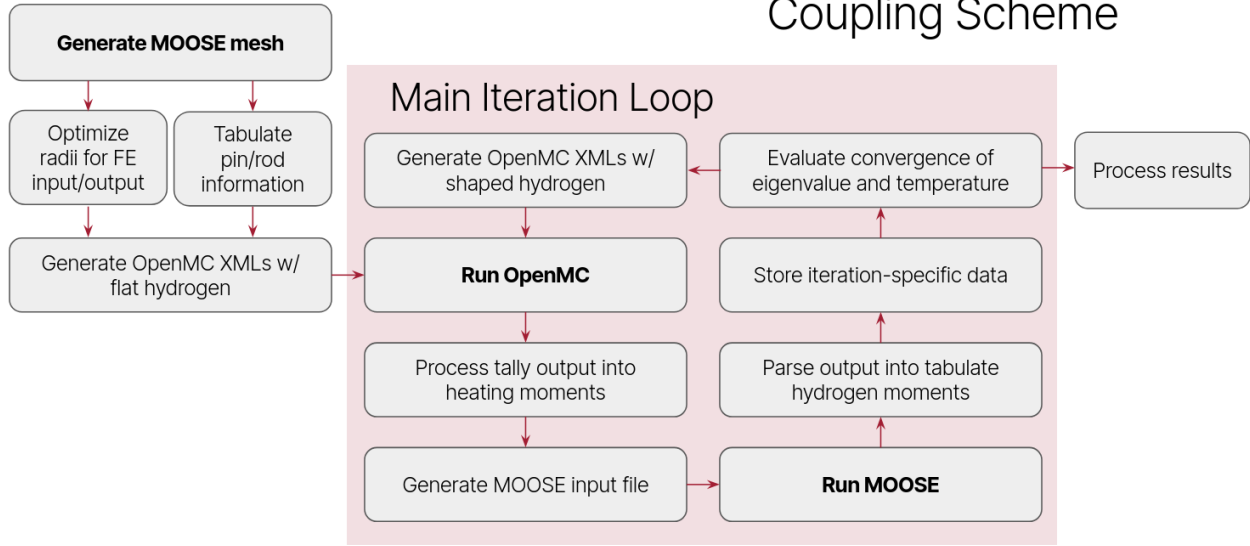


Figure 5.9: Detailed hydrogen diffusion coupled simulation scheme. All processes not in bold are performed via Python script.

were provided by the Virtual Test Bed Repository [52], [134] and were only slightly modified with regards to the assignment of mesh regions.

Optimizing radii for functional expansion input/output is a step necessary for accurately representing data used in or pulled from the mesh environment. This is described in detail in Section 5.2.4. Tabulating pin/rod information consists of generating a file that holds geometric parameters of every cylinder in the core, meaning fuel pins, moderator rods, and heat pipes. These parameters include coordinate origin, length, radius, and whether or not the cylinder is found on the boundary of the geometry. These parameters are used when generating the constructive solid geometry for the neutronics solve, as well as when defining normalization parameters for each functional expansion tally and enforcement.

Neutronic Solve

The neutronics portion of the main iteration loop begins with generating the XML files required for OpenMC operation. These XMLs contain data that informs geometry, material definitions, run settings, and tally details. The initial OpenMC run assumes a flat hydrogen profile in the moderator, uniformly throughout the core. If an iteration has already occurred, the material information for moderator rods will pull the MOOSE-calculated expansion coefficients to represent the spatial distribution of hydrogen within the ZrH. With these XML files defined, the actual OpenMC run occurs next.

After the run is complete, volumetric heating rates are prepared for the next MOOSE solve. This preparation mostly involves pulling cell heating tally results (units of eV/source particle) and converting them to volumetric heating rates (units of W/m^3), utilizing Equation 5.20. The same unit conversion process occurs for expansion coefficients, but with an additional orthonormalization constant. These volumetric heating terms are exported to a

separate file to be retrieved by the MOOSE input generation script.

$$q_i''' = \frac{H_i}{H_{\text{total}}} \frac{P}{V_i} \quad (5.20)$$

Finite Element Solve

The finite element solve picks up next. First, we generate a MOOSE input file with the previously mentioned volumetric heating terms and creating an input whose layout is detailed in Section 5.2.2. The MOOSE solver is then called. After the solve completes, the output is parsed to tabulate expansion coefficients describing the spatial hydrogen distribution. While these coefficients can be output into the console during the run, they cannot be automatically output into a generated file, unlike most MOOSE post-processors. Implementing this in the code should not be difficult, but I opted to instead parse the console output and tabulate them after-the-fact.

Following this, a single iteration is considered complete. Both MOOSE and OpenMC iteration-specific outputs (e.g. the statepoint file, output mesh, console outputs) are renamed and moved to preserve them for analysis. Next, the residual of the axial hydrogen profile in the hottest moderator rod is compared against an arbitrary convergence criterion. Convergence was found to occur quickly, usually converging after only 3-4 iterations. If it is determined that convergence has occurred, results are set to be processed, otherwise another set of OpenMC input XMLs will be generated.

5.2.2 Diffusion in MOOSE

As mentioned in the first introduction of MOOSE, one of the software's key strengths is its wide range of optional modules, featuring pre-built kernels for a wide variety of purposes. In the case of this chapter's work, a user-install of MOOSE comes with everything needed for the MOOSE-side of the simulation, minus the heat pipe representation. A graphical representation of the MOOSE input layout is included in Figure 5.10, not including executioner details, post-processors, material parameter inputs, and functions used in the aforementioned blocks. An explanation for each subset of the input, corresponding to the connected letters, is as follows:

- A **Kernel activating heat conduction and the input volumetric heating rates for every mesh volume.** All cylindrical volumes (moderator rods and fuel pins) have an associated `FunctionSeries` Function block featuring expansion coefficients generated from OpenMC output. Non-cylindrical mesh volumes are assigned their cell-average volumetric heating rate, also generated from OpenMC output.
- B **Heat pipe simulation.** First an axially-discretized surface heat flux calculation is performed on each heat pipe surface. Then a unique custom `UserObject` `MITHP` is called, linked to these fluxes as an input, in order to calculate a resulting axial temperature profile. Finally the custom boundary condition `SpatialQueryBC` samples the temperature profile generated by `MITHP` and establishes a `DirichletBC`. Again, also unique to each heat pipe.

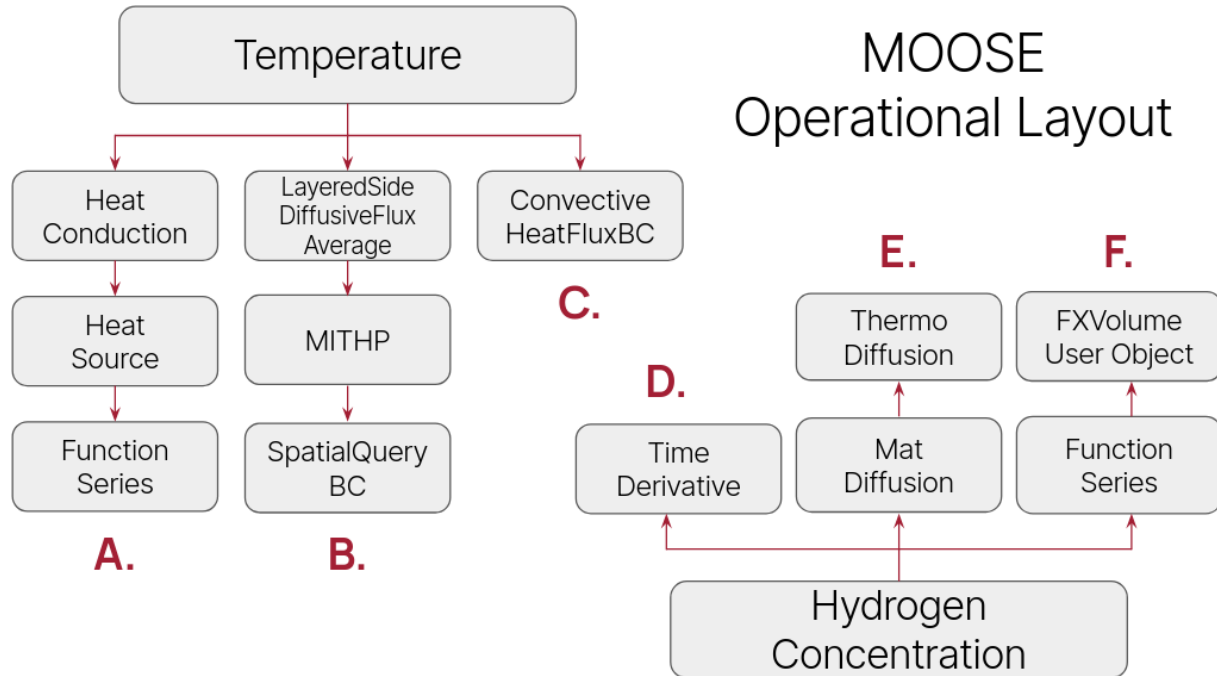


Figure 5.10: Graphical representation of the structure of the MOOSE solve, with some aspects omitted. Letters correspond to more detailed explanations that can be found in Section 5.2.2.

- C **Outer, top, and bottom boundary heat removal.** This is an arbitrary natural convection heat loss representation, mirroring the inputs used in [33] found on the Virtual Test Bed Github Repository [52]. This inclusion allows for a more realistic temperature profile for the outer air materials, given their otherwise isolated condition.
- D **Time Dependence Kernel.** This kernel includes a simple time derivative in the problem equation to solve.
- E **Diffusion kernels for hydrogen.** `MatDiffusion` is the kernel used for representing Fick’s law diffusion, and `ThermoDiffusion` is a handy pre-built kernel created specifically for Soret Effect-based diffusion.
- F **Blocks for calculating and returning expansion coefficients describing hydrogen concentration.** The `FunctionSeries` contains the mathematical framework for the convoluted cylindrical basis set, while the `FXVolumeUserObject` `UserObject` is called to actually perform the volume sweep and tabulate the requested moments.

The inclusion of a `TimeDerivative` block for hydrogen makes the spatial hydrogen solution a transient one, while the thermal solution is not time-dependent. This neglects the role hydrogen concentration plays in determining thermal conductivity of the zirconium hydride material, however its impact is relatively minor [136].

Mentioned in Section 5.1.3, leakage of hydrogen from the moderator into the reactor or a plenum of some sort is not modeled. This is mostly in order to reduce the complexity of the

simulation, as hydrogen desorption mechanisms can depend on factors like surface roughness and are instead usually approximated based on experimental results [113], [114]. However, future work on this topic should take this into account, as loss of moderator most often outweighs redistribution of moderator, and due to the risk associated with H₂ gas buildup in a reactor.

5.2.3 Continuously Varying Material Tracking

In Section 5.1.4, specifically at the end of the subsection on tallying and enforcement, I mention that OpenMC cannot receive functional expansion moments as an input for material definition purposes. This is, in fact, not an OpenMC specific constraint but instead a limitation of most Monte Carlo particle transport codes.

The Collision Distance Problem

Monte Carlo transport codes, regardless of using constructive solid geometry or CAD-based geometry, transport particles in straight lines (when simulating neutral particles like neutrons or photons), casting from their starting point to the first surface in their path. The code then needs to sample whether the particle will travel to that surface or if it will have a collision somewhere before. This distance-to-collision is sampled by evaluating the cumulative density function (CDF) describing the cumulative probability of colliding along the path:

$$F(s) = P_{\text{NC}}(s_b)H(s - s_b) + \int_0^s \frac{d\tau(s')}{ds'} e^{-\tau(s')} ds' \quad (5.21)$$

where s_b is the distance to the boundary surface, s is the distance along the particle path, and H is the Heaviside step function. $P_{\text{NC}}(s_b)$ is the probability of not colliding until the boundary, calculated via $P_{\text{NC}}(s) = e^{-\tau(s)}$ and $\tau(s)$ is the optical depth, defined as

$$\tau(s) = \int_0^s \sigma_t(s')N(s')ds' \quad (5.22)$$

where σ_t is the microscopic total cross-section and N is the number density. The actual process of sampling a distance requires inverting the CDF in Equation 5.21 and generating a random number to evaluate the stochastic collision location.

The issue arises when we remember that our plan is to replace the once-constant material definition of the moderator with a continuous functional expansion. These equations will rebuild a new distribution of hydrogen throughout our moderator, essentially changing the equation for macroscopic cross-section in material i from

$$\Sigma_{t,i}(T) = N_i \cdot \sigma_{t,i}(T) \quad \text{to} \quad \Sigma_{t,i}(\vec{r}, T) = N_i(\vec{r}) \cdot \sigma_{t,i}(T) \quad (5.23)$$

Although the convoluted Zernike and Legendre polynomials are simple and efficient to evaluate at any point in the cylindrical phase space, they are not analytically solvable in as part of an exponential along the characteristic line [30]. This means that the right-hand side of Equation 5.21 is not analytically solvable. Neither is P_{NC} for that matter, although it can be numerically integrated.

CVMT

In instances where analytical integration is not possible, the conventional solution is to numerically integrate and deal with a cost in computational speed. This isn't far off from being a very simplified explanation for the method of Continually Varying Material Tracking (CVMT). This method was developed by Brown and Martin [121], expounded up by Grisesheimer [116], and first utilized in the context of OpenMC by Ellis [30], who also optimized some of the algorithm to address 3D complexities. Ellis' work in particular was a significant source of inspiration for this chapter's work, as the use of functional expansions along with implementation of CVMT into OpenMC can be seen as following in Ellis' footsteps.

To begin with, the previously mentioned issue of P_{NC} needs to be resolved. Equation 5.22 can be approximated with a numerical integration scheme, which is particularly quick as it is a linear problem. The specific scheme used by Ellis is the three-point Newton-Cotes quadrature formula, or Simpson's rule of integration. Calculating $\tau(s_b)$ with Newton-Cotes is done by segmenting the path to the boundary into N intervals, and calculating optical depth in 2-interval steps via

$$\int_0^{2\Delta s} \Sigma_t(s') ds' = \frac{1}{3}a(2\Delta s)^3 + \frac{1}{2}b(2\Delta s)^2 + c(2\Delta s) \quad (5.24)$$

where Δs is the distance between each interval, and the constants a, b, c are

$$a = \frac{\Sigma_t(s_{n+2}) - c - 2b\Delta s}{4\Delta s^2} \quad (5.25)$$

$$b = \frac{4\Sigma_t(s_{n+1}) - \Sigma_t(s_{n+2}) - 3c}{2\Delta s} \quad (5.26)$$

$$c = \Sigma_t(s_n) \quad (5.27)$$

which when combined results in

$$\tau(s_{n+2}) = \frac{\Delta s}{3} [\Sigma_t(s_n) + 4\Sigma_t(s_{n+1}) + \Sigma_t(s_{n+2})] \quad (5.28)$$

With these equations, it is possible to numerically solve the optical depth at the boundary of the system.

In order to actually calculate the distance to collision, we can reuse the values already calculated for P_{NC} . It is important to remember that the optical depths at all intervals of the path have already been calculated, along with the polynomial constants a, b, c that describe the shape of Σ_t along each interval. In order to take advantage of these values that we already have, the CDF in Equation 5.21 can be inverted and the optical depth, $\hat{\tau}$, can be sampled instead, via

$$\hat{\tau} = -\ln(1 - (1 - P_{NC}(s_b))\zeta) \quad (5.29)$$

where ζ is the random number generated. This provides a randomly sampled $\hat{\tau}$ that can then be compared against the previously calculated interval τ_s in order to isolate the interval n in which the particle collides. To find the exact location in that space, one just calculates $\delta\tau = \hat{\tau} - \tau$ and retrieves s' from

$$\frac{1}{3}a(s')^3 + \frac{1}{2}b(s')^2 + c(s') = \delta\tau \quad (5.30)$$

in order to finally move the particle a distance of $(n-1)\Delta s + s'$ along the flight path. An algorithmic interpretation of this process can be seen in Algorithm 1. These explanations are a simplification of the complexity and optimization that exists in the true CVMT process, more detailed information can be found in [30].

Algorithm 1 Collision distance sampling using CVMT, simplified from [30]

- 1: Segment path to cell boundary into N equally-spaced intervals
 - 2: Compute $\tau(s_b)$ using Newton-Cotes method, keeping interval data
 - 3: Compute $P_{\text{NC}} = \exp(-\tau(s_b))$
 - 4: Sample random number ζ_1
 - 5: **if** $\zeta_1 \leq P_{\text{NC}}$ **then**
 - 6: Move particle to cell boundary
 - 7: **else**
 - 8: Sample random number ζ_2
 - 9: Compute $\hat{\tau} = -\ln[1 - (1 - P_{\text{NC}})\zeta_2]$
 - 10: Find interval n within N that $\hat{\tau}$ lies
 - 11: Utilize interval n 's precomputed coefficients to retrieve true collision location
 - 12: Move particle to location
 - 13: **end if**
-

This represents a significant increase in computational steps during particle transport. What was once a simple evaluation by integrating a constant cross-section across a distance is now a large number of Simpson's rule integrations. It is still faster than geometrically discretizing the Monte Carlo problem to recreate complicated shapes via 0th order moments [30].

When implementing CVMT in OpenMC, I was fortunate to rely on the work of my colleague Jiankai Yu who had previously coded CVMT into his personal branch of OpenMC. Using this resource, I updated this implementation to the current version of OpenMC and expanded some of its functionality to account for the convoluted basis set of Zernike and Legendre polynomials.

In OpenMC, the standardization parameters and expansion coefficients are included in the material definition of the moderator. "Standardization parameters" here refer to values needed for normalizing the geometric space, like axial length for Legendre polynomials and radius/origin point for Zernike polynomials. These parameters and coefficients are stored in the material definition and a flag is raised in the code whenever a particle tracks across the material. That flag triggers the CVMT process for determining collision probability and distance-to-collision. One side effect of this method of definition is that because the standardization parameters, specifically the disk origin point, are unique for each moderator rod, each moderator rod must receive a unique material.

Beyond that, this also means that the track-length estimator isn't usable for tallying, as it would require the same numerical integration process but on the tallying side of the OpenMC particle history workflow. This is why all tally results in this chapter use the

collision estimator. Additionally, eigenvalue estimation in OpenMC is typically reported via a "k-effective", a statistically optimized combination of all estimators which includes the tracklength-based eigenvalue [137]. Instead, the collision based eigenvalue must be used, and particle batch normalization must be done via that same collision eigenvalue.

Regardless of some inefficiencies, this CVMT method is extremely powerful in that it allows particle transport over materials with spatially varying properties. This allows diffused hydrogen to be represented in the moderator of the microreactor geometries studied in this chapter, and would allow accounting for phase changes in materials where thermal scattering physics will vary for future works. It is also the case that almost all real-world materials are spatially-varying in their properties, so the use of CVMT allows us to get one step closer to a true-to-life simulation of particle transport.

Other Methods

One last note on this topic, is that CVMT is not the sole method of handling spatially varying material properties in a neutronics solver. Serpent 2 features a method similar to Woodcock's delta tracking [138] where a scheme featuring rejection sampling based on the maximum total cross-section in the material is implemented [139]. In brief, the method begins with sampling a distance-to-collision by inverting and sampling the CDF in Equation 5.21, but using the maximum total cross-section to do so:

$$s = -\frac{1}{\Sigma_{\max}} \ln(\zeta_1) \quad (5.31)$$

This collision point is accepted if a second randomly generated number ζ_2 satisfies

$$\zeta_2 < \frac{\Sigma_t(\vec{r}')}{\Sigma_{\max}} \quad (5.32)$$

where $\vec{r}' = \vec{r} + s\vec{\Omega}$, and \vec{r} being the starting position. If the point is rejected, the procedure starts again by sampling a new path length starting from position \vec{r}' .

This method was not used for this work, as calculating the maximum total cross-section can be difficult when using continuous distributions of number densities [30], [140]. Furthermore, a significant part of the interest in the work of this chapter is in the use of functional expansions, for which CVMT has proven functionality in an OpenMC environment [30].

5.2.4 Function Expansion Use in Meshes

All core geometry meshes used in the work of this thesis are generated via the Reactor module of MOOSE [40]. This module expands the mesh generation capabilities of MOOSE, specifically tailored for reactor core geometries with Cartesian or hexagonal lattice patterns. Included are functionalities to mesh concentric circles, stitch different mesh cells together, define boundaries, create control drums, extrude geometries, and reassign mesh block ids, all meshed symmetrically, and all within the format of a MOOSE input file. The procedural way the mesh is generated and standardization of each mesh feature makes the output mesh highly useful for the work of this chapter.





Figure 5.11: Arbitrary functional expansion heating distribution used in conjunction with Figure 5.12. No axial shape is included.

One aspect of this meshing process is that circular geometries are polygonized, as node-to-node edges are linear. For example, because the mesh cannot use quadratic edges, an input circle of radius 1.0 would become a hexagon of side-length 1.01 to preserve area. Preserving the area (in 2D) is the most important aspect here, as that conserves mass in the system, but the polygonization means that the application of functional expansions is slightly erroneous.


As an Input

Part of the convoluted basis set, specifically the Zernike polynomials, is orthonormalized to cover a unit disk, requiring a normalization radius to leverage standardization. MOOSE samples the basis set at discrete quadrature points which lie on the aforementioned polygonal mesh. If MOOSE samples a point that lies outside of the normalization radius, the value sampled will be set to zero. This almost always happens when using a first-order mesh that conserves volume. This is a problem, as the input functional expansions for the MOOSE solve describe volumetric heating, and erroneous zero's in sampled heating rate may translate to incorrect temperature profiles or total input power. An example of this error occurring is included in Figure 5.12, with the input function expansion distribution included in Figure 5.11.

In order to reduce this error, an "optimized" radius is needed that will fit the cylindrical rebuilt power distribution to the true polygonal mesh with minimal error. Defining the error is tricky in this context, however. What we're looking for is a slightly modified normalization radius for the functional expansion, the radius used in the definition of the Zernike polynomials, that rebuilds the distribution on the polygonal mesh in a way that accurately represents the original distribution in a true cylinder. There's obviously not a one-to-one equivalent when the geometry itself has changed. Instead, nodal error, ϵ_n , is defined by comparing the



(a) MOOSE output nodal heating rate, when normalization radius = 1cm.



(b) MOOSE output nodal heating rate, when normalization radius = 1.05cm.


Figure 5.12: Figures showing error when normalization radius does not account for polygonal mesh geometry. The cylindrical radius is 1cm.

nodal value, V_n , to the functional expansion rebuilt value, $FX(z_n, r_n, \theta_n)$, where the node radius is normalized by the maximum radius of all nodes in the mesh volume, Equation 5.33. This suffices to minimize the error that occurs from not correctly "covering" each node, and can be quantified in a root mean square error (Equation 5.34).

$$\epsilon_n = V_n - FX(z_n, r_n, \theta_n) \quad (5.33)$$

$$\text{RMSE} = \sqrt{\frac{\sum_{n=0}^N \epsilon_n^2}{N}} \quad (5.34)$$

With a quantification of error defined, a Python minimization script is used to identify the optimal radius. Beginning with an exact mesh copy of the fuel pin or moderator rod, the Python script generates and runs MOOSE inputs that use an arbitrary functional expansion describing heating rates and take the normalization radius as an input. These heating rates are converted to an AuxVariable via `FunctionSeriesToAux` that way the post-run output mesh will contain the MOOSE sampled function expansion value at each node. With this, the Python script can load these nodal values and calculate an error via Equations 5.33 and 5.34. This process can be run with the normalization radius as an input, so a minimization script is run to determine the optimal normalization radius. The initial sweep along with minimization can be seen in Figure 5.13a, and the detailed optimization of the normalization radius can be seen in Figure 5.13b.



(a) Normalization radius sweeping and minimization.

(b) Normalization radius and RMSE versus iteration count.

Figure 5.13: Figures showing normalization radius convergence and error minimization for functional expansion input. The resulted in an optimal normalization radius of 1.024cm when used with the originally 1cm radius fuel pin.

This process of determining an optimal normalization radius for function expansion standardization must be done for each unique cylindrical geometry. That means once for the moderator rod geometry and once for the fuel pin geometry, for each reactor core. Given that these values do not change during the coupled MOOSE-OpenMC solve, they are pre-computed first and inserted as inputs for MOOSE input file generation.


As an Output

Similarly to when using functional expansions as an input in the previous section, when using MOOSE to return expansion coefficients to describe a spatial distribution, the normalization radius is not necessarily the original radius of the cylinder. That's because we're now trying to represent a non-cylindrical nodal data set with a set of basis functions defined over cylindrical space. To reduce the error in this representation, we'll adjust the normalization radius to better match that nodal distribution.

In this case the MOOSE nodal output is the true case. Mathematically, all that changes is the definition of ϵ_n , which becomes:

$$\epsilon_n = FX(z_n, r_n, \theta_n) - V_n \quad (5.35)$$

The error minimization script functions mostly the same, though now an arbitrary temperature gradient is forced across the mesh (which is identical to 5.12 but with units of Kelvin) and the functional expansion seeks to represent this temperature profile. The optimized normalization radius is not drastic for this case. An example of this is provided in Figure 5.14, showing the convergence of an optimal normalization radius for a moderator rod of cylindrical radius 0.975cm. The results show that a slight increase in radius from the prescribed cylindrical radius was optimal in minimizing error.



(a) Normalization radius sweeping and minimization.

(b) Normalization radius and RMSE versus iteration count.

Figure 5.14: Figures showing normalization radius convergence and error minimization for functional expansion output. The resulted in an optimal normalization radius of 0.9755cm when used with the originally 0.975cm radius moderator rod.

Because using this optimized radius only constitutes a change of RMSE of roughly 1.5 Kelvin, and using the cylindrical radius only had an average percent error of roughly 0.4%, this step can most likely be omitted. Some amount of error is guaranteed here, by nature of trying to represent a non-cylindrical distribution with a basis set of equations that are by-definition cylindrical. Defining that error is difficult, but does deserve future work. The use of second order meshing schemes may remove this error altogether, due to very minor shape differences with the cylinder.

5.2.5 Functional Expansion Use with Annular Cylinders

The work in Section 5.2.4 can be generalized as efforts to reduce the error in representing a not-quite cylindrical distribution with a cylindrical basis set, or vice versa. This process becomes more difficult when attempting to do the same, but with an annular geometry, as is the case for the fuel in the SR reactor (see Section 2.3.3). This annular fuel provides benefits in lowering center-line temperature and being advantageous for cases of severe swelling, but using functional expansions in the context of an annular geometry requires some thought.

The most obvious impact of annular geometry is that the Zernike basis functions are no longer orthogonal in their use, as the fuel phase space is from r_{inner} to r_{outer} rather than from 0 to r_{outer} . This causes an error in the 0th moment proportional to the difference in area. For the SR core fuel design, there is a difference of 26% in volume due to using annular fuel versus cylindrical, which causes a 26% decrease in Z_0^0 . A second similarly strong impact of the annular geometry is that higher order polynomials are required to correctly describe the radial shape of the solution. Because the material only begins at a finite r_{inner} , the true solution has a piece-wise effect, with a discontinuity that the continuous basis functions struggle to represent.

To quantify the total error of using Zernike FETs on an annular geometry, a 2D test case

was generated with a single annular fuel pin surrounded by water with reflective bounds. A cylindrical mesh tally was applied to the fuel, along with 0-10th order Zernike FETs. This same process was done for a contiguous cylindrical pin as well, to act as a control. The mesh tally results for both cases can be seen in Figure 5.15. Note how relatively little gradient there is from inner to outer for each geometry, as well as the symmetry in the results.

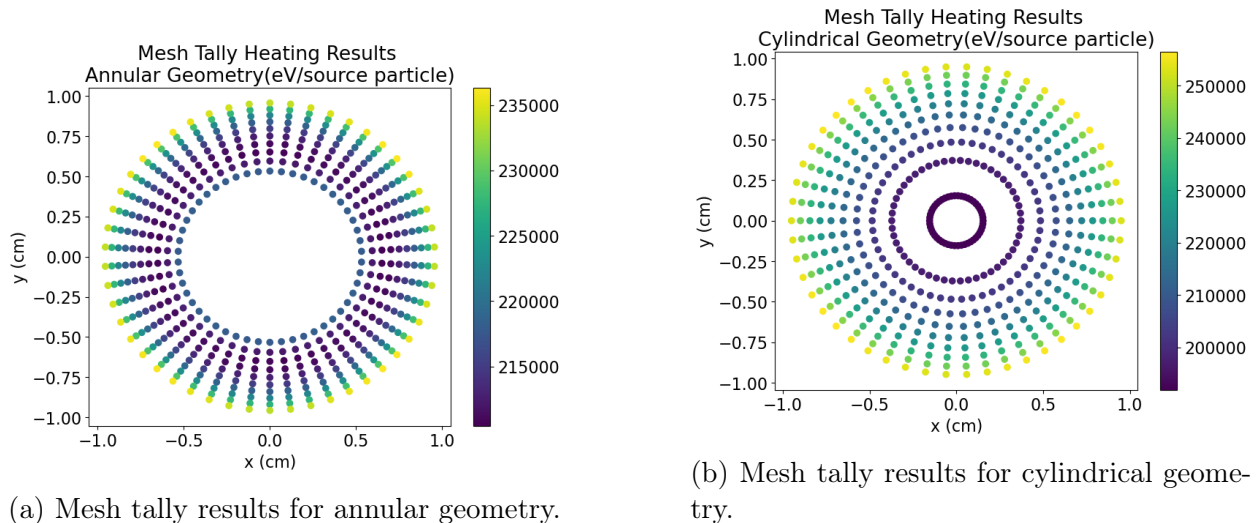


Figure 5.15: Heating mesh tally results for annular and cylindrical geometries. Units are eV per source particle.

A quantification of the L2 norm of the percent error caused by the Zernike FET use can be seen in Figure 5.16. The second line plotted shows the error excluding the innermost points, where the previously mentioned error from representing a piece-wise problem is at its largest. Also included is the same comparison but for a contiguous cylinder, showing the expected increase in accuracy. Both plots contain a comparison versus using a simple cell-average tally value.

The results in Figure 5.16 show a strong initial error due to a normalization error due to the reduced volume of the annular fuel pin, with subsequent orders decreasing that error. For this thesis, 5th order Zernike polynomials are the highest order pursued, so this would be a significant error.

$$q_i''' = \frac{H_i}{H_{\text{total}}} \frac{P}{V_i} \quad (5.36)$$

In coupled simulations, we generate a normalization factor using an additional volumetric tally of the heating rate. This takes care of the unit conversion to W/m³, via equation 5.36. The H_{total} value does not come from the 0th moment of the tally, but instead a cell tally which does not contain the previously mentioned error. This means that the FET expansion coefficients are renormalized to conserve the correct power in each volume, relying on the moments to only shape that power.

Taking the results from the aforementioned test problem and normalizing them to a hypothetical power, another comparison of error from using the functional expansion can be generation. These results can be seen in Figure 5.17. This normalization allows the 0th order

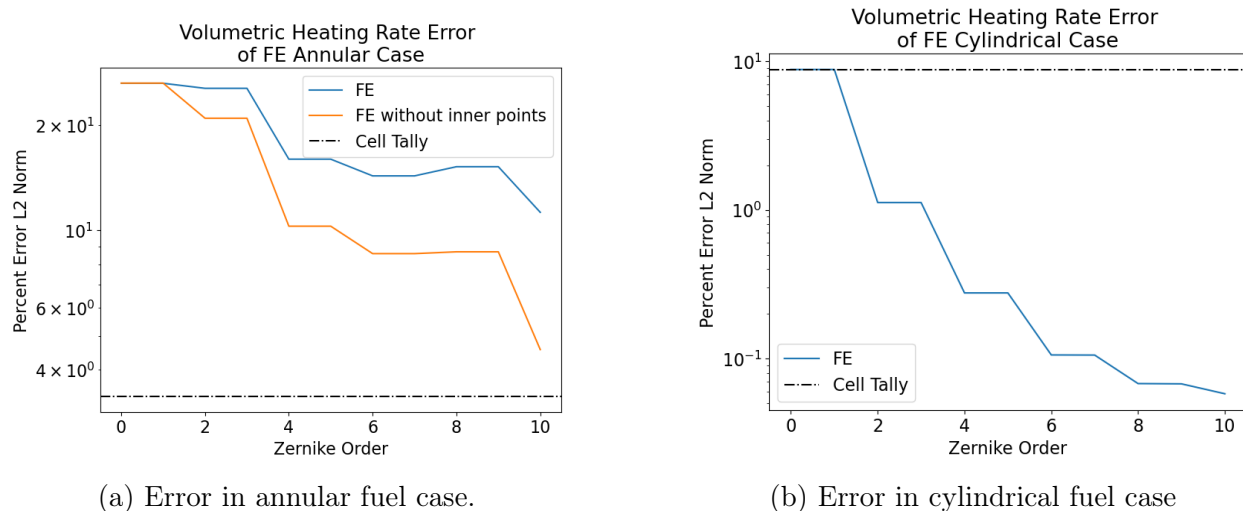


Figure 5.16: L2 norm of the percent error from using Zernike FETs versus a fine cylindrical heating mesh tally.

moment to match the cell tally value, but interestingly increasing order actually causes a large increase in error, followed by a gradual decrease that fails to reach the error from using the cell-average value. This is due to additional terms increasing the average power away from the correct value.

These results indicate that for the radial and azimuthal domain, the use of Zernike FETs will result in an introduction of error, even with power normalization applied. For this reason, the results of the SR reactor described in the results section of this chapter do not include the use of FETs to describe the spatial heating rate in the annular fuel. In the case of more strongly shaped results, or results with strong asymmetry, the use of the Zernike FETs may beat the cell-averaged case, depending on the order of polynomials used. Additionally, the thinner the annulus fuel region, the higher the order of Zernike polynomial will be required to match the geometric shape.

Note that the use of Legendre polynomials for the axial power shape has no inherent error because of the annular geometry. I decided not to use any functional expansions for the fuel heating rate, however, in order to identify the ramifications of this choice. As will be seen in the results of Section 5.4, the majority of the impact of hydrogen migration is found in mild changes to the shape of heating in the moderator and fuel, so losing the ability to spatially represent the volumetric heating of the fuel will be significant. Future work of this type with annular geometries should include the use of an FET basis set that is orthonormal over an annular geometry.




Figure 5.17: L2 norm of the percent error from using Zernike FETs with normalization versus a fine cylindrical mesh tally.




Figure 5.18: 1D test case graphical summary.

5.3 Test Cases

The test cases listed in this section are meant to introduce the effects of hydrogen diffusion in simpler contexts. This gives a baseline understanding to use when analyzing full-core effects that may otherwise be difficult to parse.

5.3.1 1D Case

This test case features a 1D geometry with a monodirectional source of neutrons at one end, and vacuum boundaries at both end. In the middle of the geometry is a span of zirconium hydride, featuring hydrogen with a modifiable shape along the single axis.


The problem set up here is applying an arbitrary Legendre polynomial cosine shape to the hydrogen distribution over the 10cm long material section. Then, a flat hydrogen distribution is imposed, and the width of the material is varied. The goal is to determine what "equivalent" length of material with the flat shape matches the reflective and transmissive result of the 10cm length of material with the cosine shape. A graphic of this setup can be seen in Figure 5.18. Note that the cosine shape has the same average hydrogen concentration as the flat shape, representing $\text{ZrH}_{1.67}$.

The results of this study can be seen in Figure 5.19. For the transmission results, the shape of the hydrogen in the material is shown to not impact transmission probability. This makes sense, as the cosine shape has the same average hydrogen concentration as the flat case, so a particle cumulatively "sees" the same amount of hydrogen in both cases. Mathematically one can prove this, as the equation for passing through the material without colliding is

$$P_{NC} = e^{-\int_{x_0}^{x_1} \Sigma_t(x) dx} \quad (5.37)$$

and the integral of the Legendre-based cosine shape is the same as the flat value, so P_{NC}

will be identical. There is a difference, however, for reflection probabilities, with the flat case requiring an additional 0.25cm of width to match the reflection probability of the cosine case. This occurs because in cases where a fast neutron collides in the material and scatters backwards, the particle will see less hydrogen on its way out in the cosine case. This is verified by tallying the current energy on the reflective side, which shows an increase in fast particles being reflected versus the flat case. This implies that hydrogen migration could make zirconium hydride a better reflector.



(a) Transmissive comparison between cosine case and flat case w/ varying widths.

(b) Reflective comparison between cosine case and flat case w/ varying widths.

Figure 5.19: 1D test case results comparing transmission and reflection probabilities.

While interesting, zirconium hydride is unlikely to be used as a reflector in any reactor context because designers generally don't want reflectors to moderate neutrons, just send neutrons back into the bulk of the core. Isotopes with higher masses are generally used for this purpose. The phenomenon recorded in this test case could still impact large moderator regions, such as for a reactor design that features large moderator blocks rather than moderator rods.

5.3.2 2D Case

The 2D test case is similar, but now includes a finite y-dimension, and the Legendre polynomial describing hydrogen density in the middle moderator is now aligned with the y-axis. This allows a more realistic test of what would be envisionable in the true reactor; Migration of hydrogen along the y-axis will change the spectrum of neutrons (traveling in the x-direction) that pass through the moderating material. This is important because a decrease in number of thermal neutrons will directly translate to a decrease in fission rate.

A simplified graphic of the workflow is attached in Figure 5.20. To summarize, a temperature gradient is fed to a MOOSE input file which solves the 2D temperature and hydrogen diffusion problem on a representative geometry. The Legendre expansion coefficients describing the migrated hydrogen profile are output and transferred for the OpenMC run. This OpenMC run uses a mono-directional planar source on one side of the moderating material, and tallies the spectrum of the neutrons that exit the other side. Both the temperature gradient and width of the moderating material are varied to build a detailed set

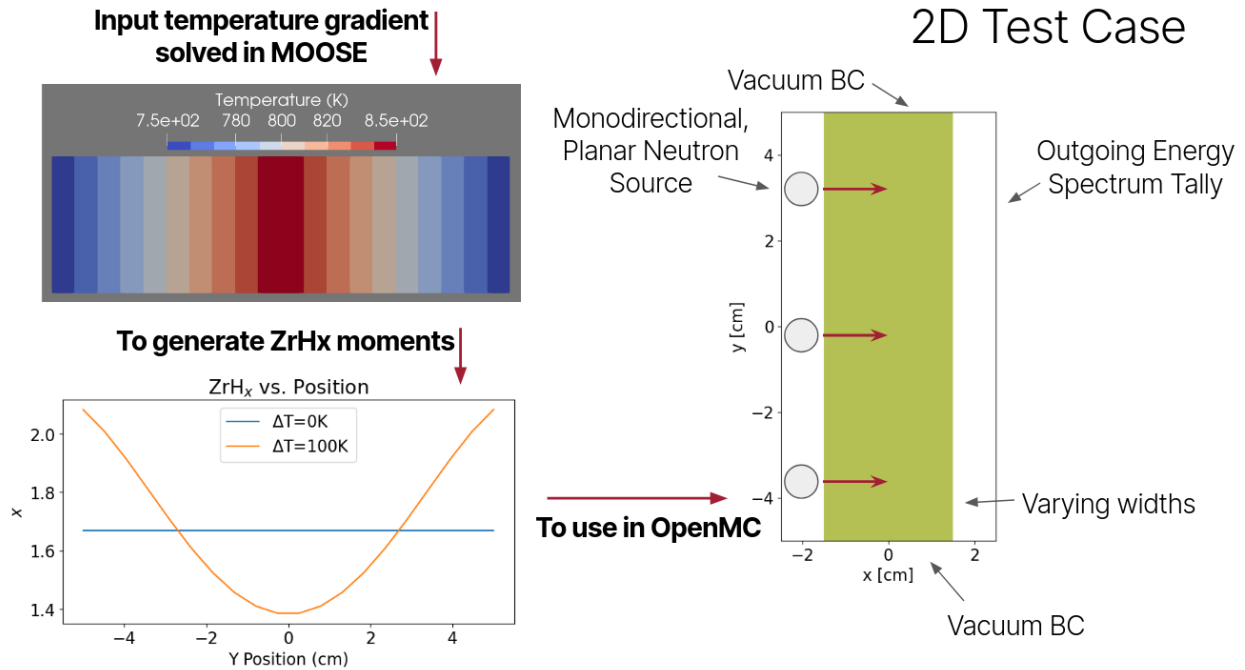


Figure 5.20: Simplified graphic of the workflow for the 2D test case.


of data. Temperature gradients are represented by a linear function peaked in the center of the geometry with an average temperature of 800 Kelvin, so a ΔT of 50 K would mean the center is 825 K and the min and max y positions are 775 K.

An example of how the temperature gradient impacts neutron moderation can be seen in Figure 5.21, which shows the impact to hydrogen distribution and thermal neutron current. Because the center peaked temperature drives hydrogen towards the top and bottom of the y -axis, the spectrum in the center hardens, and less thermal neutrons are seen by the tally. In this case "thermal" neutrons are neutrons with energies at or below 10eV.


Performing this process for a range of moderator widths and ΔT s generates Figure 5.22, a contour plot showing total thermal neutron current versus ΔT and width. Based on the direction of the gradient, it is clear that temperature gradients reduce thermal neutron current out of the moderator, and this effect is exacerbated the wider the moderator is.

What is interesting, is that equivalences can be made. For a moderator of width w and temperature gradient δT , there is a "true width", that is the width of a moderator with no gradient that would result in the same thermal current. For example, with $w = 5$ cm and $\delta T = 100$ K, the "true width" is actually 4.72 cm. The inverse of this is a "necessary width", or the width of moderator required to equal the total thermal current of the moderator before the temperature gradient was applied. For the case of $w = 5$ cm and $\delta T = 100$ K, this "necessary width" is 5.31 cm. This means a hypothetical reactor designer would need to increase moderator width by 0.31 cm to account for the migration of hydrogen lowering thermal neutron flux in the core.

Obviously this is not realistic for a real-world reactor, given a couple different reasons. Firstly, the thermal gradient here is much larger than will be seen in a reactor. While the



(a) Hydrogen concentration shift due to temperature gradient input.



(b) Axial thermal neutron current tally with temperature gradients. Units are neutrons per source particle.

Figure 5.21: Hydrogen migration result due to different ΔT values and their respective impact on axial thermal neutron current. Asymmetry is due to stochastic uncertainty.

total delta is only 100 degrees, the geometry it is applied over is only 10 cm long, so the derivative $\frac{dT}{dy}$ is a whopping 1000 Kelvin per meter. This is why the ZrH_x profile in Figure 5.21a goes beyond the actual limit of ZrH_x . Additionally, a real reactor would most likely have axial reflectors, as all geometries studied in this thesis do. This causes the thermal flux to move towards the top and bottom of the core, rather than simply decrease.

This case study does however point out the general trends that will be seen in the 3D core results, most notably a hardening of spectrum in the center of the core. While the results in Section 5.4 are not drastic enough to imply a need for wider moderators, it is good to note that this would be one possible remediation to the effects of axial hydrogen migration in the core.





Figure 5.22: Contour plot of total thermal neutron current versus temperature gradient and moderator width.



(a) "True width" versus thermal gradient input at three varying moderator widths.

(b) "Necessary width" versus thermal gradient input at three varying moderator widths.

Figure 5.23: (a) shows "True Width" versus thermal gradient. This value is the equivalent width of moderator with no temperature gradient that would result in the same total thermal neutron current leaving the moderator. (b) shows "Necessary Width" versus thermal gradient. This value is the equivalent width of moderator needed that would conserve the original $\Delta T = 0$ thermal neutron current leaving the moderator.

Table 5.2: OpenMC particle count settings for hydrogen diffusion studies.

Parameter	Count
Active Batches	500
Inactive Batches	50
Particles	80000

5.4 Results

This section details the results of performing the neutronic-thermal coupling with zirconium hydride hydrogen redistribution simulated. Results of each run can be categorized into

- **Hydrogen Results.** These being the analytics of the hydrogen shift: where hydrogen has concentrated, where it has vacated, what trends are there, etc.
- **Neutronic Impacts.** These include changes to eigenvalue, thermal fission factor, pin power, power shape, etc.
- **Thermal Impacts.** These consist of changes to the temperature profiles of the fuel and moderator.

Each reactor study also includes a heat pipe failure study wherein multiple heat pipes' heat removal effects are disabled. Note the caveats to this described at the end of Section 5.1.2, as moderator temperatures in heat pipe failure scenarios will exceed the δ -phase limits. Each heat pipe study will highlight changes to the above results, focused on the area surrounding the failed heat pipe. OpenMC particle count parameters can be seen in Table 5.2 and is consistent for each core study.

As mentioned previously convergence of the coupled simulation is based on the residual of the axial hydrogen profile in the hottest moderator rod. Typically, convergence would be evaluated against a global property like the eigenvalue, however for thermal expansion the effect on the eigenvalue is so low (as will be seen in the following sections) that uncertainty dominates the eigenvalue residual. Instead, comparing the change in axial hydrogen profile of a moderator rod allows convergence to be determined against the physical property I am most interested in. The specific value is calculated by taking the L2 norm of the relative residual of the hydrogen concentration along the moderator rod segmented into 8 equal lengths. This convergence for each reactor can be seen in Figure 5.24, showing that after the roughly 3rd iteration values are mostly converged. The criterion used in this study was criterion of $1.0e-4$, arbitrarily set to encourage extra iterations where they may not necessarily be needed. Note how the SR core has a low value even after the first iteration, due to the low migration that will be seen in Section 5.4.3.





Figure 5.24: L2 norm of axial hydrogen residuals for hottest moderator rod, for each core design during hydrogen diffusion coupled simulation.

5.4.1 MRaD Core

The MRaD core geometry, detailed in Section 2.3.1, features an input thermal power of 345.6 kW, and OpenMC-MOOSE coupled simulation results in the attached temperature plots of the core in Figure 5.25. Heat pipes remove heat vertically from the core, causing the upper regions of the core to be cooler on average. Temperatures peak in the inner assembly, with a roughly 40 Kelvin discrepancy between the inner assembly and outer assemblies. The control drums stand out as being hotter than their neighboring reflector blocks because the gas-filled gap acts as an insulator from heat pipe-based heat removal. Note that neighboring fuel and moderator rods are nearly identical in temperature, a feature that will change in the studies of the Simba and SR cores, due to the inclusion of a gap around the moderator and cladding of the fuel, respectively.



(a) Isometric view of MRaD Core temperatures after final coupled iteration.

(b) Axial mid-plane slice view of MRaD Core temperatures after final coupled iteration.

Figure 5.25: MRaD temperatures after final coupled iteration.

Hydrogen Results

By parsing the nodal data included in Figure 5.26, some statistics about the hydrogen profile can be generated. The minimum and maximum x in ZrH_x are 1.625 and 1.747, respectively. This is a -2.7% and 4.6% change in hydrogen concentration, respectively. This method can also be applied to axial slices, for example, in Table 5.3. This table shows an average movement of hydrogen away from the center towards the top and bottom, the top in particular.




Figure 5.26: Final hydrogen concentration values in MRaD moderator rods, returned by MOOSE.

The radial and azimuthal shapes of the diffused hydrogen are difficult to describe holistically, as each rod's shape responds to its surrounding thermal environment, and every rod sees a unique environment. Instead, we can focus our attention on the central hottest fuel pins and moderator rods, located in the inner-most assembly. This constitutes three fuel pins alternating with three moderator rods, concentrically surrounding a heat pipe. Figures 5.27 and 5.28 feature X-Y slices of these fuel pins and moderator rods, located axially at the midpoint ($z = 80$ cm) and the top ($z = 160$ cm), respectively.


These results clearly show the radial/azimuthal response to the thermal effects of the surrounding fuel and central heat pipe. It is interesting to note that these effects are most strongly felt in the axial center ($z=80$ cm). At the top of the core ($z=160$ cm), the temperature gradients are significantly less and the overall temperature is lower. This causes a hydrogen profile that is flat in radial and azimuthal shape and on average higher than in the mid-plane. This is consistent with the general trend, that axial migration is the dominant migration form that occurs in the core.

Table 5.3: Final axial values for x in ZrH_x , sampled from all MRaD moderator rods. The layer under each result are the percent change from original $ZrH_{1.67}$


Axial Location (cm)	Average x	Min. x	Max. x
z=0	1.685	1.658	1.726
Δ	0.9%	-0.7%	3.3%
z=80	1.662	1.626	1.716
Δ	-0.5%	-2.6%	2.8%
z=160	1.705	1.673	1.747
Δ	2.1%	0.2%	4.6%

To describe the axial response of the hydrogen, the MOOSE-output hydrogen expansion coefficients for the hottest moderator rod can be numerically integrated in axial segments. When divided by volume, these values represent the axially-averaged ZrH_x that OpenMC will sample during particle transport. These values for the hottest moderator rod in the MRaD core are included in Figure 5.29, showing a clear axial migration towards the top of the core. A decrease in average hydrogen stoichiometric ratio occurs in the center of the length, peaking at roughly 1% decrease. This is directly due to hydrogen migrating away from areas of higher temperature.

Note that when looking at the outer moderator rods in Figure 5.26, this same shape appears, constituting a decrease of H/Zr ratio in the center and an increase at top of the rod. The magnitude of this shift is far less however. These rods are only approximately 40 Kelvin cooler than the peak moderator rods, indicating that hydrogen migration rapidly increases in magnitude as the temperature approaches the 860 °K range.




(a) Fuel and moderator temperature profiles.




(b) Moderator x in ZrH_x profile.

Figure 5.27: Temperature and H/Zr stoichiometric ratio of hottest moderator and fuel, sampled via axial slice at the core mid-plane.



(a) Fuel and moderator temperature profiles.



(b) Moderator x in ZrH_x profile.

Figure 5.28: Temperature and H/Zr stoichiometric ratio of hottest moderator and fuel, sampled via axial slice at the core top.




Figure 5.29: Axially-averaged hydrogen profile for MRaD hottest moderator rod.


Table 5.4: Global neutronic results and differences from initial flat-hydrogen case for MRaD core. "AR" is Absorption Rate, "Thermal FF" is the ratio of fissions caused by thermal neutrons compared to total fissions. Absorption rates are in units of "per source particle". \pm values are one standard deviation.

Quantity	Result Value	Difference	Pct. Dif. (%)
Eigenvalue	1.02689 ± 0.00021	-0.00032 ± 0.00025	-
Thermal FF	$0.79989 \pm 5.0e-5$	$-0.00028 \pm 6.0e-5$	-0.034 ± 0.008
Fuel AR	$0.65513 \pm 1.2e-4$	$0.00003 \pm 1.4e-4$	0.004 ± 0.021
Heat Pipe AR	$0.10636 \pm 2.9e-5$	$-0.00012 \pm 3.5e-5$	-0.112 ± 0.033
Moderator AR	$0.04373 \pm 9.0e-6$	$-0.00016 \pm 1.1e-5$	-0.371 ± 0.024
Reflector AR	$0.13731 \pm 7.0e-5$	$0.00022 \pm 8.0e-5$	0.160 ± 0.060

Neutronic Impact

While quantifying the shift in hydrogen is interesting to note, what is of importance is the downstream effects that shift has on steady state reactor operation. These effects begin with an impact to core neutronics, as redistribution of moderator means redistribution of flux and flux spectrum.

Some of the important global neutronic final-iteration quantities and changes are listed in 5.4. The general takeaway is that there's essentially no impact to global neutronic behavior. This can be confirmed further in Figure 5.30, showing the change in fuel pin power as a result of hydrogen migration. There's no clear, dominating increase or decrease in nominal pin power in a particular section of the core. Additionally, standard deviations of respective pin power changes make some of the larger changes uncertain, particularly in outer assemblies.




(a) MRaD percent change in pin powers from initial to final iteration.

(b) Standard deviation of percent change in pin powers.

Figure 5.30: Percent change in pin powers of MRaD core due to hydrogen migration, and their respective standard deviations.

While global effects are hard to pin down, local effects centered around the highest

temperature region in the core are much more evident. Figures 5.31 and 5.32 display axial profiles of volumetric heating rate for the hottest fuel rod and moderator rod. Each sees a minor decrease in power in the center, and a slightly larger increase in power at the top, mirroring the hydrogen movement from Figure 5.29. Note that the percent change in power is relative to that axial segment's initial power. A 2% increase in power at the final segment of the moderator rod may represent less overall power change than a 0.5% decrease in power in a central segment of the rod, for example. Radial and azimuthal effects of the hydrogen




(a) Volumetric power versus axial position for fuel pin, in initial and final iteration case.


(b) Percent comparison of volumetric power versus axial position for fuel pin.

Figure 5.31: Axial segment-averaged volumetric heating rate of MRaD hottest fuel pin from initial and final iteration, and the percent change between the two.

change were minimal in magnitude, and were dwarfed by the axial changes. For that reason they were omitted from this section of results.



(a) Volumetric power versus axial position for moderator rod, in initial and final iteration case.



(b) Percent comparison of volumetric power versus axial position for moderator rod.

Figure 5.32: Axial segment-averaged volumetric heating rate of MRaD hottest moderator rod from initial and final iteration, and the percent change between the two.

Thermal Impacts

Temperature changes follow any change in power distribution. As noted in the previous subsections, for the case of hydrogen migration in the MRaD core geometry changes to power distribution are mostly axial and located at the hottest region of the core. By comparing the temperature outputs between the initial and final iterations, Figure 5.33 was generated.

This figure shows a decrease in temperature in the mid-plane due to the migration of hydrogen, peaking at a decrease of roughly 0.3 Kelvin. This decrease is noticeable in the simulation, but would hardly register experimentally, given the inherent uncertainty of thermocouples.

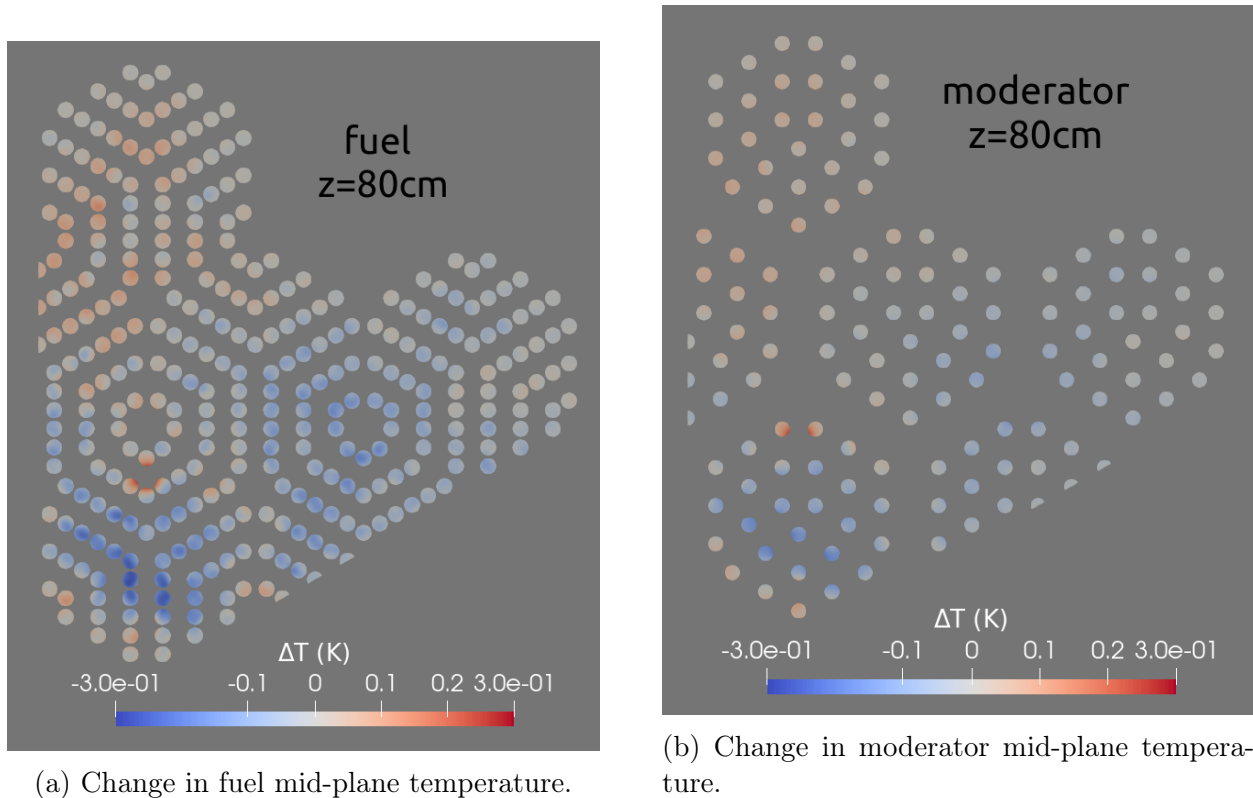




Figure 5.33: Change in MRaD fuel and moderator mid-plane temperature due to hydrogen migration.

By sampling nodal values of the MOOSE output mesh, temperature differences between initial and final iteration solves can be compared. Figure 5.34 displays the result of this process, applied to the hottest fuel pin and moderator rod, with nodes averaged by axial layer. These results show a decrease in temperature in the center of the core, and a very minor increase towards the top of the core. This aligns with the hydrogen shape in Figure 5.29 and power shapes in Figures 5.31 and 5.32. Again, while noticeable, the magnitude of these changes in temperature are well below what would register as a change in material cross-section or threaten safety margins.



(a) Axial nodal-average temperature of MRaD hottest fuel pin from initial and final iteration, along with comparison.



(b) Axial nodal-average temperature of MRaD hottest moderator rod from initial and final iteration, along with comparison.

Figure 5.34: Comparisons of initial and final iteration axial temperature of MRaD hottest fuel pin and moderator rod.

Heat Pipe Failure

In order to simulate the extreme case of long-term heat pipe failure, four heat pipes had their heat removal capability disabled. These heat pipes were located in hottest region of the core, the center of the inner-most assembly. The fuel pins selected and resultant mid-plane temperature profile as a result of these disabled heat pipes can be seen in Figure 5.35.

Note that temperature effects of these heat pipe failures are relatively isolated to where the failures occur. This is a result of one of the benefits of using heat pipes, that their variable heat removal allows the functioning heat pipes to "pick up the slack". Because of this mostly isolated impact, the following analysis of hydrogen migration and its impact is focused on the hottest fuel pin and moderator rod in this region. The global eigenvalue of the system

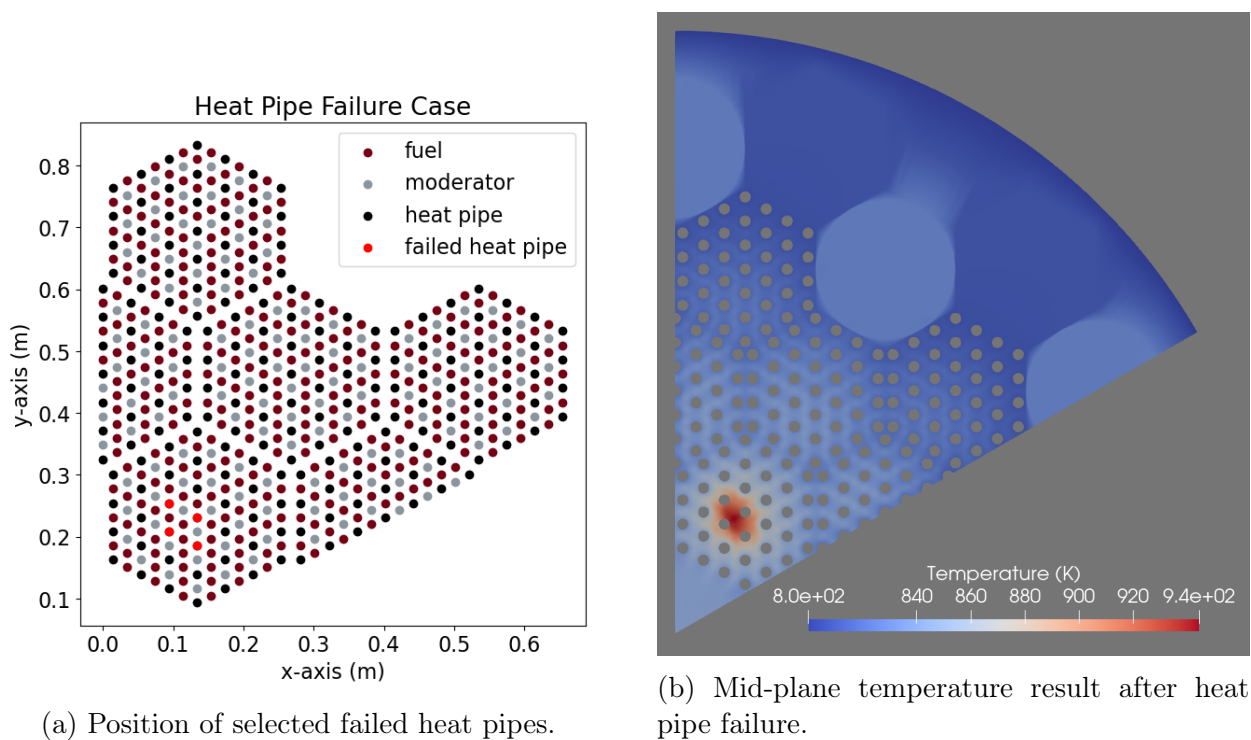


Figure 5.35: Selected heat pipe failure locations and the resulting mid-plane temperature effect on the MRaD core.

saw a decrease of 68 ± 26 pcm. This, coupled with the non-heat pipe failure result, indicates that the hydrogen diffusion most-likely has a slight negative effect on eigenvalue. Given the uncertainties of these values, however, this cannot be quantified with true certainty. No other significant changes to global neutronic properties were identified. It is important to remember that temperature values have not been transferred to OpenMC, so this change in eigenvalue is due to solely the change in hydrogen profiles in the system.

Axial hydrogen profile of the hottest moderator rod, seen in Figure 5.36, show a proportionally significant increase in both the increase of hydrogen at the top of the rod and the decrease of hydrogen in the center of the rod. This leads to the results in Figure 5.37, with the axial power shape change for the hottest fuel pin being exacerbated. That power shape

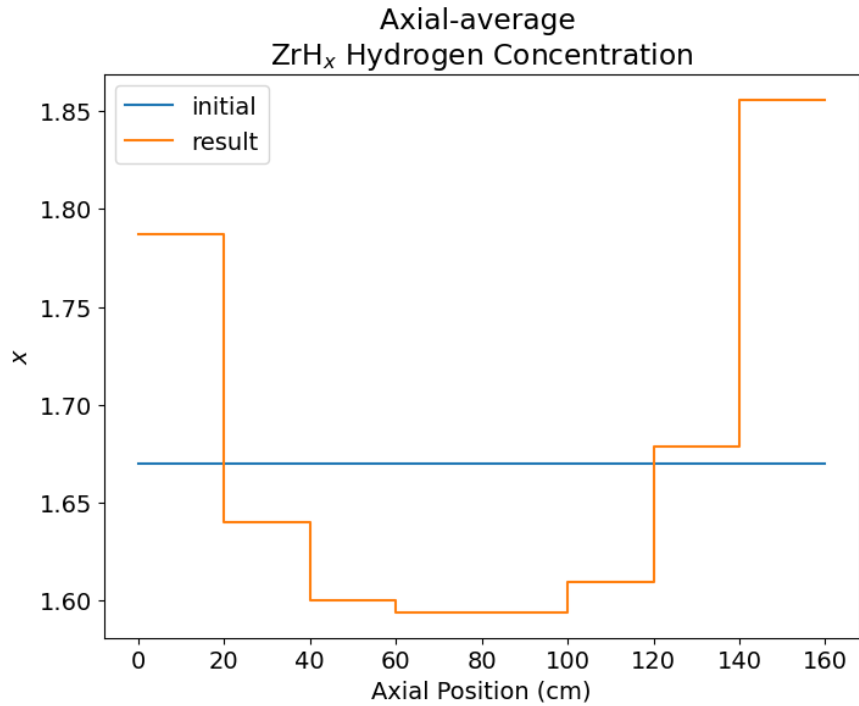




Figure 5.36: Axially-averaged hydrogen profile for MRaD hottest moderator rod after heat pipe failure.

then results in a strong hydrogen migration-derived temperature change, seen in Figure 5.38. From an original maximum deviation of 0.3 Kelvin, now the fuel records a maximum 2 Kelvin decrease because of the hydrogen migration. In this case of increased migration due to heat pipe failure, the fuel pin also sees an increase in temperature at the top of the pin. Though these values are still below what would be a concern for a reactor designer, they are clear enough to note a trend.




(a) Volumetric power versus axial position for fuel pin, in initial and final iteration case with heat pipe failures.




(b) Percent comparison of volumetric power versus axial position for fuel pin in heat pipe failure case.

Figure 5.37: Axial segment-averaged volumetric heating rate of MRaD hottest fuel pin from initial and final iteration with failed heat pipes. Includes the percent change between the two cases.



(a) Change in fuel mid-plane temperature.




(b) Axial fuel pin temperature and temperature change.


Figure 5.38: Change in MRaD hottest fuel pin caused by hydrogen migration with heat pipe failure.

5.4.2 Simba Core

With the Simba core, defined in Section 2.3.2, having a prescribed power of 2 MW thermal for the entire core, the simulated 1/12th case has a power of 166.67 kW. This results in the thermal profiles seen in Figure 5.39 after convergence of the OpenMC-MOOSE coupling. Because of the Simba core design of a gas-filled gap between each assembly, the assembly is somewhat thermally isolated. Similarly, each moderator rod features a thin gas-filled gap where heat transfer is radiative. This causes a diminished heat transfer to the moderator rods, and will affect the results shown further in this section, which I will note.



(a) Isometric view of the Simba Core temperatures after final coupled iteration.



(b) Axial mid-plane slice view of Simba Core temperatures after final coupled iteration.

Figure 5.39: Simba temperature distributions after final coupled iteration.




Figure 5.40: Final hydrogen concentration values in Simba moderator rods, returned by MOOSE.

Hydrogen Results


The converged stoichiometric ratio of hydrogen in ZrH is shown for all moderator rods in Figure 5.40. Note that the maximum ratio of H/Zr is higher here than was seen for the MRaD results, as well as Simba moderator not having the more visible increase in hydrogen at the bottom of the rods. Because hydrogen migration is so closely tied to thermal profile, especially the gradient of temperature in the rod, the effect of the gap around the moderator is likely to be the cause of this difference. The gap reduces thermal conductivity and increases the temperature of the moderator, but also dampens the temperature gradients *within* the rod, especially in the radial and azimuthal spaces.

This can be seen in Figures 5.41 and 5.42, showing the hottest section of the core and the fuel and moderator rods located there. A radial/azimuthal thermal gradient in the moderator is nearly non-existent, and the gradient in H/Zr is muted because of it (note the scale for Figures 5.41b and 5.42b).


Taking the moderator-average H/Zr values at axial intervals, listed in Table 5.5, the axial movement of hydrogen is still quite visible. An average increase of 3.9% in hydrogen content at the top of the moderator rods aligns with Figure 5.40.

Table 5.5: Final axial values for x in ZrH_x , sampled from all Simba moderator rods. The layer under each result is the percent change from original $ZrH_{1.67}$

Axial Location (cm)	Average x	Min. x	Max. x
$z=0$	1.672	1.650	1.714
Δ	0.1%	-1.2%	2.6%
$z=80$	1.657	1.648	1.671
Δ	-0.8%	-1.3%	0.3%
$z=160$	1.736	1.699	1.765
Δ	3.9%	1.8%	5.7%




(a) Fuel and moderator temperature profiles.




(b) Moderator x in ZrH_x profile.

Figure 5.41: Temperature and H/Zr stoichiometric ratio of hottest moderator and fuel, sampled via axial slice at the core mid-plane.

The axial hydrogen shift of the hottest moderator rod is visible in Figure 5.43, showing a peak value of roughly $ZrH_{1.73}$ at the top of the moderator rod. Again, for the Simba core the bottom of the moderator rod doesn't see the net-increase in hydrogen concentration that occurred in MRaD moderator rods.



(a) Fuel and moderator temperature profiles.



(b) Moderator x in ZrH_x profile.

Figure 5.42: Temperature and H/Zr stoichiometric ratio of hottest moderator and fuel, sampled via axial slice at the core top.




Figure 5.43: Axially-averaged hydrogen profile for Simba hottest moderator rod.


Table 5.6: Global neutronic results and differences from initial flat-hydrogen case for Simba core. "AR" is Absorption Rate, "Thermal FF" is the ratio of fission caused by thermal neutrons compared to total fission. Absorption rates are in units of "per source particle".

Quantity	Result Value	Difference	Pct. Dif. (%)
Eigenvalue	1.35536 ± 0.00020	-0.00007 ± 0.00024	-
Thermal FF	$0.40887 \pm 9.0e-5$	$-0.00077 \pm 1.0e-4$	-0.187 ± 0.025
Fuel AR	$0.90357 \pm 1.3e-4$	$0.00020 \pm 1.5e-4$	0.022 ± 0.017
Heat Pipe AR	$0.02425 \pm 1.4e-5$	$-0.00016 \pm 1.7e-5$	-0.070 ± 0.070
Moderator AR	$0.01854 \pm 6.0e-6$	$-0.00006 \pm 0.7e-5$	-0.320 ± 0.040
Reflector AR	$0.02425 \pm 8.0e-6$	$0.00016 \pm 1.0e-5$	0.130 ± 0.080


Neutronic Impacts

For the Simba core, although the hydrogen migration seen is greater than occurred in the MRaD geometry, the global effect to neutronic response is actually lower in magnitude. Some of these global neutronic values are tabulated in Table 5.6. The calculated change in eigenvalue is well beneath the uncertainty for the value, and all changes to material absorption rate are low enough to not be mentioned.

The only value of note is that the thermal fission factor had a noticeable decrease, though still well below a single percent. The decrease of thermal fission factor, a measure of the number of fissions caused by thermal neutrons divided by total fission number, may be a result of less thermalization in the center of the core due to migration of hydrogen vertically. The eigenvalue doesn't see the same impact seen in MRaD results, likely because the Simba core actually relies on more fast fission than thermal fission. The lack of impact to pin power, seen in Figure 5.44, is a continuation of this trend of minimal global effect of the hydrogen migration. Although it seems like the hotter fuel pins in the core see an increase in pin power, and the cooler see a decrease, the standard deviations of each change make a true characterization difficult. Focusing on the hottest moderator rod and fuel pin, a clear difference in the axial profile of heating can be seen. Figure 5.45 shows a 6% increase in heating rate at the top of the fuel pin, and Figure 5.46 shows an equal 6% increase in heating rate at the top of the moderator rod. This is a trend throughout each core geometry studied in this chapter, that hydrogen migration changes the shape of power in the fuel and moderator, but doesn't really diminish or move that power elsewhere.




(a) Simba percent change in pin powers from initial to final iteration.




(b) Standard deviation of percent change in pin powers.

Figure 5.44: Percent change in pin powers of Simba core due to hydrogen migration, and their respective standard deviations.




(a) Volumetric power versus axial position for the fuel pin, in initial and final iteration case.




(b) Percent comparison of volumetric power versus axial position for fuel pin.

Figure 5.45: Axial segment-averaged volumetric heating rate of Simba hottest fuel pin from initial and final iteration, and the percent change between the two.



(a) Volumetric power versus axial position for the moderator rod, in initial and final iteration case.



(b) Percent comparison of volumetric power versus axial position for moderator rod.

Figure 5.46: Axial segment-averaged volumetric heating rate of MRaD hottest moderator rod from initial and final iteration, and the percent change between the two.

Thermal Impacts

The Simba core sees a minor increase in temperature due to hydrogen migration in the areas of hottest temperature, as opposed to the decrease seen in the MRaD core. Figure 5.47 shows this difference at difference axial slices. Figure 5.48 shows a hint for why this is, featuring the same shape of temperature change seen in the MRaD case, but this time with a higher overall magnitude.

I believe that the core sees a more net positive change in temperature because the Simba core fission rate is dominated by fast fissions, as previously noted. Because a lack of hydrogen in the center of the core doesn't strongly impact fission count there, there is almost no decrease in temperature in the center or below for the fuel pin studied (see Figure 5.48a). Conversely, the upper region of the core has less fast fissions due to leakage, so an increase in hydrogen concentration and the resulting increase in thermal fission will register as a larger net increase in temperature.

This explanation is my conjecture, though proving it would require a more in-depth study. Comparing between Simba and MRad has its limitations, as they have different materials, spectra, and geometry.

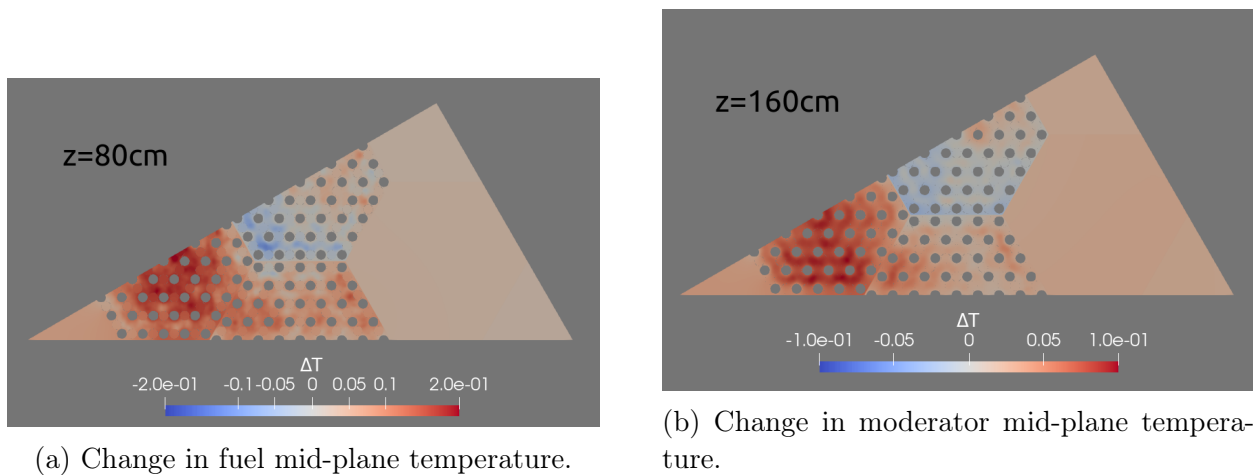




Figure 5.47: Change in fuel and moderator mid-plane temperature due to hydrogen migration in the Simba core.



(a) Axial temperature of fuel pin from initial and final iteration, with comparison.



(b) Axial temperature of moderator rod from initial and final iteration, with comparison.

Figure 5.48: Comparison of initial and final iteration axial nodal-average temperatures of Simba core hottest fuel pin and moderator rod.

Heat Pipe Failure

Like for the previous heat pipe failure case, four heat pipes located in the hottest region of the Simba core were selected to be "failed". These heat pipes and the resulting mid-plane temperature profile of the core can be seen in Figure 5.49. Interestingly, the effects of these

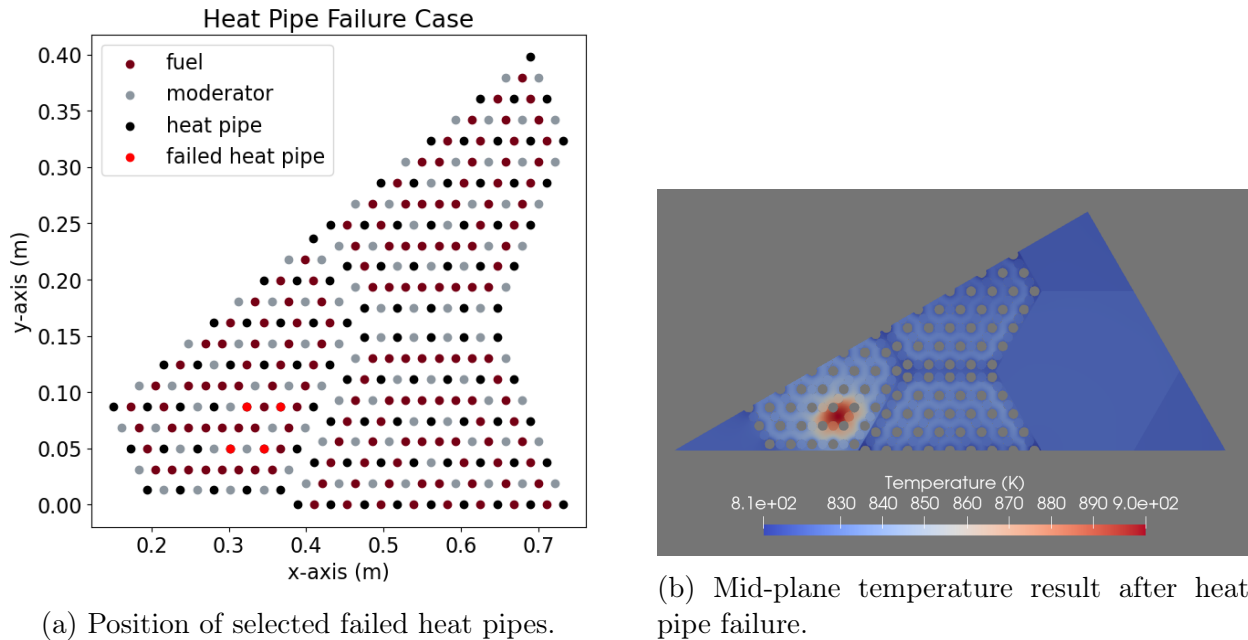


Figure 5.49: Selected heat pipe failure locations and the resulting mid-plane temperature effect on the Simba core.

heat pipe failures are less pronounced on the axial H/Zr profile in the hottest moderator rod, Figure 5.50, than was seen in the MRaD heat pipe failure case. This is likely due to the gap around the moderator insulating some of the effects of these failures.

After convergence, an eigenvalue change of $-69 \pm 23\text{pcm}$ was recorded, still low-enough to not stand out, though the fact that the MRaD case and this Simba case have a negative change is beginning to be a pattern. The heat pipe failure-driven hydrogen migration resulted in a decrease of the thermal fission factor by $-0.2\% \pm 0.025$. The hottest fuel pin sees a more magnified increase and decrease in power, depending on axial position, seen in Figure 5.51. This new heating rate results in the change in temperature slice and axial temperature change comparison in Figure 5.52. This axial comparison shows a decrease of 2 Kelvin in the center and an increase of 3 Kelvin at the top, the largest changes in temperature due to hydrogen migration recorded yet.


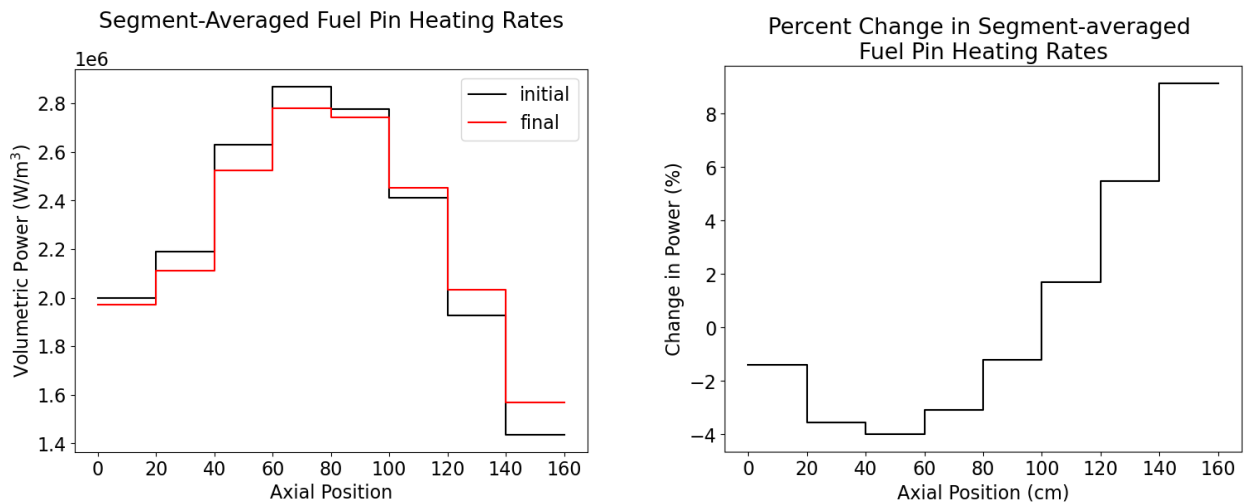



Figure 5.50: Axially-averaged hydrogen profile for Simba hottest moderator rod after heat pipe failures




(a) Volumetric power versus axial position for fuel pin after initial and final case.

(b) Percent comparison of volumetric power versus axial position for fuel pin.

Figure 5.51: Axial segment-averaged volumetric heating rate of Simba hottest fuel pin from initial and final iteration with failed heat pipes. Includes the percent change between the two cases.



(a) Change in fuel mid-plane temperature.

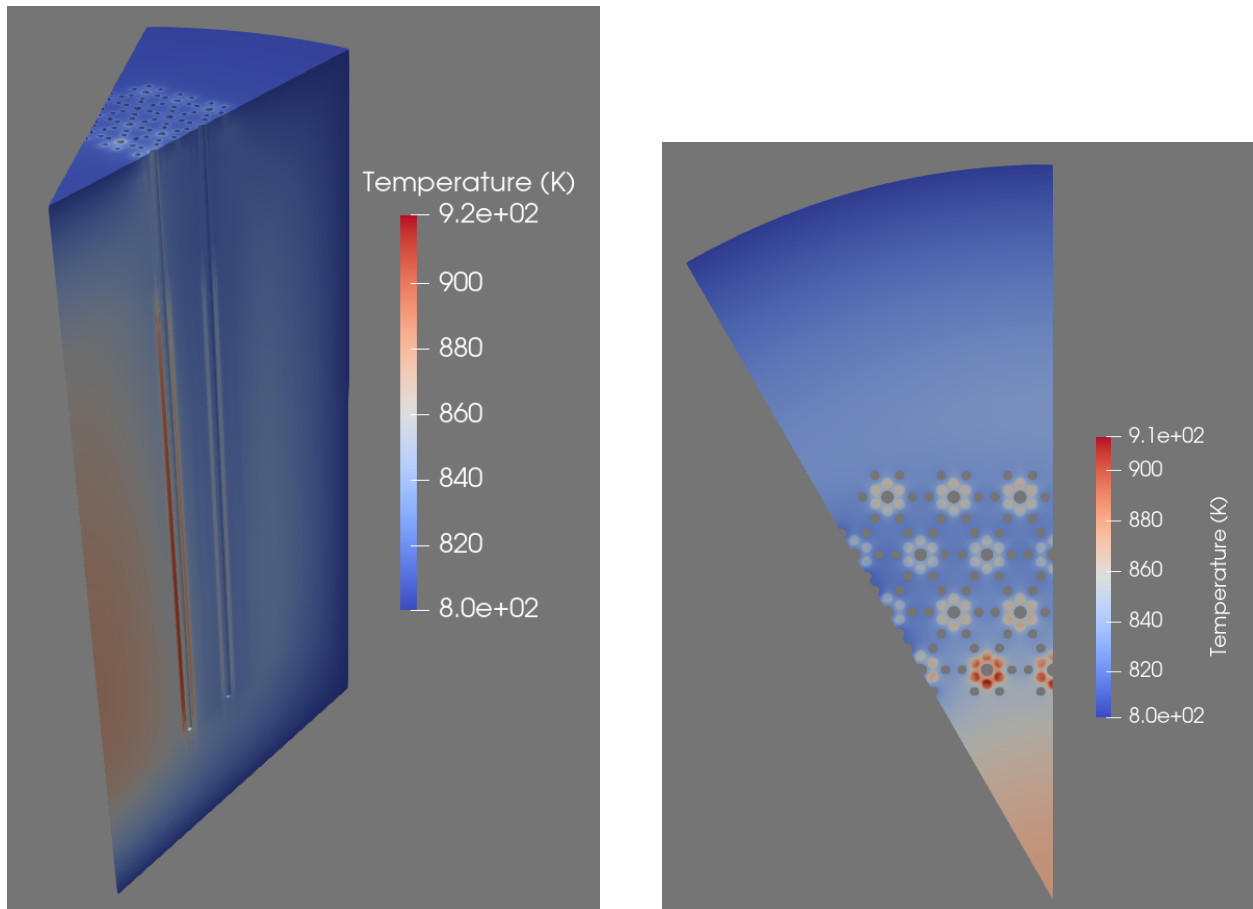


(b) Axial fuel pin temperature and temperature change.

Figure 5.52: Change in Simba hottest fuel pin temperature caused by hydrogen migration with heat pipe failure.

5.4.3 SR Core

Like how the Simba core differed from the MRaD core because of the gap around the moderator rods, the SR core, defined in Section 2.3.3, is unique amongst the geometries studied as cladding for the fuel pins is included, along with the use of a uniform distribution for fuel pin heating, as explained in Section 5.2.5. The moderator rods still have heating and hydrogen profile described via functional expansions. The core has the same power input as the Simba core, 166.67 kW thermal, but because the fewer fuel pins means higher power density and the fuel pins are cladded, the fuel has a much higher temperature than has been seen in the cores so far. The hydrogen migration in the moderator of this core, as well as the neutronic and thermal effects of that hydrogen migration, are presented in the subsections below.



(a) Isometric view of 1/16th of the SR core temperatures after final coupled iteration.

(b) Axial mid-plane slice view of Simba core temperatures after final coupled iteration.

Figure 5.53: SR core temperatures after final coupled iteration.

Table 5.7: Final axial values for x in ZrH_x , sampled from all SR moderator rods. The layer under each result is the percent change from original $\text{ZrH}_{1.67}$

Axial Location (cm)	Average x	Min. x	Max. x
z=0	1.670 0.0%	1.653 -1.0%	1.691 1.2%
z=70	1.669 0.0%	1.646 -1.5%	1.689 1.2%
z=150	1.684 0.9%	1.672 0.1%	1.712 2.5%

Hydrogen Results

The hydrogen migration in the SR core is noticeably more muted, comparatively to the results of the MRaD and Simba cores. When looking at Figure 5.54, an isometric view of all core moderator rods and their hydrogen ratio, and Table 5.7, an axial average of these moderator rods' ratios, the axial migration of hydrogen can be seen to be less pronounced compared to MRaD and Simba results.

A main cause of this is due to the geometry of the core, specifically being where the peak power occurs in the core. So far in the results for MRaD and Simba, it has been clear that hydrogen migration occurs where thermal gradients are large, and with these cores being all fresh fuel with no burnable absorbers, power and therefore temperature gradient are highly focused in one region of the core. In the case of the MRaD and Simba geometry, the hottest fuel pin was located on the inside of an assembly, generally surrounded by multiple heat pipes. The SR core geometry doesn't have conventional assemblies, and instead the hottest fuel pin is found at the inner edge of the core, relatively isolated. There are only roughly 2 moderator rods that directly "feel" the peak pin's high temperature in the SR core geometry. Additionally, the lack of an axial shape in the fuel pin power causes a notable flattening of thermal gradient in both the fuel and moderator. This less pronounced gradient further decreases the magnitude of hydrogen migration witnessed.

This effect can be seen by comparing the isometric H/Zr figures across core designs (Figures 5.54 versus 5.40 and 5.26). The SR core result has roughly 2-3 moderator rods that feature core-maximum hydrogen content, while both MRaD and Simba have 8-10. By averaging across these moderators, it makes sense that the averages in Table 5.7 will be lower than any other core studied.

The effect of the fuel cladding is apparent in Figures 5.55 and 5.56, axial slices of the hottest region of the core featuring moderator and fuel temperature and moderator hydrogen stoichiometric ratio. The fuel temperature rises but moderator temperature remains relatively low, only seeing a roughly 10 degree Kelvin difference between mid-plane and top.

Even when looking at the hottest moderator rod in the core, the axial profile of hydrogen (Figure 5.57) is less pronounced than in the case of the previous geometries. It is similar in shape to the MRaD results of Figure 5.29, but lacks a more clearly defined migration towards the top of the reactor.






Figure 5.54: Final hydrogen concentration values in SR moderator rods, returned by MOOSE.




(a) Fuel and moderator temperature profiles.

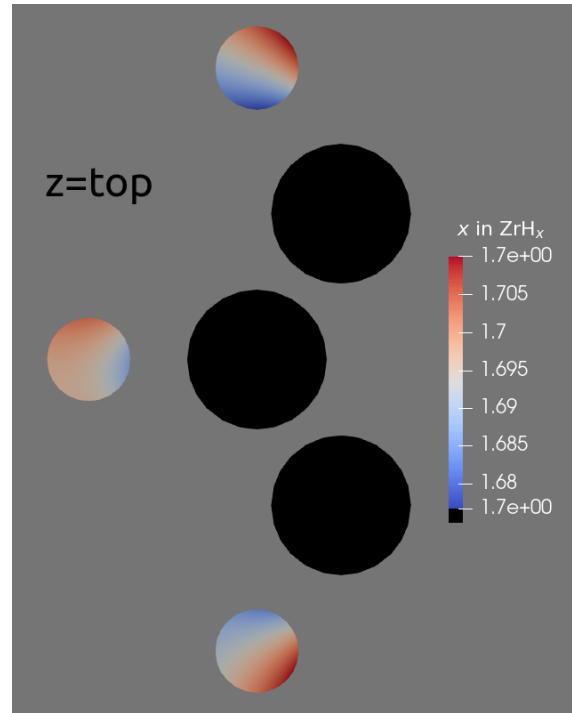


(b) Moderator x in ZrH_x profile.

Figure 5.55: Temperature and H/Zr stoichiometric ratio of hottest moderator and fuel, sampled via axial slice at the core mid-plane. Moderator rods are the smaller pins on the left side.



(a) Fuel and moderator temperature profiles.



(b) Moderator x in ZrH_x profile.

Figure 5.56: Temperature and H/Zr stoichiometric ratio of hottest moderator and fuel, sampled via axial slice at the core top. Moderator rods are the smaller pins on the left side.




Figure 5.57: Axially-averaged hydrogen profile for SR hottest moderator rod.

Table 5.8: Global neutronic results and differences from initial flat-hydrogen case for SR core. "AR" is Absorption Rate, "Thermal FF" is the ratio of fissions caused by thermal neutrons compared to total fissions. Absorption rates are in units of "per source particle".

Quantity	Result Value	Difference	Pct. Dif. (%)
Eigenvalue	1.01371 ± 0.00019	0.00048 ± 0.00027	-
Thermal FF	$0.79193 \pm 5.0e-5$	$0.00010 \pm 8.0e-5$	0.012 ± 0.010
Fuel AR	$0.73482 \pm 1.1e-4$	$0.00027 \pm 1.6e-4$	0.036 ± 0.022
Heat Pipe AR	$0.07945 \pm 2.5e-5$	$0.00002 \pm 4.0e-5$	0.020 ± 0.05
Moderator AR	$0.02427 \pm 6.0e-6$	$0.00001 \pm 8.0e-6$	0.043 ± 0.034
Reflector AR	$0.04197 \pm 2.4e-5$	$-0.00005 \pm 3.3e-5$	-0.130 ± 0.080

Neutronic Impacts

Global neutronic effects of the hydrogen migration follow the same pattern as in the case of MRaD and Simba, with one minor difference. Eigenvalue shows a potential positive change on the order of 40-50 pcm, but the uncertainties are large enough to make a drawing a firm conclusion impossible. Likewise, pin power changes, seen in Figure 5.58, remain mostly without any apparent trend besides a minor movement outwards in power, and the standard deviations for these values nearly equal the value in many cases.

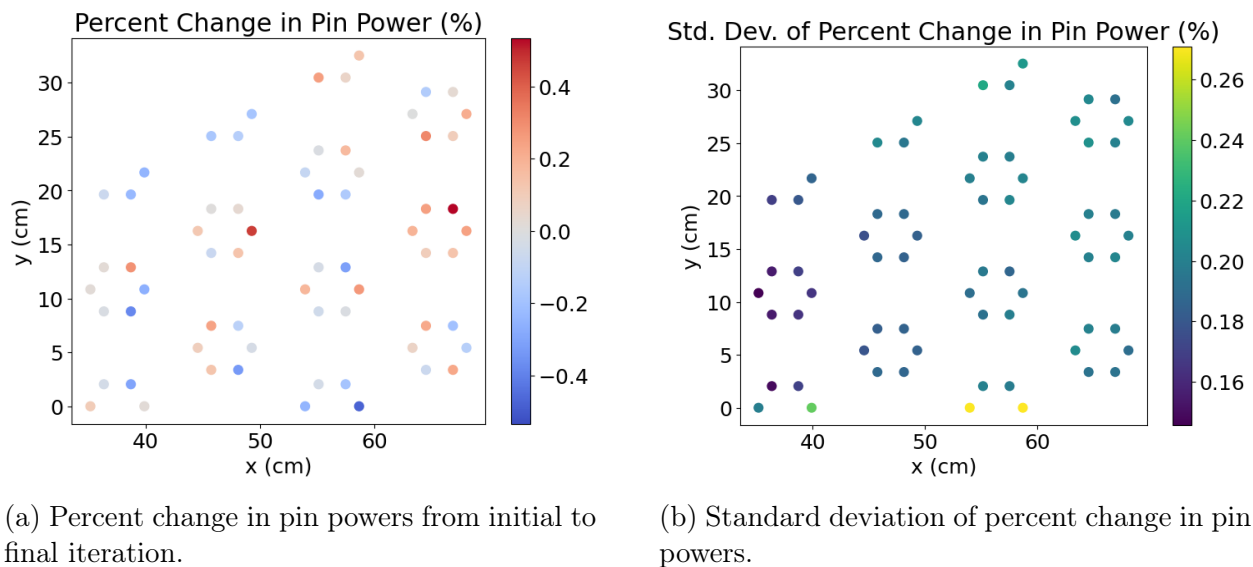

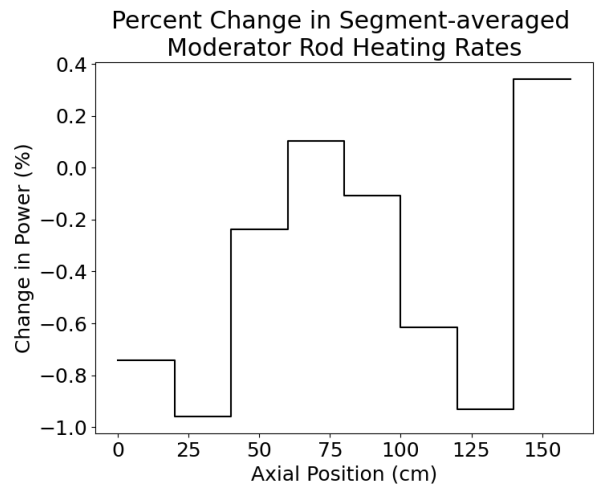


Figure 5.58: Percent change in pin powers of SR core due to hydrogen migration and their respective standard deviations.

Looking at the results in Figure 5.59, the moderator power profile sees an odd "w" like change in axial power profile due to hydrogen migration. This is different from the MRaD and Simba results, and is almost certainly due to the lack of a true axial peak in fuel power. With the fuel power being evenly distributed resulting in a less shaped temperature profile for the moderator, the hydrogen only strongly shows movement at axial ends of the moderator rods.



(a) Volumetric power versus axial position for the moderator rod, in initial and final iteration case.




(b) Percent comparison of volumetric power versus axial position for moderator rod.

Figure 5.59: Axial segment-averaged volumetric heating rate of SR hottest fuel pin from initial and final iteration, including the percent change between the two.


Thermal Impacts

Given the fact that the hydrogen migration and neutronic impacts for the SR core have been notably lesser, the thermal impacts from hydrogen diffusion should be equally as muted. This is the case for axial temperature differences for the hottest moderator and fuel, seen in Figure 5.61. Temperature differences indicate a decrease in the lower portion of the core and an increase in the upper, similar to results seen for the previous reactors studied, however the magnitude of these changes are negligible.

A mid-plane comparison of temperature difference caused by hydrogen migration can be seen for fuel and moderator in Figure 5.60. These results indicate a slight decrease in temperature in the inner of the core and increase in the outer, but the temperature changes are minor, and without a true axial shape to fuel power they are questionable at best.




(a) Change in fuel mid-plane temperature.




(b) Change in moderator mid-plane temperature.

Figure 5.60: Change in fuel and moderator mid-plane temperature due to hydrogen migration in the SR core.



(a) Axial temperature of hottest fuel pin from initial and final iteration, with comparison.



(b) Axial temperature of hottest moderator rod from initial and final iteration, with comparison.

Figure 5.61: Comparison of initial and final iteration axial nodal-average temperatures of SR core hottest fuel pin and moderator rod.

Heat Pipe Failure

For the SR core, three heat pipes were selected to "fail", located in the hottest region of the core. These pipes can be seen in Figure 5.62, along with the mid-plane temperature profile of the core after coupling iteration. The global effect of hydrogen migration after heat pipe failure remains below the inherent uncertainties, including for the change in eigenvalue. The increased axial migration of hydrogen in the hottest moderator rod, seen in Figure

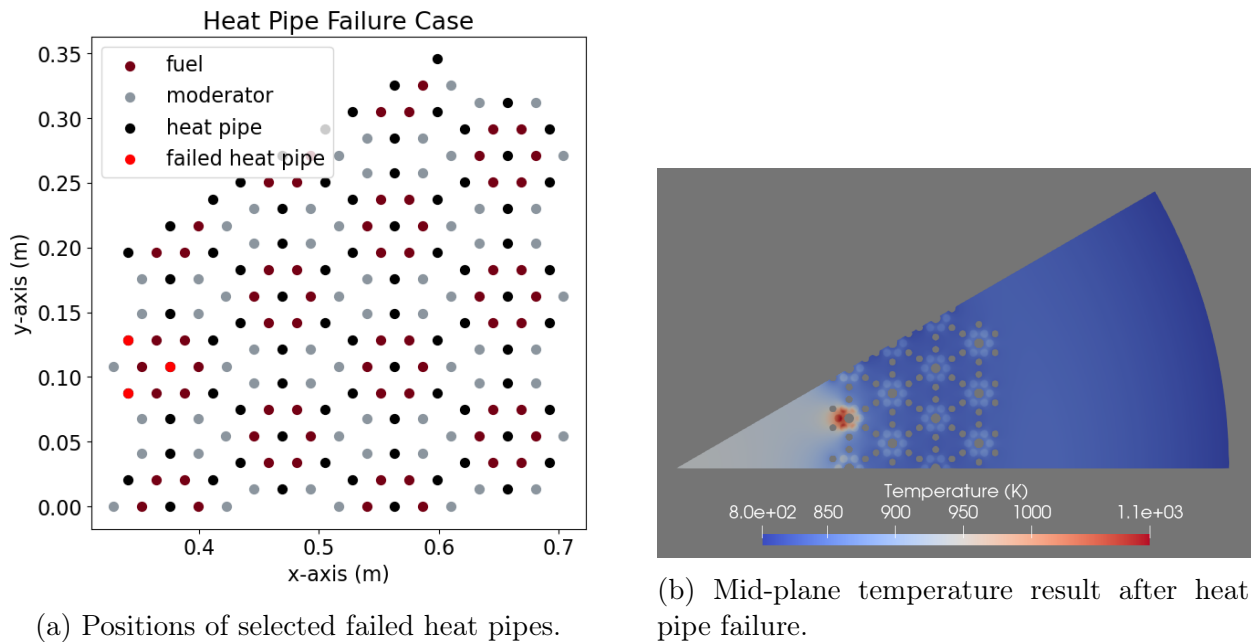


Figure 5.62: Selected heat pip failure locations and the resulting mid-plane temperature effect on the SR core.

5.63, causes the hottest moderator rod axial heating rates, seen in Figure 5.64 to respond accordingly. The previously mentioned "w" shape remains, again responding to the lack of a strong axial shape from the fuel pins. Figure 5.65 show the mid-plane temperature change due to hydrogen migration for all fuels. Comparing this figure to the heat pipe failure cases for MRad (Fig. 5.38 and Simba (Fig. 5.52), a similar trend of a decrease in mid-plane temperature in hot regions is observed. The axial profile of the hottest fuel pin temperature, seen in Figure 5.65b, shows a very noisy profile, only really indicating a net decrease in temperature, though minor enough to not register on any instrumentation.






Figure 5.63: Axially-averaged hydrogen profile for SR hottest moderator rod after heat pipe failures.




(a) Volumetric power versus axial position for moderator rod after initial and final case.

(b) Percent comparison of volumetric power versus axial position for moderator rod.

Figure 5.64: Axial segment-averaged volumetric heating rate of SR hottest moderator rod from initial and final iteration with heat pipes. Includes the percent change between the two cases.



(a) Change in fuel mid-plane temperature.



(b) Axial fuel pin temperature and temperature change.

Figure 5.65: Change in SR hottest fuel pin temperature caused by hydrogen migration with heat pipe failure.

5.5 Conclusion

This chapter revolved around the simulation of hydrogen migration and quantifying the effect of including hydrogen migration with neutronic-thermal simulation of three heat pipe microreactors. Coupling was performed in tight Picard scheme by alternating OpenMC and MOOSE solves, with OpenMC generating volumetric power for MOOSE, and MOOSE generating hydrogen concentrations for OpenMC. Both sets of information were transferred as expansion coefficients for a cylindrical basis set made by convoluting Legendre and Zernike polynomials.

The results of each coupled simulation showed that axial hydrogen migration was more impactful than radial or azimuthal migration, likely due to the radius of the moderator rods being lower than the mean free path of a neutron. That migration became far more active as temperatures increased beyond the 830 Kelvin area. The effect of migration was limited to the area where the migration occurred, with power and temperature seeing a decrease in the axial center and an increase at the top of the cores. These changes were relatively minor, however, with the strongest change in temperature due to hydrogen migration being on the order of a couple of degrees Kelvin. Globally, the effect was within uncertainties, although some results suggest that large axial migration can cause a decrease in thermal fission factor as moderation drops in the center of the core.

Heat pipe failures exacerbated these trends, across the board. Although the still-functioning heat pipes kept core temperatures from rising too much on average, in the regions with failed heat pipes, hydrogen migration resulted in further decreased center temperatures and increased axial-top temperatures.

In general, this work suggests that the migration of hydrogen in the zirconium hydride moderator results in a decrease in power and temperature in the axial center of the core, and an increase in power and temperature at the axial top. The magnitudes of these effects are likely minor, potentially below what would even be detectable by instrumentation. In cases of heat pipe failure or general temperature spikes, hydrogen will further migrate and increase central temperature dampening, a potentially beneficial behavior for reactor accident scenarios.

The different reactor designs studied are not comparable in a 1-to-1 sense, but some takeaways are possible. Based on the Simba geometry, gaps around the moderator rods, as exist in real-life reactors, will insulate the rod from some strong thermal gradients, like in the radial and azimuthal spaces. It also, however, reduces heat transfer out of the rod, causing higher temperatures and potentially more migration. Based on the SR geometry, cladding on the fuel may reduce some of the thermal gradient seen by the moderator, as the increased thermal resistivity diffuses some of the temperature gradient native to the fuel pin. Additionally, the SR results show a very strong need for axial power resolution in the core fuel pins in order to accurately capture the effects of hydrogen diffusion. Without this axial profile, results were overall muted in magnitude and lacked significance compared to the MRaD and Simba results. I recommend that similar work of this nature with annular geometry include the use of a basis set that is orthogonal over the annular geometry, as explained in Section 5.2.5.

This study is not without its weaknesses, however, and the results should be taken as

approximations of real life behavior. For example, all fuel material in this study was fresh with no burnable absorbers, causing all cores to have very peaked power. Reactor designers typically try to flatten this shape as much as they can, which could potentially lead to a lower thermal gradient and therefor less hydrogen migration in the moderator. Additionally, the application of δ -phase ZrH behavior to all temperatures and concentrations hides the impacts that may exist as the ZrH moderator transitions into ϵ -phase. This phase is characterized with drastically increased diffusion, and would potentially change the migration characteristics seen in this work.

Finally, uncertainty for expansion coefficients, specifically for those describing heating rates, were neglected in their application for this work. The uncertainties for the power shape expansion coefficients are tabulated in the OpenMC output, but translating these uncertainties to an uncertainty in MOOSE reported temperatures is non trivial. I believe it is possible, but was not within my abilities for this work. These uncertainties do exist, however, and quantifying their value for axial heating profiles and on the resulting temperature changes may affect the conclusion drawn about this study.

An area of work that may increase the impactfulness of hydrogen migration is during transients. In the case of a fast transient, such as a control rod ejection, a 4% hydrogen gradient may cause much larger temperature variations than this study saw. Transients are difficult to represent using Monte Carlo particle transport, however, so certain methods used in this study may not be optimal for a study on the topic of fast transients.

When comparing these results to the results that exist in literature, there is clear, strong agreement. Because the studies that currently exist only really characterize the hydrogen shift and its impact on global core characteristics, only these facets can be compared [33], [37], [108]. To that regard, they line up well, as these contemporary studies show the same movement of hydrogen to the cooler top of the core, and a effect on eigenvalue in the order of less than 50 pcm.

Chapter 6

Conclusions

Based on the results of neutronic-thermal coupling with the separate addition of thermal expansion and hydrogen migration, some general conclusions can be drawn about their respective total impact.

For thermal expansion, inclusion in the simulation of feedback shows a clear decrease in leakage and accompanying increase in core eigenvalue. Beyond that, however, impacts are low enough to be difficult to discern beyond their uncertainties due to stochastic randomness. Nominal pin powers do not change, and no clear change in temperature is witnessed. A study of the axial or radial effects may show differences in thermal expansion's inclusion, but this study's use of DAGMC restricted heating rates to being cell-average only. Heat pipe failure did not affect thermal expansion rates to a degree large-enough to impact these results.

For hydrogen diffusion, inclusion of hydrogen migration in neutronic-thermal simulation resulted in essentially no global effects, with nominal pin powers remaining constant. The axial profile of these powers did, however, see distinct changes, with power increasing towards the upper cooler region of the core and decreasing in the hotter lower regions. These power shifts resulted in very minor temperature changes, below what would be noticeable in operation. During heat pipe failure, however, the axial shifts in power of the fuel pins closest to the failures were exacerbated. This caused a decrease of 2 Kelvin and increase of 1.5 Kelvin recorded at different axial locations in the fuel. While still below what would be a concern for safety margins, this does prove large enough to establish a clear trend.

Beyond the nominal value of having these results to inform potential design choices or to use for comparison of the next study on these topics, these results can also be used to give an opinion on their usefulness in a coupled simulation-sense. This is important to note because coupling different physics systems incurs a cost. Often that cost is computational, as the more complex the problem the more core-hours are required to solve that problem. Table 6.1 shows the difference in simulation rate of OpenMC depending on what method is applied. Notice that utilising DAGMC with OpenMC causes a significant penalty in simulation rate, while CVMT has relatively little impact. Note that each value is the result of a single run, so some variation is expected due to stochasticity.

However that cost is also in simulation quality. For example, in order to avoid simulation failure during mesh deformation, the MRaD core mesh had to be altered and simplified in different ways. This allows the coupling of thermal expansion, but at a cost to the quality of the thermal and neutronic results due to the lower quality geometry. Evaluating whether

Table 6.1: Comparison of OpenMC particle simulation rates versus method and geometry. Rates are particles per second.

Core Geometry	Method	Inactive Rate	Active Rate
MRaD	Default	8162.8	3972.8
	CVMT	7894.3	3308.03
	DAGMC	1926.4	1895.8
Simba	Default	17982.9	6545.8
	CVMT	19690.1	5123.4
	DAGMC	-	-
SR	Default	8642.6	4671.1
	CVMT	12319	5408.5
	DAGMC	2620.2	2583.6

these costs are worth paying in order to have the whole coupled simulation is an important takeaway of this work.

Based on the results of the thermal expansion aspect of this work, there is little reason to couple thermal expansion to neutronic-thermal simulation. Thermal expansion results in a global decrease in leakage and couple hundred pcm increase to eigenvalue, and while that change is important to quantify for predicting criticality, thermal expansion does not appear to cause any distinct shifts in power magnitude or location. The complexity of simulation solid mechanics is also very high, and mesh deformation requires careful design of the core mesh. In my opinion, thermal expansion, and the other very important mechanical effects in the core such as contact forces and yield prediction, deserve their own detailed study, not as an addition to neutronic-thermal simulation.

There is the potential to use the general method applied for other deformation topics that may be more impactful in a neutronic-thermal simulation context, however. In particular, deformation that is asymmetric in the core may produce interesting results in coupled neutronic-thermal simulation. Asymmetric changes in reactor geometry will drive asymmetric changes in flux, meaning potential shifts in power. These changes may also arise from stronger thermal gradients than are seen in this work, along with different structural deformation regimes.

Hydrogen diffusion, on the other hand, has a stronger case to be made as an important addition to neutronic-thermal coupling. The diffusion of hydrogen on the MOOSE side is easy to compute, and the only cost on the neutronic side is the use of Continually Varying Material Tracking. For analyzing regular, steady-state operation of the core, representing the diffusion of hydrogen will not make a large impact to neutronic or thermal results. However, given the heat pipe failure study results, hydrogen migration is likely to cause notable impacts during failure scenarios or during cases of strong thermal changes, such as during fast transients where high temperatures will accelerate hydrogen migration. Additionally, no leakage or phase change was simulated in this study, and both of these would likely increase hydrogen's diffusion behavior and reduction in concentration in the core. In my opinion, hydrogen diffusion should be used in coupled simulation with the neutronic-thermal solve when evaluating extreme cases such as failure in reactors that feature zirconium hydride.

Finally, a note should be made that thermal expansion and hydrogen diffusion have the capability to be solved in the same simulation. Thermal expansion requires the use of DAGMC to transfer the deformed mesh for particle transport, but nothing specific about the use of functional expansions or CVMT would inhibit their use on this DAGMC-based geometry. This would require, however, a more detailed generation of normalization parameters for functional expansion usage. For example, if a moderator rod expands axially by 3 centimeters, the length parameter used to standardize the basis set needs to reflect this. A more complicated case is how to handle the radial parameter. Section 5.2.4 details the complexity in using the cylindrical basis set as a representation for a tessellated object. It is easy to imagine the difficulty in doing the same if that object had different amounts of radial swelling along its axial profile. The easy solution to this would be to inhibit deformation of the moderator rods and fuel pins, but then the thermal expansion simulation would hardly be considered realistic. It is for these reasons that this thesis approaches the simulation of each separately.

This thesis set out to perform neutronic-thermal simulation of three heat pipe microreactors with the addition of hydrogen diffusion and thermal expansion. This was accomplished, along with heat pipe failure studies for each model studied. There is significant work that can potentially be done to extend the results shown in this thesis, as further complexity can be added to both the thermal expansion and hydrogen diffusion physics simulated.

References

- [1] *Microreactors*, <https://inl.gov/trending-topics/microreactors/>, Accessed: 2024-06.
- [2] *eVinci Microreactor*, <https://westinghousenuclear.com/energy-systems/evinci-microreactor>, Accessed: 2024-06.
- [3] J. A. Christensen, W. L. Moe, and P. Jordan, “Determining the appropriate emergency planning attributes for microreactors,” Idaho National Lab.(INL), Idaho Falls, ID (United States), Tech. Rep., 2020.
- [4] *Micro-Reactor*, <https://www.rolls-royce.com/innovation/novel-nuclear/micro-reactor.aspx>, Accessed: 2024-06.
- [5] *Terrestrial Micro RX*, <https://www.bwxt.com/what-we-do/advanced-technologies/terrestrial-micro-rx>, Accessed: 2024-07.
- [6] K. A. G. Black D. Shropshire and A. van Heek, “Prospects for nuclear microreactors: A review of the technology, economics, and regulatory considerations,” *Nuclear Technology*, vol. 209, no. sup1, S1–S20, 2023. DOI: [10.1080/00295450.2022.2118626](https://doi.org/10.1080/00295450.2022.2118626). eprint: <https://doi.org/10.1080/00295450.2022.2118626>. URL: <https://doi.org/10.1080/00295450.2022.2118626>.
- [7] “Microreactors In Alaska: Use Case Analysis,” The University of Alaska, Tech. Rep., 2020.
- [8] R. Macdonald and J. E. Parsons, *The Value of Nuclear Microreactors in Providing Heat and Electricity to Alaskan Communities*. JSTOR, 2021.
- [9] J. McCallum, “Nuclear Costs in Context,” The Nuclear Energy Institute, Tech. Rep., 2023.
- [10] P. R. M. David I. Poston Marc A. Gibson and R. G. Sanchez, “Results of the krusty warm critical experiments,” *Nuclear Technology*, vol. 206, no. sup1, S78–S88, 2020. DOI: [10.1080/00295450.2020.1727287](https://doi.org/10.1080/00295450.2020.1727287). eprint: <https://doi.org/10.1080/00295450.2020.1727287>. URL: <https://doi.org/10.1080/00295450.2020.1727287>.
- [11] T. J. Tautges, P. P. H. Wilson, J. A. Kraftcheck, B. M. Smith, and D. L. Henderson, “Acceleration techniques for direct use of CAD-based geometries in Monte Carlo radiation transport,” in *M&C 2009 Saratoga Springs NY*, 2009.

- [12] P. K. Romano, N. E. Horelik, B. R. Herman, A. G. Nelson, B. Forget, and K. Smith, “Openmc: A state-of-the-art monte carlo code for research and development,” *Annals of Nuclear Energy*, vol. 82, pp. 90–97, 2015, Joint International Conference on Supercomputing in Nuclear Applications and Monte Carlo 2013, SNA + MC 2013., ISSN: 0306-4549. DOI: <https://doi.org/10.1016/j.anucene.2014.07.048>. URL: <https://www.sciencedirect.com/science/article/pii/S030645491400379X>.
- [13] A. Davis, P. Shriwise, and X. Zhang, “Dag-openmc: Cad-based geometry in openmc,” Jan. 2020, pp. 395–398. DOI: [10.13182/T122-32104](https://doi.org/10.13182/T122-32104).
- [14] G. Giudicelli, A. Lindsay, L. Harbour, *et al.*, “3.0 - MOOSE: Enabling massively parallel multiphysics simulations,” *SoftwareX*, vol. 26, p. 101690, 2024, ISSN: 2352-7110. DOI: <https://doi.org/10.1016/j.softx.2024.101690>. URL: <https://www.sciencedirect.com/science/article/pii/S235271102400061X>.
- [15] M. N. Cinbiz, C. N. Taylor, E. Luther, H. Trelue, and J. Jackson, “Considerations for hydride moderator readiness in microreactors,” *Nuclear Technology*, vol. 209, no. sup1, S136–S145, 2023. DOI: [10.1080/00295450.2022.2121583](https://doi.org/10.1080/00295450.2022.2121583). eprint: <https://doi.org/10.1080/00295450.2022.2121583>. URL: <https://doi.org/10.1080/00295450.2022.2121583>.
- [16] D. Olander, E. Greenspan, H. D. Garkisch, and B. Petrovic, “Uranium–zirconium hydride fuel properties,” *Nuclear Engineering and Design*, vol. 239, no. 8, pp. 1406–1424, 2009, Hydride Fueled LWRs, ISSN: 0029-5493. DOI: <https://doi.org/10.1016/j.nucengdes.2009.04.001>. URL: <https://www.sciencedirect.com/science/article/pii/S0029549309001745>.
- [17] J. D. Galloway, V. J. Lawdensky, D. I. Poston, H. R. Trelue, and M. E. Blood, “Effects of heat pipe failures in microreactors,” Los Alamos National Lab.(LANL), Los Alamos, NM (United States), Tech. Rep., 2020.
- [18] *Standard Heat Pipes*, <https://myheatsinks.com/heat-pipe-solutions/standard-heat-pipes/>, Accessed: 2024-08.
- [19] K. A. Woloshun, M. A. Merrigan, and E. D. Best, “Htpipe: A steady-state heat pipe analysis program: A user’s manual,” Nov. 1988. URL: <https://www.osti.gov/biblio/6726117>.
- [20] J. E. Hansel, R. A. Berry, D. Andrs, M. S. Kunick, and R. C. Martineau, “Sockeye: A one-dimensional, two-phase, compressible flow heat pipe application,” *Nuclear Technology*, vol. 207, no. 7, pp. 1096–1117, 2021. DOI: [10.1080/00295450.2020.1861879](https://doi.org/10.1080/00295450.2020.1861879). eprint: <https://doi.org/10.1080/00295450.2020.1861879>. URL: <https://doi.org/10.1080/00295450.2020.1861879>.
- [21] C. Fletcher and R. Schultz, “Relap5/mod3 code manual,” Nuclear Regulatory Commission, Washington, DC (United States), Tech. Rep., 1992.
- [22] M. Stålek and C. Demazière, “Development and validation of a cross-section interface for parcs,” *Annals of Nuclear Energy*, vol. 35, no. 12, pp. 2397–2409, 2008.
- [23] S. Nikonov, I. Pasichnyk, and K. Velkov, “Application of system code athlet for sub-channel analysis,” in *Proceedings of the 27th AER Symposium on VVER Reactor Physics and Reactor Safety, Germany, Munich*, 2017.

- [24] U. Grundmann, S. Mittag, and U. Rohde, “Dyn3d2000/m1 for the calculation of reactivity initiated transients in lwr with hexagonal and quadratic fuel elements—code manual and input data description for release,” *Inc., September*, 2001.
- [25] F. S. D’Auria, A. BOUSBIA SALAH, G. Galassi, J. Vedovi, F. Reventós, A. Cuadra, J. Gago, A. Sjoberg, M. Yitbarek, O. Sandervag, *et al.*, “Neutronics/thermal-hydraulics coupling in lwr technology—crissue-s wp3: Achievements and recommendations report,” 2004.
- [26] V. Kyriakopoulos, M. E. Tano, and J. C. Ragusa, “Development of a single-phase, transient, subchannel code, within the moose multi-physics computational framework,” *Energies*, vol. 15, no. 11, p. 3948, 2022.
- [27] C. Lee, Y. S. Jung, H. Park, E. R. Shemon, J. Ortensi, Y. Wang, V. M. Labouré, and Z. M. Prince, “Griffin software development plan,” Idaho National Lab.(INL), Idaho Falls, ID (United States); Argonne National . . . , Tech. Rep., 2021.
- [28] A. Novak, P. Shriwise, P. Romano, R. Rahaman, E. Merzari, and D. Gaston, “Coupled monte carlo transport and conjugate heat transfer for wire-wrapped bundles within the moose framework,” *Nuclear Science and Engineering*, vol. 197, no. 10, pp. 2561–2584, 2023.
- [29] A. Novak, P. Romano, B. Wendt, R. Rahaman, E. Merzari, L. Kerby, C. Permann, R. Martineau, and R. N. Slaybaugh, “Preliminary coupling of openmc and nek5000 within the moose framework,” in *Proceedings of PHYSOR*, vol. 4, 2018.
- [30] M. S. Ellis, “Methods for Including Multiphysics Feedback in Monte Carlo Reactor Physics Calculations,” Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA, USA, 2017. URL: <http://hdl.handle.net/1721.1/112381>.
- [31] S. Harper, “Tally derivative based surrogate models for faster monte carlo multi-physics,” Ph.D. dissertation, Massachusetts Institute of Technology, 2020.
- [32] N. E. Stauff, K. Mo, Y. Cao, J. W. Thomas, Y. Miao, C. Lee, C. Matthews, and B. Feng, “Preliminary applications of neams codes for multiphysics modeling of a heat pipe microreactor,” *Proceedings of the American Nuclear Society Annual 2021 Meeting*, 2021.
- [33] N. Stauff, K. Mo, Y. Cao, J. Thomas, Y. Miao, L. Zou, D. Nunez, E. Shemon, B. Feng, and K. Ni, “Detailed analyses of a triso-fueled microreactor: Modeling of a micro-reactor system using neams tools,” Argonne National Lab.(ANL), Argonne, IL (United States), Tech. Rep., 2021.
- [34] Y.-H. Kim, K.-S. Kim, and J.-M. Noh, “Reactivity-equivalent physical transformation for homogenization of double-heterogeneous fuels,” in *Proceedings of the Korean Nuclear Society Conference*, Korean Nuclear Society, 2005, pp. 115–116.
- [35] A. P. Shivprasad, T. E. Cutler, J. K. Jewell, V. K. Mehta, S. W. Paisner, C. A. Taylor, C. Taylor, H. R. Trelue, D. Wootan, and E. P. Luther, “Advanced moderator material handbook,” Los Alamos National Laboratory (LANL), Los Alamos, NM (United States), Tech. Rep., 2020.

- [36] R. J. McConn, C. J. Gesh, R. T. Pagh, R. A. Rucker, and R. Williams III, “Compendium of material composition data for radiation transport modeling,” Pacific Northwest National Lab.(PNNL), Richland, WA (United States), Tech. Rep., 2011.
- [37] S. Terlizzi and V. Labouré, “Asymptotic hydrogen redistribution analysis in yttrium-hydride-moderated heat-pipe-cooled microreactors using direwolf,” *Annals of Nuclear Energy*, vol. 186, p. 109735, 2023.
- [38] E. Gateau, N. Todreas, and J. Buongiorno, “Consequence-based security for microreactors,” *Nuclear Engineering and Technology*, vol. 56, no. 3, pp. 1108–1115, 2024, ISSN: 1738-5733. DOI: <https://doi.org/10.1016/j.net.2024.01.053>. URL: <https://www.sciencedirect.com/science/article/pii/S173857332400055X>.
- [39] C. S. Crawford, “Conceptual design of a nuclear microreactor transportation cask,” Ph.D. dissertation, Massachusetts Institute of Technology, 2024.
- [40] E. Shemon, Y. Miao, S. Kumar, K. Mo, Y. S. Jung, A. Oaks, S. Richards, G. Giudicelli, L. Harbour, and R. Stogner, “Moose reactor module: An open-source capability for meshing nuclear reactor geometries,” *Nuclear Science and Engineering*, vol. 197, no. 8, pp. 1656–1680, 2023. DOI: [10.1080/00295639.2022.2149231](https://doi.org/10.1080/00295639.2022.2149231). eprint: <https://doi.org/10.1080/00295639.2022.2149231>. URL: <https://doi.org/10.1080/00295639.2022.2149231>.
- [41] D. P. Adhikary, C. Jayasundara, R. K. Podgorney, and A. H. Wilkins, “A robust return-map algorithm for general multisurface plasticity,” *International Journal for Numerical Methods in Engineering*, vol. 109, pp. 218–234, Jan. 2016. DOI: [10.1002/nme.5284](https://doi.org/10.1002/nme.5284).
- [42] R. J. LeVeque, *Finite Volume Methods for Hyperbolic Problems* (Cambridge Texts in Applied Mathematics). Cambridge University Press, 2002.
- [43] M. Vinokur, “An analysis of finite-difference and finite-volume formulations of conservation laws,” *Journal of computational physics*, vol. 81, no. 1, pp. 1–52, 1989.
- [44] M. Hubbard, M. J. Baines, and P. Jimack, “Consistent dirichlet boundary conditions for numerical solution of moving boundary problems,” *Applied Numerical Mathematics*, vol. 59, no. 6, pp. 1337–1353, 2009.
- [45] T. J. Hughes, G. Engel, L. Mazzei, and M. G. Larson, “The continuous galerkin method is locally conservative,” *Journal of Computational Physics*, vol. 163, no. 2, pp. 467–488, 2000.
- [46] *10. Heating and Energy Deposition*, https://docs.openmc.org/en/stable/methods/energy_deposition.html, Accessed: 2024-06.
- [47] J. S. Rosenthal, “Parallel computing and monte carlo algorithms,” *Far east journal of theoretical statistics*, vol. 4, no. 2, pp. 207–236, 2000.
- [48] W. McKinney *et al.*, “Data structures for statistical computing in python.,” in *SciPy*, vol. 445, 2010, pp. 51–56.
- [49] T. D. Blacker, W. J. Bohnhoff, and T. L. Edwards, “Cubit mesh generation environment. volume 1: Users manual,” Sandia National Lab.(SNL-NM), Albuquerque, NM (United States), Tech. Rep., 1994.

- [50] M. A. Abdou, C. W. Maynard, and R. Q. Wright, “Mack: Computer program to calculate neutron energy release parameters (fluence-to-kerma factors) and multigroup neutron reaction cross sections from nuclear data in endf format,” Jul. 1973. URL: <https://www.osti.gov/biblio/4471714>.
- [51] D. Muir, R. Boicourt, and A. Kahler, “The njoy nuclear data processing system, version 2012,” *Los Alamos, USA: Los Alamos National Laboratory*, 2012.
- [52] *The National Reactor Innovation Center’s (NRIC) Virtual Test Bed Repository*, https://github.com/idaholab/virtual_test_bed/tree/devel, Accessed: 2021-01.
- [53] *Thermal Expansion*, https://en.wikipedia.org/wiki/Thermal_expansion, Accessed: 2024-06.
- [54] J. E. Lennard-Jones, “Cohesion,” *Proceedings of the Physical Society*, vol. 43, no. 5, p. 461, Sep. 1931. DOI: [10.1088/0959-5309/43/5/301](https://doi.org/10.1088/0959-5309/43/5/301). URL: <https://dx.doi.org/10.1088/0959-5309/43/5/301>.
- [55] A. F. Bower, *Applied mechanics of solids*. CRC press, 2009.
- [56] M. Edenius, “Studies of the reactivity temperature coefficient in light water reactors,” Chalmers Tekniska Hoegskola, Tech. Rep., 1976.
- [57] J. Latche, F. Lamare, and M. Cranga, “Computing reactivity initiated accidents in pwr’s,” 1995.
- [58] Y. Yun, “Thermal expansion,” Paul Scherrer Institut, Tech. Rep., 2015.
- [59] Y. Sumi, L. Keer, and S. Nemat-Nasser, “Thermally induced radial cracking in fuel element pellets,” *Journal of Nuclear Materials*, vol. 96, no. 1-2, pp. 147–159, 1981.
- [60] J. Rees, “Wigner energy in irradiated graphite and post-closure safety,” Environmental Agency (United Kingdom), Tech. Rep., 2002.
- [61] M. Abdullah, M. Arshad, S. Pervez, and K. Akhtar, “Beryllium reflector elements for parr-1,” Pakistan Inst. of Nuclear Science and Technology, Tech. Rep., 1996.
- [62] B. Marsden, “Nuclear graphite for high temperature reactors,” AEA Technology, Tech. Rep., 2001.
- [63] A. Massih, “Irradiation induced swelling of nuclear fuel,” ASEA-ATOM AB, Tech. Rep., 1988.
- [64] A. F. Lietzke, *Simplified analysis of nuclear fuel pin swelling*. National Aeronautics and Space Administration, 1970.
- [65] T. Drezewiecki, J. Schmidt, C. Van Wert, and P. Clifford, “Fuel Qualification for Advanced Reactors NUREG-2246,” U.S. Nuclear Regulator Commission, Tech. Rep., 2021.
- [66] N. C. Andrews, “Development of fission gas swelling and release models for metallic nuclear fuels,” Ph.D. dissertation, Massachusetts Institute of Technology, 2012.

- [67] A. Aitkaliyeva, L. He, H. Wen, B. Miller, X. Bai, and T. Allen, “7 - irradiation effects in generation iv nuclear reactor materials,” in *Structural Materials for Generation IV Nuclear Reactors*, P. Yvon, Ed., Woodhead Publishing, 2017, pp. 253–283, ISBN: 978-0-08-100906-2. DOI: <https://doi.org/10.1016/B978-0-08-100906-2.00007-0>. URL: <https://www.sciencedirect.com/science/article/pii/B9780081009062000070>.
- [68] H. Chen, H. Xiao, C. Long, and X. Leng, “A model for calculating the irradiation swelling of aged absorber in nuclear control rods,” *Nuclear Engineering and Technology*, vol. 56, no. 2, pp. 552–557, 2024, ISSN: 1738-5733. DOI: <https://doi.org/10.1016/j.net.2023.10.032>. URL: <https://www.sciencedirect.com/science/article/pii/S1738573323004886>.
- [69] R. Daum, Y. Chu, and A. Motta, “Identification and quantification of hydride phases in zircaloy-4 cladding using synchrotron x-ray diffraction,” *Journal of Nuclear Materials*, vol. 392, no. 3, pp. 453–463, 2009, ISSN: 0022-3115. DOI: <https://doi.org/10.1016/j.jnucmat.2009.04.004>. URL: <https://www.sciencedirect.com/science/article/pii/S0022311509005467>.
- [70] C. Taylor, T. Smith, M. Hahn, J. Torres, E. Tegtmeier, T. Nizolek, K. Bohn, C. Emblerley, A. Shivprasad, S. Widgeon-Paisner, *et al.*, “Properties of uranium-zirconium hydride moderated nuclear fuel synthesized by powder metallurgy,” Tech. Rep. LA-UR-22-29969, Los Alamos National Laboratory, Tech. Rep., 2022.
- [71] Coreform, *Neutronics on exact CAD geometry: advances in the Coreform Cubit DAGMC workflow*, <https://coreform.com/news/neutronics-on-exact-cad-geometry-advances-in-the-coreform-cubit-dagmc-workflow/>, Accessed: 2024-07-01.
- [72] *Significant Slowdown in Simulation Time in DAG-OpenMC*, <https://github.com/svalinn/DAGMC/discussions/953>, Accessed: 2024-03.
- [73] G. Hu, R. Hu, J. Kelly, and J. Ortensi, “Multi-physics simulations of heat pipe micro reactor,” Argonne National Lab.(ANL), Argonne, IL (United States), Tech. Rep., 2019.
- [74] Y. Ma, W. Han, B. Xie, H. Yu, M. Liu, X. He, S. Huang, Y. Liu, and X. Chai, “Coupled neutronic, thermal-mechanical and heat pipe analysis of a heat pipe cooled reactor,” *Nuclear Engineering and Design*, vol. 384, p. 111473, 2021, ISSN: 0029-5493. DOI: <https://doi.org/10.1016/j.nucengdes.2021.111473>. URL: <https://www.sciencedirect.com/science/article/pii/S0029549321004258>.
- [75] B. Yan, C. Wang, and L. Li, “The technology of micro heat pipe cooled reactor: A review,” *Annals of Nuclear Energy*, vol. 135, p. 106948, 2020, ISSN: 0306-4549. DOI: <https://doi.org/10.1016/j.anucene.2019.106948>. URL: <https://www.sciencedirect.com/science/article/pii/S0306454919304359>.
- [76] C. Chen, H. Mei, Z. Wang, S. Zhang, S. Liu, H. Wang, S. Chen, M. He, and T. Li, “Study of the thermal expansion effects of a space nuclear reactor with an integrated honeycomb core design using openmc and ansys,” *Annals of Nuclear Energy*, vol. 191, p. 109901, 2023, ISSN: 0306-4549. DOI: <https://doi.org/10.1016/j.anucene.2023.109901>. URL: <https://www.sciencedirect.com/science/article/pii/S0306454923002207>.

- [77] C. Pope and E. Lum, “Nuclear reactor thermal expansion reactivity effect determination using finite element analysis coupled with monte carlo neutron transport analysis,” in *Finite Element Methods and Their Applications*, M. Baccouch, Ed., Rijeka: IntechOpen, 2020, ch. 7. DOI: [10.5772/intechopen.93762](https://doi.org/10.5772/intechopen.93762). URL: <https://doi.org/10.5772/intechopen.93762>.
- [78] W. Xiao, X. Li, P. Li, T. Zhang, and X. Liu, “High-fidelity multi-physics coupling study on advanced heat pipe reactor,” *Computer Physics Communications*, vol. 270, p. 108 152, 2022, ISSN: 0010-4655. DOI: <https://doi.org/10.1016/j.cpc.2021.108152>. URL: <https://www.sciencedirect.com/science/article/pii/S0010465521002642>.
- [79] M. J. Jeong, J. Im, S. Lee, and H. K. Cho, “Multiphysics analysis of heat pipe cooled microreactor core with adjusted heat sink temperature for thermal stress reduction using openfoam coupled with neutronics and heat pipe code,” *Frontiers in Energy Research*, vol. 11, 2023, ISSN: 2296-598X. DOI: [10.3389/fenrg.2023.1213000](https://doi.org/10.3389/fenrg.2023.1213000). URL: <https://www.frontiersin.org/journals/energy-research/articles/10.3389/fenrg.2023.1213000>.
- [80] E. Walker, S. Skutnik, W. Wieselquist, A. Shaw, and F. Bostelmann, “Scale modeling of the fast spectrum heat pipe reactor,” Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States), Tech. Rep., 2022.
- [81] F. Aldebie, K. Fernandez-Cosials, and Y. Hassan, “Thermal-mechanical safety analysis of heat pipe micro reactor,” *Nuclear Engineering and Design*, vol. 420, p. 113 003, 2024, ISSN: 0029-5493. DOI: <https://doi.org/10.1016/j.nucengdes.2024.113003>. URL: <https://www.sciencedirect.com/science/article/pii/S0029549324001031>.
- [82] J. W. Sterbentz, J. E. Werner, M. G. McKellar, A. J. Hummel, J. C. Kennedy, R. N. Wright, and J. M. Biersdorf, “Special purpose nuclear reactor (5 mw) for reliable power at remote sites assessment report,” Idaho National Lab.(INL), Idaho Falls, ID (United States), Tech. Rep., 2017.
- [83] B. S. Kirk, J. W. Peterson, R. H. Stogner, and G. F. Carey, “Libmesh: A c++ library for parallel adaptive mesh refinement/coarsening simulations,” vol. 22, no. 3–4, pp. 237–254, Dec. 2006, ISSN: 0177-0667.
- [84] *Macroscopic Slowing Down Power – MSDP*, <https://www.nuclear-power.com/glossary/macroscopic-slowing-down-power-msdp/>, Accessed: 2024-05-5.
- [85] X. Hu, P. Mouche, D. Schappel, and D. J. Sprouster, “Technology enabling zero-epz micro modular reactors: Fabrication of zirconium hydride with controlled hydrogen loading,” Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States); Stony Brook . . . , Tech. Rep., 2020.
- [86] Z. Wang, F. Liu, Z. Guo, J. Zhang, L. Wang, and G. Yan, “Advance in and prospect of moderator materials for space nuclear reactors,” *International Journal of Energy Research*, vol. 45, no. 8, pp. 11 493–11 509, 2021.
- [87] *Zirconium Hydride Powder*, <https://www.us-nano.com/inc/sdetail/14745/>, Accessed: 2024-07-5.

- [88] H. G. Rickover, L. D. Geiger, and B. Lustman, “History of the development of zirconium alloys for use in nuclear reactors,” Mar. 1975. DOI: [10.2172/4240391](https://doi.org/10.2172/4240391). URL: <https://www.osti.gov/biblio/4240391>.
- [89] O. Courty, A. T. Motta, and J. D. Hales, “Modeling and simulation of hydrogen behavior in zircaloy-4 fuel cladding,” *Journal of Nuclear Materials*, vol. 452, no. 1, pp. 311–320, 2014, ISSN: 0022-3115. DOI: <https://doi.org/10.1016/j.jnucmat.2014.05.013>. URL: <https://www.sciencedirect.com/science/article/pii/S0022311514002803>.
- [90] J. Gylfe, H. Rood, J. Greenleaf, K. Balkwill, L. Prem, and L. Goldfisher, “Evaluation of zirconium hydride as moderator in integral, boiling water-superheat reactors,” Atomics International. Div. of North American Aviation, Inc., Canoga Park . . . , Tech. Rep., 1962.
- [91] D. Parkison, M. Tunes, T. Nizolek, T. Saleh, P. Hosemann, and C. Kohnert, “Fabrication of bulk delta-phase zirconium hydride from zircaloy-4 for use as moderators in microreactors,” *Scripta Materialia*, vol. 239, p. 115 771, Jan. 2024, ISSN: 1359-6462. DOI: [10.1016/j.scriptamat.2023.115771](https://doi.org/10.1016/j.scriptamat.2023.115771). URL: <http://dx.doi.org/10.1016/j.scriptamat.2023.115771>.
- [92] F. Long, N. Badr, Z. Yao, and M. Daymond, “Towards resolving a long existing phase stability controversy in the zr-h, ti-h systems,” *Journal of Nuclear Materials*, vol. 543, p. 152 540, 2021, ISSN: 0022-3115. DOI: <https://doi.org/10.1016/j.jnucmat.2020.152540>. URL: <https://www.sciencedirect.com/science/article/pii/S002231152031148X>.
- [93] Z. Zhao, J.-P. Morniroli, A. Legris, A. Ambard, Y. Khin, L. Legras, and M. Blat-Yrieix, “Identification and characterization of a new zirconium hydride,” *Journal of Microscopy*, vol. 232, no. 3, pp. 410–421, 2008. DOI: <https://doi.org/10.1111/j.1365-2818.2008.02136.x>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1365-2818.2008.02136.x>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1365-2818.2008.02136.x>.
- [94] K. Barraclough and C. Beevers, “Some observations on the phase transformations in zirconium hydrides,” *Journal of Nuclear Materials*, vol. 34, no. 2, pp. 125–134, 1970, ISSN: 0022-3115. DOI: [https://doi.org/10.1016/0022-3115\(70\)90112-1](https://doi.org/10.1016/0022-3115(70)90112-1). URL: <https://www.sciencedirect.com/science/article/pii/0022311570901121>.
- [95] L. Lanzani and M. Ruch, “Comments on the stability of zirconium hydride phases in zircaloy,” *Journal of Nuclear Materials*, vol. 324, no. 2, pp. 165–176, 2004, ISSN: 0022-3115. DOI: <https://doi.org/10.1016/j.jnucmat.2003.09.013>. URL: <https://www.sciencedirect.com/science/article/pii/S0022311503004409>.
- [96] B. W. Veal, D. J. Lam, and D. G. Westlake, “X-ray photoemission spectroscopy study of zirconium hydride,” *Phys. Rev. B*, vol. 19, pp. 2856–2863, 6 Mar. 1979. DOI: [10.1103/PhysRevB.19.2856](https://doi.org/10.1103/PhysRevB.19.2856). URL: <https://link.aps.org/doi/10.1103/PhysRevB.19.2856>.
- [97] M. P. Cassidy and C. M. Wayman, “The crystallography of hydride formation in zirconium: II. the $\delta \rightarrow \epsilon$ transformation,” *Metallurgical Transactions A*, vol. 11, no. 1, pp. 57–67, Dec. 1980. DOI: [10.1007/BF02700438](https://doi.org/10.1007/BF02700438).

- [98] R. C. Bowman, E. L. Venturini, B. D. Craft, A. Attalla, and D. B. Sullenger, “Electronic structure of zirconium hydride: A proton nmr study,” *Phys. Rev. B*, vol. 27, pp. 1474–1488, 3 Feb. 1983. DOI: [10.1103/PhysRevB.27.1474](https://doi.org/10.1103/PhysRevB.27.1474). URL: <https://link.aps.org/doi/10.1103/PhysRevB.27.1474>.
- [99] A. Trainer, “Hydrogen distribution in metal hydrides and its effect on reactor physics calculations,” Ph.D. dissertation, Massachusetts Institute of Technology, 2023.
- [100] T. Maimaitiyili, A. Steuwer, J. Blomqvist, C. Bjerken, M. S. Blackmur, O. Zanellato, J. Andrieux, and F. Ribeiro, “Observation of the δ to ϵ zr-hydride transition by in-situ synchrotron x-ray diffraction,” *Crystal Research and Technology*, vol. 51, no. 11, pp. 663–670, 2016.
- [101] R. Van Houten, “Selected engineering and fabrication aspects of nuclear metal hydrides (li, ti, zr, and y),” *Nuclear Engineering and Design*, vol. 31, no. 3, pp. 434–448, 1974, ISSN: 0029-5493. DOI: [https://doi.org/10.1016/0029-5493\(75\)90178-8](https://doi.org/10.1016/0029-5493(75)90178-8). URL: <https://www.sciencedirect.com/science/article/pii/0029549375901788>.
- [102] A. Shivprasad, D. Frazer, V. Mehta, M. Cooper, T. Saleh, J. White, J. Wermer, E. Luther, and D. Rao, “Elastic moduli of high-density, sintered monoliths of yttrium dihydride,” *Journal of Alloys and Compounds*, vol. 826, p. 153 955, 2020, ISSN: 0925-8388. DOI: <https://doi.org/10.1016/j.jallcom.2020.153955>. URL: <https://www.sciencedirect.com/science/article/pii/S0925838820303182>.
- [103] D. Zhao, A. Würger, and X. Crispin, “Ionic thermoelectric materials and devices,” *Journal of Energy Chemistry*, vol. 61, pp. 88–103, 2021, ISSN: 2095-4956. DOI: <https://doi.org/10.1016/j.jechem.2021.02.022>. URL: <https://www.sciencedirect.com/science/article/pii/S2095495621001054>.
- [104] G. Majer, W. Renz, and R. G. Barnes, “The mechanism of hydrogen diffusion in zirconium dihydrides,” *Journal of Physics: Condensed Matter*, vol. 6, no. 15, p. 2935, Apr. 1994. DOI: [10.1088/0953-8984/6/15/015](https://doi.org/10.1088/0953-8984/6/15/015). URL: <https://dx.doi.org/10.1088/0953-8984/6/15/015>.
- [105] J. Huang, B. Tsuchiya, K. Konashi, and M. Yamawaki, “Estimation of hydrogen redistribution in zirconium hydride under temperature gradient,” *Journal of Nuclear Science and Technology*, vol. 37, no. 10, pp. 887–892, 2000. DOI: [10.1080/18811248.2000.9714969](https://doi.org/10.1080/18811248.2000.9714969). eprint: <https://doi.org/10.1080/18811248.2000.9714969>. URL: <https://doi.org/10.1080/18811248.2000.9714969>.
- [106] A. Steuwer, J. Santisteban, M. Preuss, M. Peel, T. Buslaps, and M. Harada, “Evidence of stress-induced hydrogen ordering in zirconium hydrides,” *Acta Materialia*, vol. 57, no. 1, pp. 145–152, 2009, ISSN: 1359-6454. DOI: <https://doi.org/10.1016/j.actamat.2008.08.061>. URL: <https://www.sciencedirect.com/science/article/pii/S1359645408006332>.
- [107] C. Matthews, A. P. Shivprasad, and M. W. D. Cooper, “Metal hydride simulations using swift,” Los Alamos National Laboratory (LANL), Los Alamos, NM (United States), Tech. Rep., 2021.

- [108] V. K. Mehta, J. Armstrong, D. V. Rao, and D. Kotlyar, “Capturing multiphysics effects in hydride moderated microreactors using marm,” *Annals of Nuclear Energy*, vol. 172, p. 109 067, 2022.
- [109] B. Kammenzind, B. Berquist, R. Bajaj, P. Kreyns, and D. Franklin, “The long range migration of hydrogen through zircaloy in response to tensile and compressive stress gradients,” Bettis Atomic Power Lab., West Mifflin, PA (United States), Tech. Rep., 1998.
- [110] P. R. McClure, D. I. Poston, V. R. Dasari, and R. S. Reid, “Design of megawatt power level heat pipe reactors,” Los Alamos National Laboratory (LANL), Los Alamos, NM (United States), Tech. Rep., 2015.
- [111] J. Yanez, M. Kuznetsov, and A. Souto-Iglesias, “An analysis of the hydrogen explosion in the fukushima-daiichi accident,” *International Journal of Hydrogen Energy*, vol. 40, no. 25, pp. 8261–8280, 2015, ISSN: 0360-3199. DOI: <https://doi.org/10.1016/j.ijhydene.2015.03.154>. URL: <https://www.sciencedirect.com/science/article/pii/S0360319915008009>.
- [112] V. K. Mehta, P. McClure, and D. Kotlyar, “Selection of a space reactor moderator using lessons learned from snap and anp programs,” in *AIAA Propulsion and Energy 2019 Forum*, 2019, p. 4451.
- [113] F. Xiaogang, Q. Bo, M. Haoran, Z. Jinqun, and L. Bin, “A Study on Hydrogen Emission of Zirconium Hydride,” *Journal of Nuclear Engineering and Radiation Science*, vol. 5, no. 1, p. 014 501, Jan. 2019, ISSN: 2332-8983. DOI: [10.1115/1.4041274](https://doi.org/10.1115/1.4041274). eprint: https://asmedigitalcollection.asme.org/nuclearengineering/article-pdf/5/1/014501/6070916/ners_005_01_014501.pdf. URL: <https://doi.org/10.1115/1.4041274>.
- [114] D. Wongsawaeng and S. Jaiyen, “High-temperature absolute hydrogen desorption kinetics of zirconium hydride under clean and oxidized surface conditions,” *Journal of Nuclear Materials*, vol. 403, no. 1, pp. 19–24, 2010, ISSN: 0022-3115. DOI: <https://doi.org/10.1016/j.jnucmat.2010.05.025>. URL: <https://www.sciencedirect.com/science/article/pii/S0022311510002370>.
- [115] D. P. Griesheimer, W. R. Martin, and J. P. Holloway, “Convergence properties of monte carlo functional expansion tallies,” *Journal of Computational Physics*, vol. 211, no. 1, pp. 129–153, 2006, ISSN: 0021-9991. DOI: <https://doi.org/10.1016/j.jcp.2005.05.023>. URL: <https://www.sciencedirect.com/science/article/pii/S0021999105002652>.
- [116] D. P. Griesheimer, “Functional expansion tallies for Monte Carlo simulations,” Ph.D. dissertation, University of Michigan, Jan. 2005.
- [117] *Section 10.16: Taylor Series*, <https://tutorial.math.lamar.edu/classes/calci/taylorseries.aspx>, Accessed: 2024-07-01.
- [118] S. O. Rice, “Mathematical analysis of random noise,” *The Bell System Technical Journal*, vol. 23, no. 3, pp. 282–332, 1944. DOI: [10.1002/j.1538-7305.1944.tb00874.x](https://doi.org/10.1002/j.1538-7305.1944.tb00874.x).
- [119] W. L. Chadsey, C. W. Wilson, and V. W. Pine, “X-ray photoemission calculations,” *IEEE Transactions on Nuclear Science*, vol. 22, no. 6, pp. 2345–2350, 1975. DOI: [10.1109/TNS.1975.4328131](https://doi.org/10.1109/TNS.1975.4328131).

- [120] A. Noel and H. Wio, “A new series-expansion approach in monte carlo: Application to neutron shielding,” *Annals of Nuclear Energy*, vol. 11, no. 5, pp. 225–227, 1984, ISSN: 0306-4549. DOI: [https://doi.org/10.1016/0306-4549\(84\)90053-7](https://doi.org/10.1016/0306-4549(84)90053-7). URL: <https://www.sciencedirect.com/science/article/pii/0306454984900537>.
- [121] F. Brown and W. Martin, “Direct sampling of monte carlo flight paths in media with continuously varying cross-sections,” *ANS Mathematics and Computation Topical Meeting*, Jan. 2003.
- [122] D. Griesheimer and W. Martin, “Estimating the global scalar flux distribution with orthogonal function expansions,” *Transactions-American Nuclear Society*, pp. 462–464, 2003.
- [123] D. P. Griesheimer and W. R. Martin, “Two dimensional functional expansion tallies for monte carlo simulations,” *ANS PHYSOR 2004-Conference Proceedings*, 2004.
- [124] L. Kerby, A. Tumulak, J. Leppänen, and V. Valtavirta, “Preliminary serpent-moose coupling and implementation of functional expansion tallies in serpent,” in *International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering, M&C 2017*, 2017.
- [125] B. Wendt, L. Kerby, A. Tumulak, and J. Leppänen, “Advancement of functional expansion capabilities: Implementation and optimization in serpent 2,” *Nuclear Engineering and Design*, vol. 334, pp. 138–153, 2018, ISSN: 0029-5493. DOI: <https://doi.org/10.1016/j.nucengdes.2018.05.004>. URL: <https://www.sciencedirect.com/science/article/pii/S0029549318305144>.
- [126] B. L. Wendt, “Functional expansions methods: Optimizations, characterizations, and multiphysics practices generalized data representation and transfer solutions in multiphysics simulations through the characterization and advancement of functional expansion implementations,” Ph.D. dissertation, Idaho State University, 2018.
- [127] A. royale des sciences (France), *Mémoires de mathématique et de physique, présentés à l’Académie royale des sciences, par divers sçavans & lûs dans ses assemblées*. Paris, 1768, vol. Vol 5, p. 772, <https://www.biodiversitylibrary.org/bibliography/4360>. URL: <https://www.biodiversitylibrary.org/item/88512>.
- [128] A. Janssen, *A generalization of the zernike circle polynomials for forward and inverse problems in diffraction theory*, 2011. arXiv: [1110.2369 \[math-ph\]](https://arxiv.org/abs/1110.2369). URL: <https://arxiv.org/abs/1110.2369>.
- [129] Z. Han, “Performance analysis of functional expansion tallies on 2d pwr pin cell,” Ph.D. dissertation, Massachusetts Institute of Technology, 2020.
- [130] B. Wendt, A. Novak, L. Kerby, and P. Romano, “Integration of functional expansion methodologies as a moose module,” in *PHYSOR 2018: Reactor Physics paving the way towards more efficient systems*, Apr. 2018.
- [131] V. K. Mehta, J. Armstrong, D. V. Rao, and D. Kotlyar, “Capturing multiphysics effects in hydride moderated microreactors using marm,” *Annals of Nuclear Energy*, vol. 172, p. 109 067, 2022.

- [132] G. Abaqus, “Abaqus 6.11,” *Dassault Systemes Simulia Corporation, Providence, RI, USA*, vol. 3, 2011.
- [133] J. K. Shultis and R. E. Faw, “An mcnp primer,” Kansas State University, Tech. Rep., 2011.
- [134] G. L. Giudicelli, A. Abou-Jaoude, A. J. Novak, *et al.*, “The virtual test bed (vtb) repository: A library of reference reactor models using neams tools,” *Nuclear Science and Engineering*, vol. 197, no. 8, pp. 1–17, 2023. DOI: [10.1080/00295639.2022.2142440](https://doi.org/10.1080/00295639.2022.2142440). URL: <https://doi.org/10.1080/00295639.2022.2142440>.
- [135] C. Matthews, A. P. Shivprasad, and M. W. D. Cooper, “Metal hydride simulations using swift,” Los Alamos National Laboratory (LANL), Los Alamos, NM (United States), Tech. Rep., 2021.
- [136] M. Ito, “Studies on physical properties of metal hydrides,” Ph.D. dissertation, Osaka University, 2008.
- [137] R. E. P. Todd J. Urbatsch R. Arthur Forster and R. J. Beckman, “Estimation and interpretation of keff confidence intervals in mcnp,” *Nuclear Technology*, vol. 111, no. 2, pp. 169–182, 1995. DOI: [10.13182/NT95-A35128](https://doi.org/10.13182/NT95-A35128). eprint: <https://doi.org/10.13182/NT95-A35128>. URL: <https://doi.org/10.13182/NT95-A35128>.
- [138] E. Woodcock, T. Murphy, P. Hemmings, and S. Longworth, “Techniques used in the gem code for monte carlo neutronics calculations in reactors and other systems of complex geometry,” in *Proc. Conf. Applications of Computing Methods to Reactor Problems*, Argonne National Laboratory, vol. 557, 1965.
- [139] J. Leppänen, “Modeling of nonuniform density distributions in the serpent 2 monte carlo code,” en, *Nuclear Science and Engineering*, vol. 174, no. 3, pp. 318–325, Jul. 2013, ISSN: 0029-5639, 1943-748X. DOI: [10.13182/NSE12-54](https://doi.org/10.13182/NSE12-54).
- [140] T. Viitanen and J. Leppänen, “Explicit treatment of thermal motion in continuous-energy monte carlo tracking routines,” en, *Nuclear Science and Engineering*, vol. 171, no. 2, pp. 165–173, Jun. 2012, ISSN: 0029-5639, 1943-748X. DOI: [10.13182/NSE11-36](https://doi.org/10.13182/NSE11-36).