ORNL/SPR-2023/3080

Enhanced LWR High Burnup Transient Simulation Capabilities to Support AOO Margin Identification

Aaron Graham
Mehdi Asgari
Shane Henderson

September 2023

OAK RIDGE National Laboratory

ORNL IS MANAGED BY UT–BATTELLE, LLC, FOR THE U.S. DEPARTMENT OF ENERGY

ORNL/SPR-2023/3080

Nuclear Energy and Fuel Cycle Division

ENHANCED LWR HIGH BURNUP TRANSIENT SIMULATION CAPABILITIES TO SUPPORT AOO MARGIN IDENTIFICATION

Aaron Graham
Mehdi Asgari
Shane Henderson

September 2023

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, TN 37831-6283
managed by
UT-BATTELLE, LLC
for the
US DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABBREVIATIONS

API	application programming interface
BWR	boiling-water reactor
CMFD	coarse mesh finite difference
DOE	Department of Energy
FFRD	fuel fragmentation, relocation, and dispersion
H1C1	Hatch unit 1 cycle 1
HBu	high-burnup
HDF5	Hierarchical Data Format 5
LOCA	loss-of-coolant accident
LWR	light-water reactor
NEAMS	Nuclear Energy Advanced Modeling and Simulation
PKE	point kinetics equations
PLR	part-length rod
PWR	pressurized water reactor
RMSE	root mean square error
STH	simplified thermal hydraulics
TH	thermal hydraulics
TIP	transverse in-core probe
TML	transient multilevel
VERA	Virtual Environment for Reactor Applications

ABSTRACT

As part of ongoing efforts to support the Nuclear Energy Advanced Modeling and Simulation (NEAMS) program's development of a fuel fragmentation, relocation, and dispersal (FFRD) screening methodology, a number of improvements are required for the NEAMS core simulation capabilities, namely the Virtual Environment for Reactor Applications (VERA). Three areas of improvement were identified in VERA which are important for continued development and application of the FFRD screening methodology. First, the FFRD screening methodology will soon be extended to boiling water reactors (BWRs), requiring development and validation of the VERA BWR capabilities. Second, the screening methodology occasionally requires that VERA be used to simulate a transient in addition to nominal operations. Thus, improvements to both accuracy and performance of the VERA transient capabilities are necessary. Third, the VERAOneWay component of VERA is used to develop BISON fuel performance inputs using the rod-by-rod histories calculated by VERA. Prior use of VERAOneWay exposed significant accuracy, robustness, and performance issues with VERAOneWay; these must be addressed for it to be an effective tool in the NEAMS FFRD methodology. This report documents the efforts in FY23 in each of these three areas to enable successful use of VERA and VERAOneWay for FFRD calculations in FY24 and following years.

1. INTRODUCTION

Light-water reactor (LWR) operators are currently considering transitions from conventional 18-month cycles to 24-month cycles to reduce fuel costs and outage time, improving the economic position of the commercial fleet. One concern that comes with such a change is the fuel integrity of fuel rods, which would operate in the core 33% longer than before. To support this mission, the US Department of Energy (DOE)'s Nuclear Energy Advanced Modeling and Simulation (NEAMS) program is developing a methodology to predict the likelihood of fuel fragmentation, relocation, and dispersion (FFRD) for high-burnup (HBu) core designs [1].

The FFRD screening methodology consists of three major components [2–4]. First, the core simulation is performed using a high-fidelity light-water reactor (LWR) core simulator, Virtual Environment for Reactor Applications (VERA) [5], to generate rod-by-rod power and burnup histories. Next, the results of these calculations can be fed into the TRACE systems analysis code [6] to simulate the system response during accident scenarios, which an important part of the thermal hydraulics (TH) qualification of fuel and cycle designs. Finally, the VERA results can also be fed into the BISON fuel performance code [7–9] for a detailed simulation of individual fuel rods during steady-state or transient conditions.

So far, applications of this methodology within the NEAMS program have focused on loss-of-coolant accident (LOCA) analysis, specifically for pressurized water reactors (PWRs). For this application, the core simulation is primarily used to simulate the depletion of the fuel during nominal operation and generate decay heat curves for the more detailed TRACE analysis. Because a LOCA is a TH-driven transient that occurs after core SCRAM, it is not necessary to run the core simulator for the transient because very little neutronics activity occurs.

In FY24 and the years that follow, the scope of the FFRD methodology will be expanded to boiling-water reactors (BWRs) and other types of transients. With this in view, additional developments are needed in the VERA core simulator to continue supporting the FFRD analysis for the full breadth of scenarios. This report documents progress toward that goal. The efforts presented here are threefold. First, the VERA BWR capabilities are quite new compared to the PWR capability, and they still experienced some robustness and accuracy issues before the work documented in this report. Effort was made to resolve these issues to bring

the BWR capabilities closer to the same quality as the PWR capabilities. Second, some anticipated transients of interest will require greater use of the VERA transient capability due to more significant neutronic feedback on the TH and fuel performance. Thus, a variety of usability and robustness improvements were pursued for the transient capability to support its future use. Finally, previous work identified performance and robustness issues with VERAOneWay, the package used to generate BISON input files from VERA output files. This package is crucial to the FFRD workflow, so an investigation into the these issues was conducted, and some initial progress was made on improving that package.

2. BWR IMPROVEMENTS

2.1 COUPLING DEFECT FIX

Before this work, BWR models in VERA were inconsistent in their accuracy. For some cycles, the errors compared to the transverse in-core probe (TIP) detectors were quite good, at around 5%; for subsequent cycles, those errors were then observed to be as high as 20%. These errors seemed to be correlated with transition cycles, during which a reactor's fuel is replaced with another type. However, the issue's cause was unclear because individual fuel bundle calculations matched TH validation well.

To explore this issue, a series of studies were performed, including neutronics, TH, and multiphysics studies. Many of these studies produced no clues about the issue and are not discussed here. Most recently, while developing additional Monte Carlo models for neutronics validation, an error was discovered in the application of the TH solution to the neutronics mesh. Specifically, regions of water directly above the tops of part-length rods (PLRs) were not being updated properly; the density of the water in those regions remained close to the inlet density (0% void) instead of begin reduced significantly (50% or greater void). This resulted in over-moderation in the top of the core, explaining why transition cycles were worse, since newer fuel types tend to have a larger number of PLRs.

Upon rerunning the validation cases after resolving the defect, the resulting errors were much more consistent than before, generally under 10% and frequently around 5%. Additionally, the specific trends of the errors within an assembly, across the core, and throughout a cycle made much more sense and were more regular across cycles. This finding confirms that the misapplication of water densities was the primary problem causing the large errors and that it has been fixed properly.

2.2 BYPASS HEATING MODEL

With the model defects ironed out, results were significantly improved. However, there was still a trend of the VERA-predicted axial power being less bottom-peaked early in the cycle compared with the measurements. One possibility to explain this was the fact that the axial profile of the moderator density in the bypass and water rod regions was constant, and therefore too much moderation was occurring in the top portion of the core and keeping the power higher than it should be. To address this, a bypass heating model was developed to obtain axial density profiles in the moderator.

Much of this model was developed in FY22, but the effects from the model were obscured by the defect discussed in the previous section. It was warranted to revisit the bypass heating model. There are 3 sources of heat to the bypass:

1. direct neutron heating to the channel box, bypass flow, and control blades;
2. conduction of heat through the channel box;
3. and direct gamma heating of the channel box, bypass flow, and control blades.

There are two ways to model this. First, a fraction of the total reactor power can simply be moved to the bypass flow based on known data and measurements; this option is referred to as *fixed heating*. For this option, VERA will take a specified fraction of the power produced by the assembly surrounded by one region of bypass flow. This ensures that some spatial distribution of the power is retained both radially and axially. This type of approach has been used in the past by multiple other codes. The second approach is to calculate the heating explicitly and is referred to as *explicit heating*. This can also be done in VERA. In this case, the neutron flux is used to calculate heat generation in the bypass components due to down-scatter and absorption, and a 1D slab conduction model is used to calculate the heat conduction between the bypass

flow and the in-channel coolant flow. Gamma transport is not supported in 3D calculations at this time, so direct heating due to gammas is neglected.

The fixed heating can be enabled by setting `bypass_treatment` to `fixed_heating` in the `[COUPLING]` block and then setting the amount of heat using the `bypass_fraction`; setting `bypass_fraction 0.03` would place 3% of the heat into the bypass. The explicit heating is enabled by setting `bypass_treatment` to `explicit_heating`. In this case, the `bypass_fraction` input is not required, but if it is still set to a non-zero value, then a hybrid approach is taken. In this case, the explicit heat calculations are done, and then a fixed fraction of *additional* heat is added to that distribution based on the value of `bypass_fraction`. The intention behind this is to allow the user to account for the fact that the code cannot capture gamma heating, so a small fixed fraction can be added to the explicit heating calculation to make up for that missed heating.

2.3 HATCH UNIT 1 CYCLE 1 RESULTS

Hatch unit 1 is a BWR that went online in 1974 and has public data available for the first 2 cycles. This section will present results for cycle 1. Hatch unit 1 cycle 1 (H1C1) is made up of 7×7 assemblies; being an older core, it had no water rods or PLRs like many modern BWR cores. This makes it much simpler to model than newer cores. However, it also had a very inconsistent operational history, which makes its operation difficult to model. Historically, it has been very difficult to obtain accurate simulation results for this core.

The base case presented here used saturated water in the bypass region and the simplified thermal hydraulics (STH) thermal hydraulics solver for the channel. Several additional permutations were also tested and will be compared with the base case: 1% fixed heating, 1.5% heating, explicit heating, explicit heating + 3% fixed heating, and explicit heating with no channel box conduction. All of these cases are compared to the TIP measurement data available for H1C1.

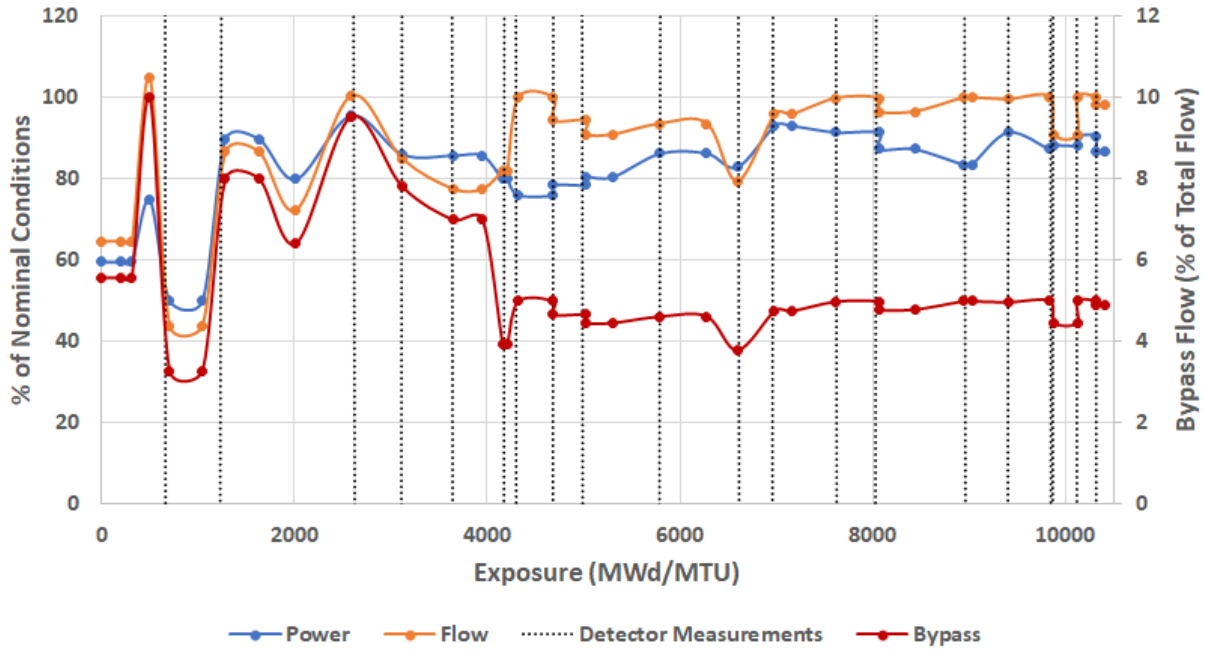

Figure 1. Operation history of Hatch unit 1 cycle 1 (H1C1) for power, flow, and bypass; detector measurements times are denoted with vertical dash lines

Figure 1 shows the power, flow, and bypass conditions for H1C1. There were many changes in power and flow curing cycle 1, which is common for the first cycle of operation in old reactors. A total of 40 states were selected for VERA simulation to approximate the actual operational history. These states are the points in Figure 1; the lines connecting the points are not directly simulated by VERA. Additionally, there were a total of 20 TIP measurements performed throughout the cycle; the VERA model was set up in such a way as to ensure that all 20 measurement times were also simulated by VERA.

2.3.1 Saturated Bypass Conditions

Figure 2 shows several detector map results for saturated bypass conditions, with the measurements in red and the VERA calculations in blue. The earliest state points were low power and low flow, so saturated conditions in the bypass are a good approximation. Thus, the total root mean square error (RMSE) for the first map at 0.702 GWDMT is 5.6%. This is a good result for comparison to BWR TIP data for three reasons. First, the measurement data itself is often quite noisy with non-physical artifacts, which affects the statistics negatively. Second, the two-phase flow in BWRs is notoriously difficult to resolve accurately enough to reproduce the true power history. Third, the VERA model was run using quarter symmetry but the physical reactor was not quarter symmetric; testing was done with VERA which indicated this is likely a small issue, but it is still worth noting. The second map, at 2.584 GWDMT, is taken at close to 100% nominal operating conditions. The results are much worse, with a total RMSE of 17.5%. It is clear in this map that the measured power is much more bottom-peaked than in the first map. This is due to increased power causing extensive boiling and voiding of the coolant in the upper portion of the core. This in turn reduces moderator and drives the power toward the bottom of the core. This is characteristic of all BWRs, but is clearly not predicted accurately by VERA. The third map is taken around the middle of the cycle at 5.024 GWDMT, and the fourth is taken at the end of the cycle at 10.311 GWDMT. The errors only get worse. The end of the cycle map is important because it demonstrates a power shift from the bottom of the core to the top of the core. This occurs because for much of the cycle the power was restricted to the bottom of the core due to voiding in the top. During this time, the under-moderated top of the core builds in extensive ^{239}Pu; once the fuel in the bottom of the core burns out, the voiding at the top of the core reduces and the ^{239}Pu that build up throughout the cycle begins to drive the power in the upper half of the core. However, VERA misses this effect because its simulations did not accurately predict the bottom-peaked power in the earlier parts of the core.

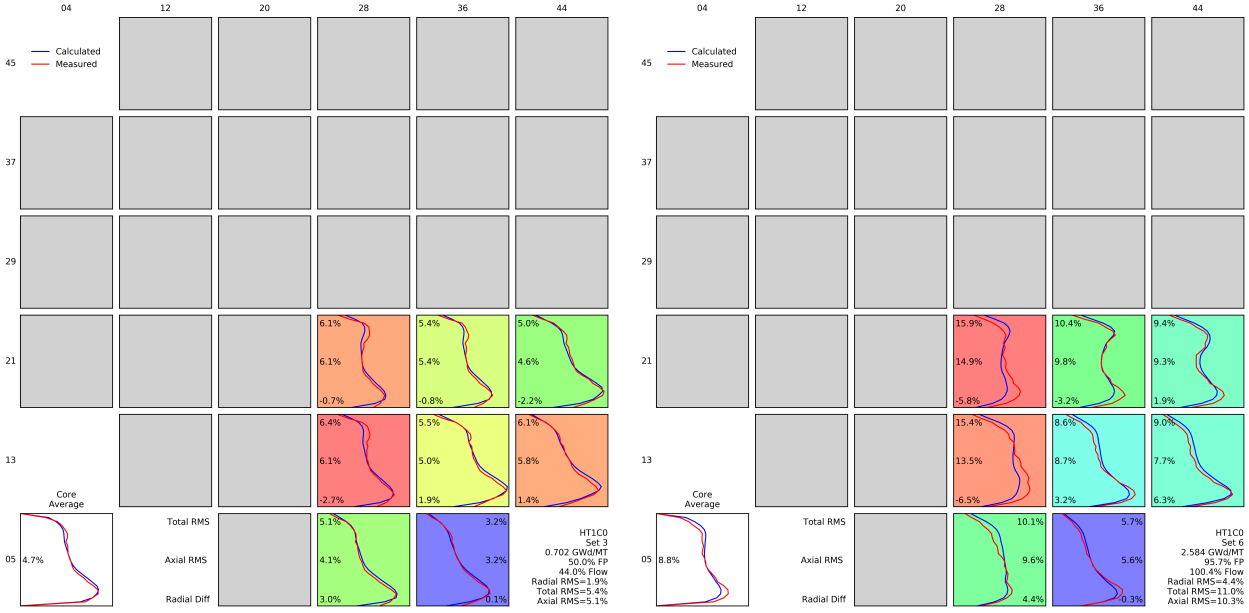

(a) Map 1: 0.702 GWDMT

(b) Map 3: 2.584 GWDMT

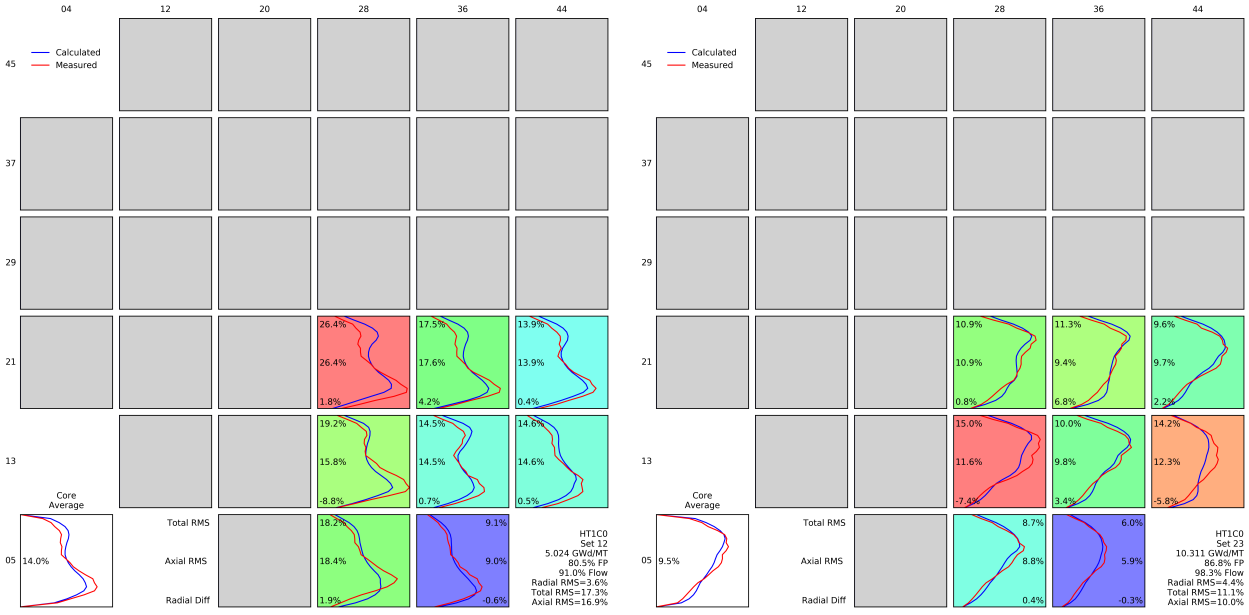

(c) Map 9: 5.024 GWDMT

(d) Map 20: 10.311 GWDMT

Figure 2. Selected detector measurement maps for Hatch unit 1 cycle 1 (H1C1) with saturated bypass conditions

2.3.2 1% Fixed Bypass Heating

The first set of fixed bypass heating calculations placed 1% of the heat into the bypass flow. At each axial level and each assembly, 1% of the heat was moved into the bypass that immediately surrounded that location. Thus, there is still some spatial shape to the bypass heating based on the spatial shape of the fuel heating. Figure 3 shows the same maps as before, but with the 1% fixed heating. The 2.584 GWDMT map total RMSE reduces from 17.5% to 9.8%, and the 5.024 GWDMT map total RMSE reduces from 22.1% to 13.3%. These are both major improvements. It is clear from both maps that the heating in the bypass provides an axial density profile, which results in moving the power downward in the core. As a result, the 10.311 GWDMT map looks completely different than it did before. The bottom of the core was properly burned, so VERA now captures the upward swing in power that occurs late in the cycle. This reduces the total RMSE from an unacceptable 40.9% to a respectable 6.5%. Visual inspection shows that the shapes are resolved quite well at each location. Despite these improvements, the 0.702 GWDMT map actually got much worse, with the total RMSE increasing from 5.6% to 18.4%. For this case, VERA over-predicts the bottom peaking of the power. This occurs because that state point was a low power state point. Thus the in-channel coolant would have been relatively high density compared with other state points, causing the power to remain a bit higher in the core. The higher density in the channel would have also prevented as much neutron leakage from the channel to the bypass, reducing the bypass heating. Thus, for nominal operation 1% heating seems to be a fairly good approximation, but for lower power operation it is too high and performs much worse than saturation. This indicates a need for something more robust than a fixed fraction of heat being placed in the bypass.

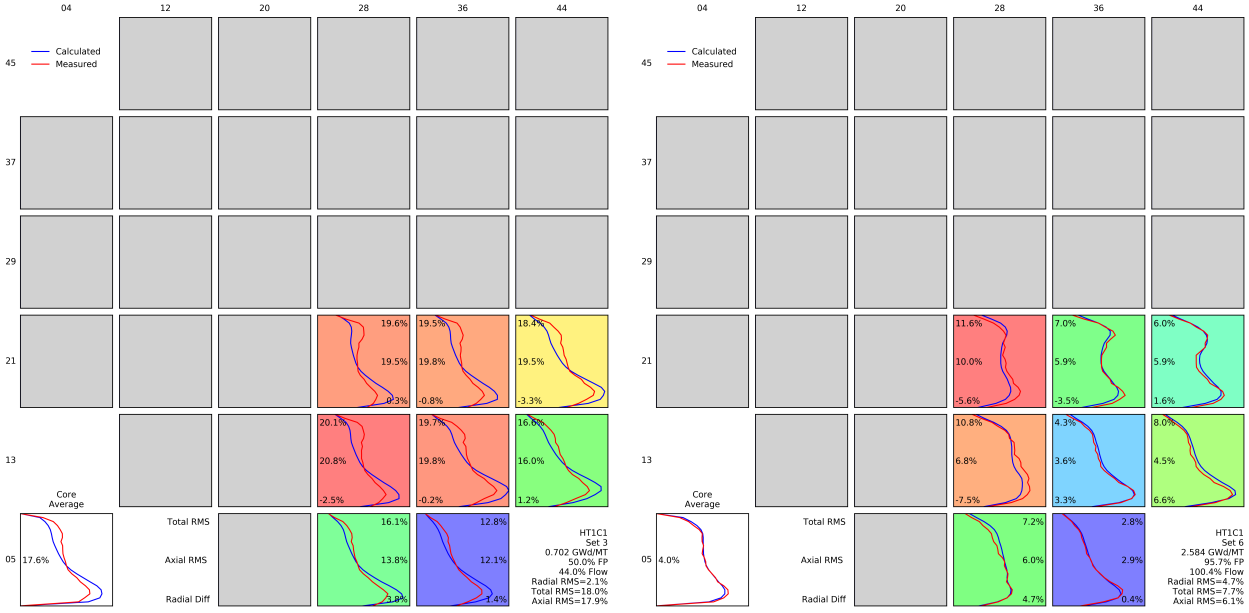

(a) Map 1: 0.702 GWDMT

(b) Map 3: 2.584 GWDMT

(c) Map 9: 5.024 GWDMT

(d) Map 20: 10.311 GWDMT

Figure 3. Selected detector measurement maps for Hatch unit 1 cycle 1 (H1C1) with 1% fixed bypass heating

2.3.3 Explicit Heating

H1C1 was also simulated using the explicit bypass heating model in an attempt to take advantage of the fixed heating benefits while still maintaining accuracy for off-nominal operating conditions. The results are shown in Figure 4. The 0.702 GWDMT map shows good results, with a total RMSE of 6.1%. The 2.584 GWDMT result has a total RMSE of 13.1%; this is worse than the fixed heating case but better than the saturated conditions. The same is true for the 5.024 GWDMT map, which has a total RMSE of 15.6%. Finally, the 10.311 GWDMT map gives a total RMSE of 10.3%, which is again better than the saturated conditions but not quite as good as the 1% heating. Overall, the explicit heating does a good job resolving the low power condition that was problematic for the 1% fixed heating.

For the 0.702 GWDMT state point, VERA calculates that 0.89% of the heat goes to the bypass; this number is adjusted downward to 0.76% because of conduction through the channel box. This shows that the 1% heating was simply too high, which is confirmed visually by the fact that the VERA calculations were more bottom-peaked than the measurements. This trend is consistent throughout the remainder of the cycle; the total bypass heating is typically between 0.9% and 0.95% before the channel box conduction and between 0.75% and 0.83% after. This indicates that 1% heating may have been too high for the early low-power state points, but that it was probably about right for the later state points. It seems that the channel box conduction is actually making the answers worse, but it is not clear why that would be the case.

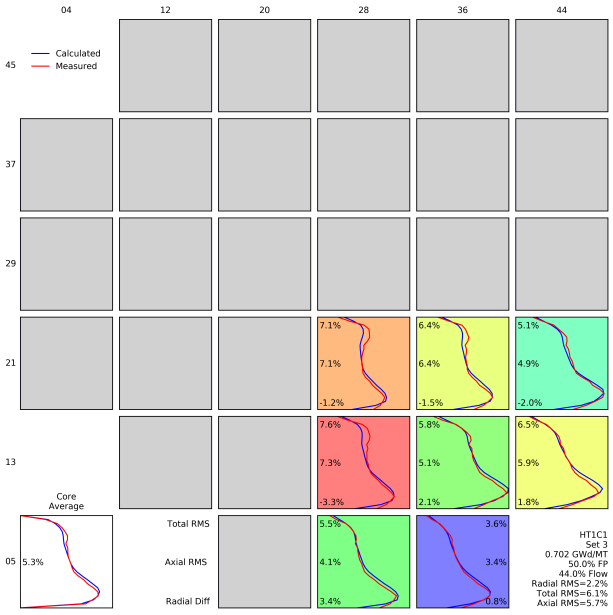

(a) Map 1: 0.702 GWDMT

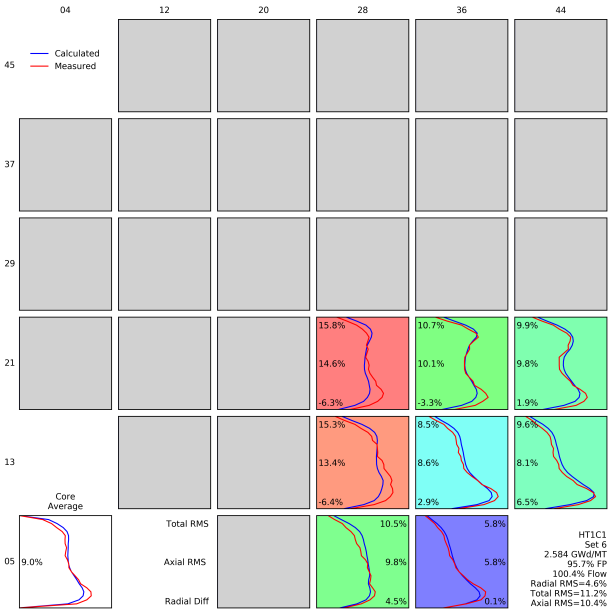

(b) Map 3: 2.584 GWDMT

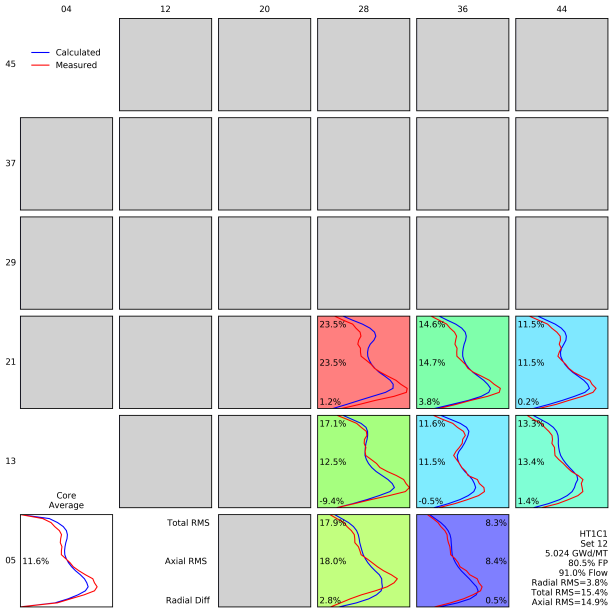

(c) Map 9: 5.024 GWDMT

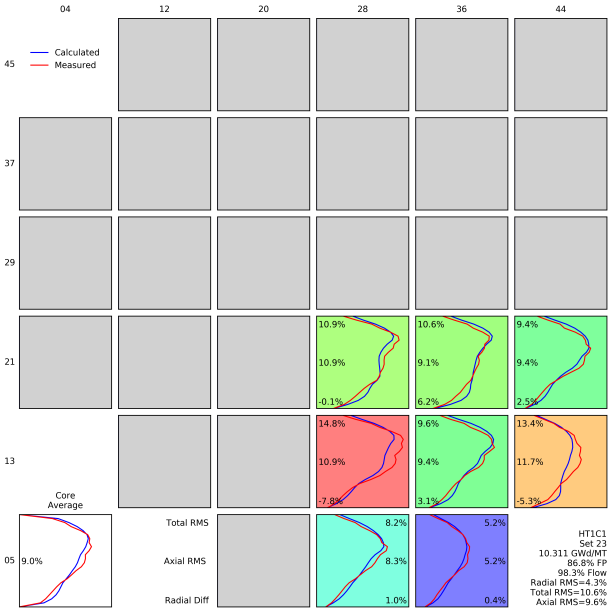

(d) Map 20: 10.311 GWDMT

Figure 4. Selected detector measurement maps for Hatch unit 1 cycle 1 (H1C1) with explicit heating calculation

2.3.4 Explicit Heating + 3%

Because the explicit heating model was not putting quite enough heat into the bypass to capture the axial power shape, another calculation was run that added extra fixed heating in addition to the explicit bypass heating. A fixed amount of 3% was added. This is a lot compared with the cases that only used fixed heating, but as shown in Figure 5 it produced surprisingly good answers. A closer inspection of the VERA output shows that even though nearly 4% of the heat was placed in the bypass (3% fixed plus nearly 1% explicit), most of that heat conducted from the bypass to the channel via the channel box. The effective heat that actually remained in the bypass region was generally between 1% and 1.3%. This confirms two facts. First, for this model, only around 1% of the heating should end up in the bypass; more than that will overestimate the axial shift in power. Second, there is likely a problem with the channel box conduction. It does not seem incorrect to be shifting so much heat back to the channel. This must be investigated.

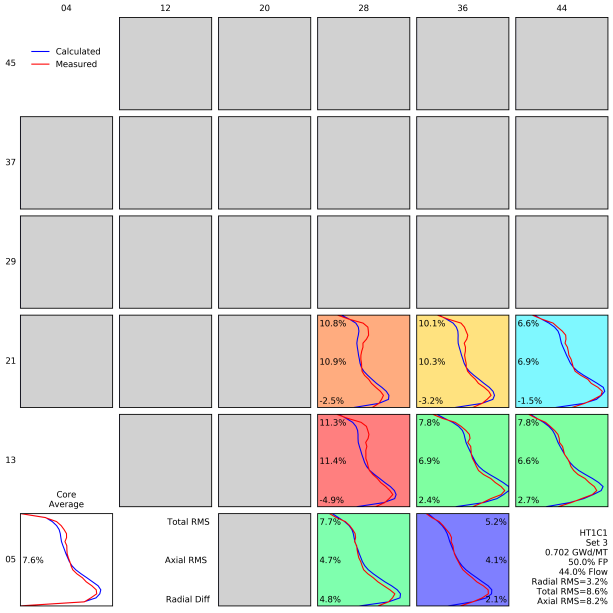

(a) Map 1: 0.702 GWDMT

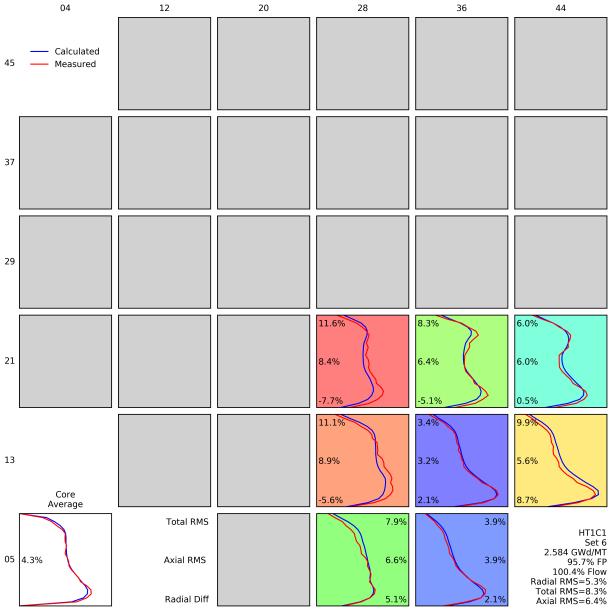

(b) Map 3: 2.584 GWDMT

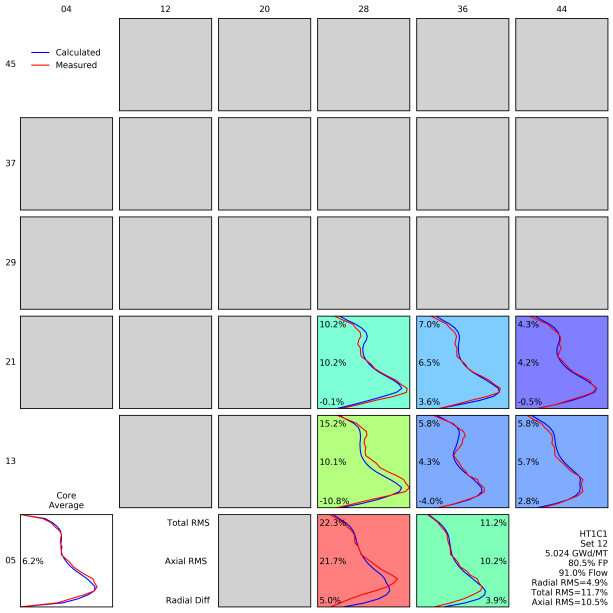

(c) Map 9: 5.024 GWDMT

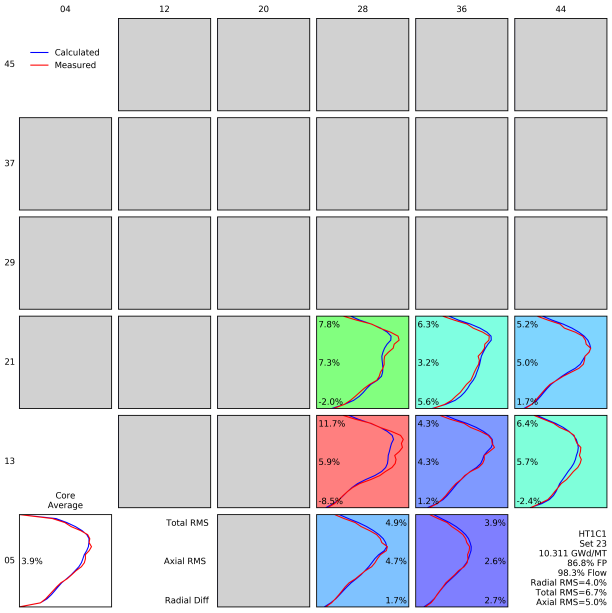

(d) Map 20: 10.311 GWDMT

Figure 5. Selected detector measurement maps for Hatch unit 1 cycle 1 (H1C1) with explicit heating + 3% calculation

2.3.5 Explicit Heating, No Conduction

To further test the explicit heating, another calculation was performed using no additional fixed heating and disabling the channel box conduction. The results for this run are shown in Figure 6. The early flux maps look like this approach may accomplish its goal. The first state point looks much better than for the 1% heating, matching well with the explicit heating. The 5.024 GWDMT map looks similar to the explicit heating map; this map combined with several other middle-of-cycle maps that are not shown here indicate that the power is not bottom-peaked enough, especially compared to the 1% heating. The end of cycle map is slightly worse than the 1% heating. Overall, the results for this approach show better average results across the cycle compared to the explicit heating. This indicates that more heat needs to be placed into the bypass, but that the current channel box conduction model is not doing that correctly.

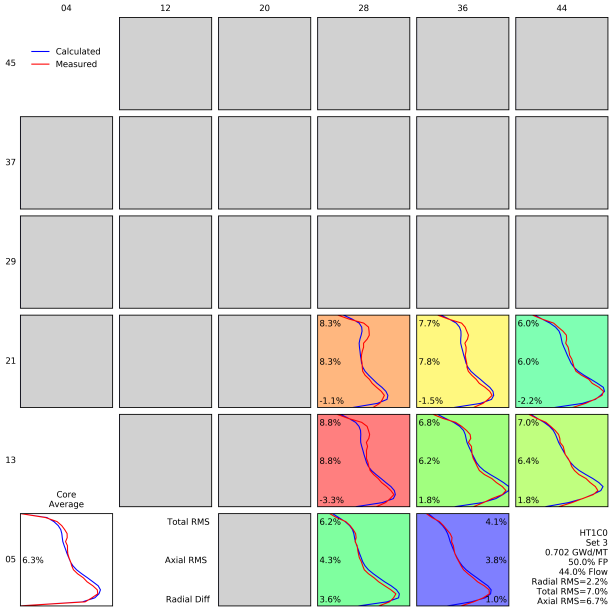

(a) Map 1: 0.702 GWDMT

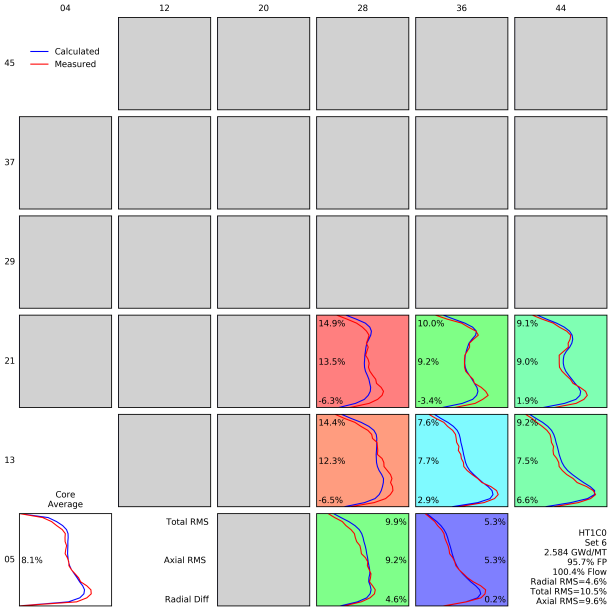

(b) Map 3: 2.584 GWDMT

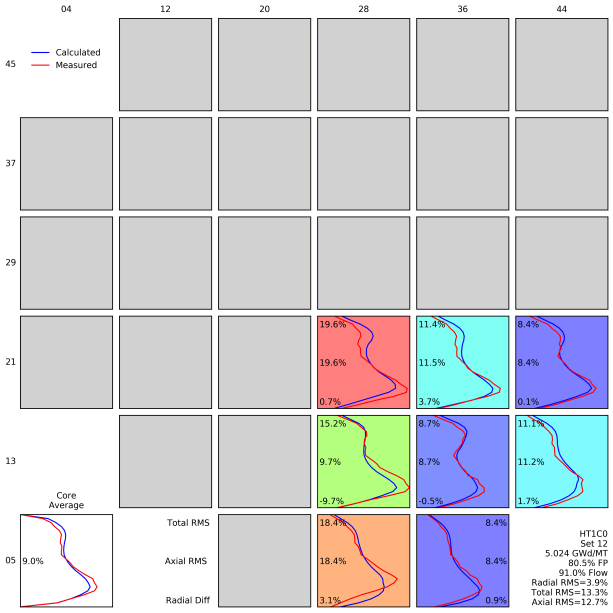

(c) Map 9: 5.024 GWDMT

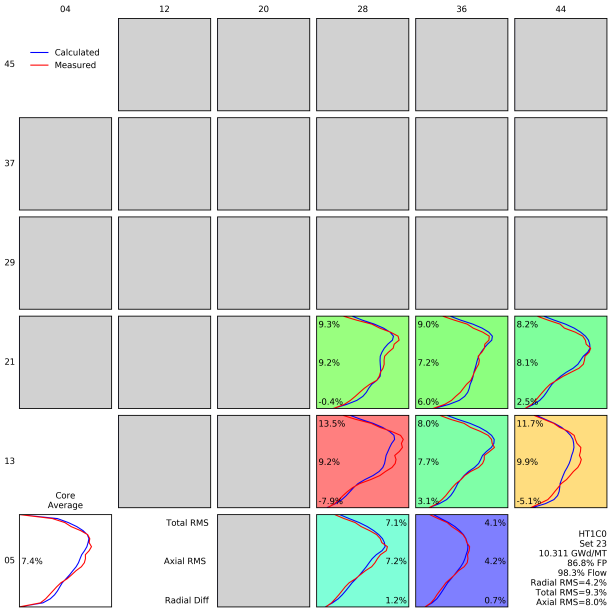

(d) Map 20: 10.311 GWDMT

Figure 6. Selected detector measurement maps for Hatch unit 1 cycle 1 (H1C1) with explicit heating with no channel box conduction

2.3.6 Hatch Results Summary

Figure 7 shows plots of k_{eff} over time for each of the Hatch simulations, and Table 1 shows the k_{eff} summary for the cases. Prediction of BWR k_{eff} is notoriously difficult compared to PWRs, as evidenced by the wide range of results shown in the plot. All cases are subcritical at most or all points. In general, the more heat removed to the bypass regions, the lower k_{eff} becomes. However, as shown by the detailed results and summarized later in this section, the lower amounts of bypass heating tend to give better predictions for the flux shape. Thus, the calculations which gave the best flux predictions also gave the best k_{eff} predictions as well, indicating that those calculations are doing the best job of accurately simulating the reactor.

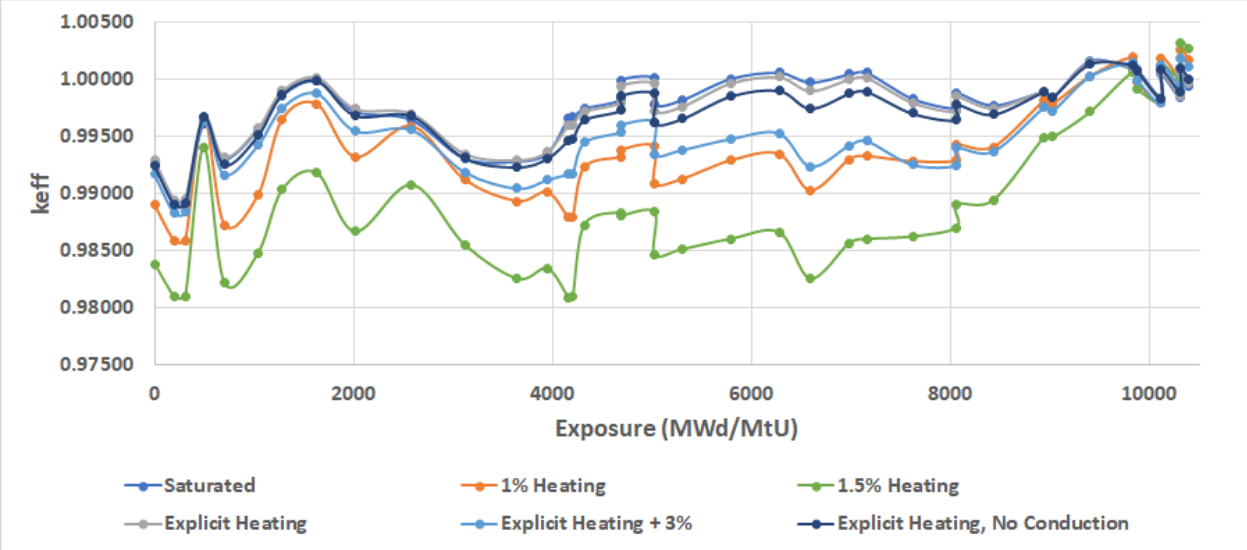

Figure 7. K_{eff} over time for various bypass heating treatments for Hatch unit 1 cycle 1 (H1C1)

Table 1. K_{eff} summary for various bypass heating treatments for Hatch unit 1 cycle 1 (H1C1); all results are pcm difference from critical; minimum and maximum errors are defined as being closest to and furthest from 1.0, respectively, regardless of sign

Case	Average	Std. Dev.	Min.	Max.
Saturated	271	305	5	-1085
1% Fixed Heating	645	454	17	-1419
1.5% Fixed Heating	1109	656	15	-1907
Explicit Heating	274	293	-1	-1057
Explicit Heating + 3%	515	348	-11	-1171
Explicit Heating, No Conduction	327	303	5	-1095

Table 2 shows summaries of the radial, axial, and total TIP detector response comparisons for various bypass models. Some additional models are shown here that were not shown in detail in previous sections for the sake of brevity, but the general trends are easily understood in light of what results were shown. In addition to the models which had detailed results in previous sections, 1.5% fixed heating is also shown in the summary to demonstrate that much more than 1% heating is too much. The first thing to note from the table is that the radial results are practically the same for all treatments, confirming that the primary difficulty for this model is in resolving the axial shape. Second, the axial shapes are significantly improved for the 1% fixed heating and the explicit heating + 3% case. The 1.5% fixed heating was not shown in detail because its results are

much worse, indicating that 1.5% is far too much heat in the bypass for this model. The explicit heating case is about the same as the saturated case; this indicates that either some physics is missing, the channel box conduction model is removing too much heat from the bypass, or both. This was confirmed by running an additional calculation that disabled the channel box conduction.

Table 2. Summary of Hatch unit 1 cycle 1 results for various bypass treatments; results are averaged over all 20 transverse in-core probe measurements

Case	Radial		Axial		Total	
	RMSE	**Max.**	**RMSE**	**Max.**	**RMSE**	**Max.**
Saturated	4.60%	8.09%	13.47%	18.91%	15.32%	20.49%
1% Fixed Heating	4.61%	8.63%	5.56%	10.95%	8.90%	13.52%
1.5% Fixed Heating	4.77%	9.07%	12.44%	18.64%	14.08%	19.11%
Explicit Heating	4.70%	8.61%	12.46%	17.69%	14.40%	19.45%
Explicit Heating + 3%	5.26%	10.18%	6.87%	12.53%	10.17%	15.81%
Explicit Heating, No Conduction	4.69%	8.73%	10.81%	15.93%	12.95%	17.84%

Figures 8, 9, and 10 show the RMSE and maximum errors for the radial, axial, and 3D TIP comparisons for each bypass treatment. As stated before, the axial and total plots look very similar because the effects of the coolant and moderator voiding are primarily axial. This is further confirmed by the fact that all 6 treatments look very similar for both radial plots. Overall, the 1% heating performs the, though it is one of the worst for the early power comparisons. It is also clear that the explicit heating treatment is much more similar to the saturated conditions than expected, further confirming that treatment needs additional work. This trend continues throughout the cycle, even if the channel box conduction is disabled, proving that some heat is not accounted for correctly.

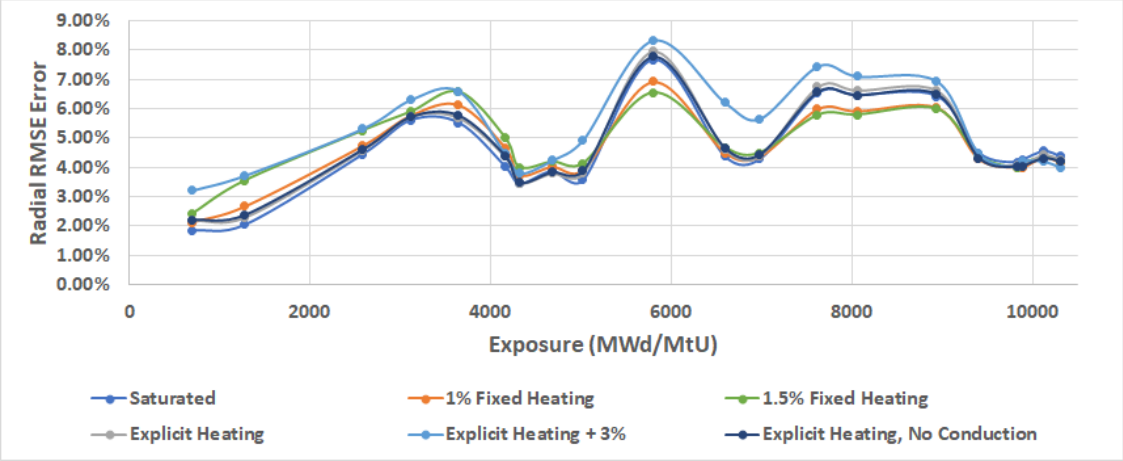

(a) root mean square error

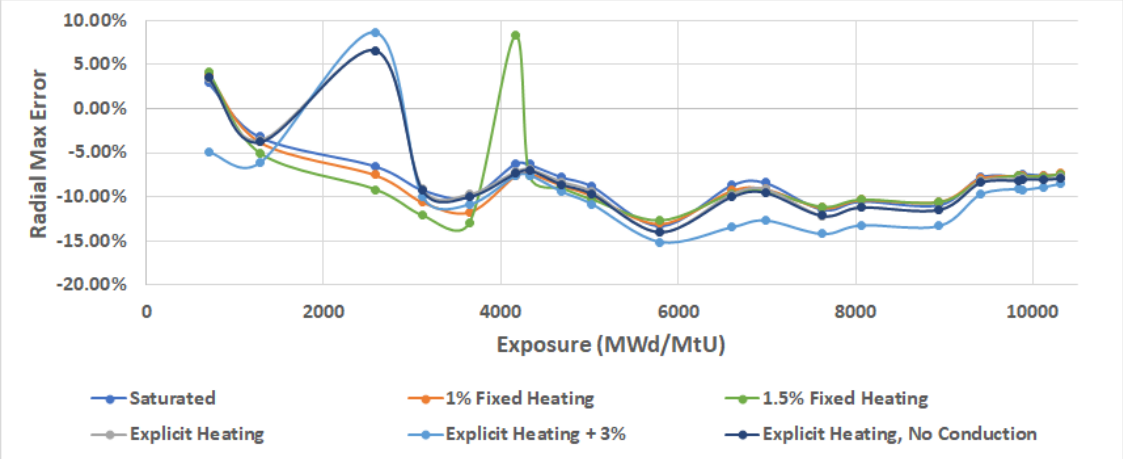

(b) Maximum

Figure 8. Radial TIP detector comparisons for various H1C1 bypass heating treatments

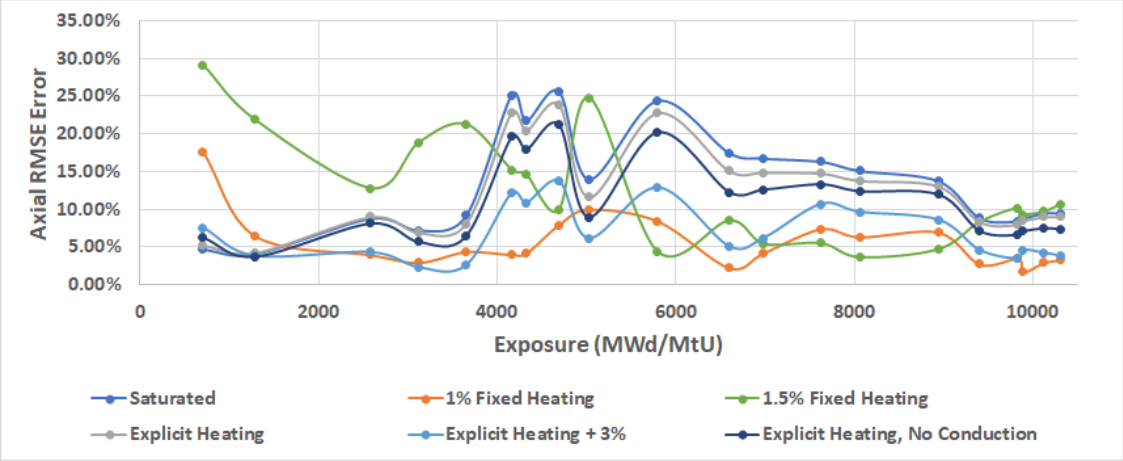

(a) root mean square error

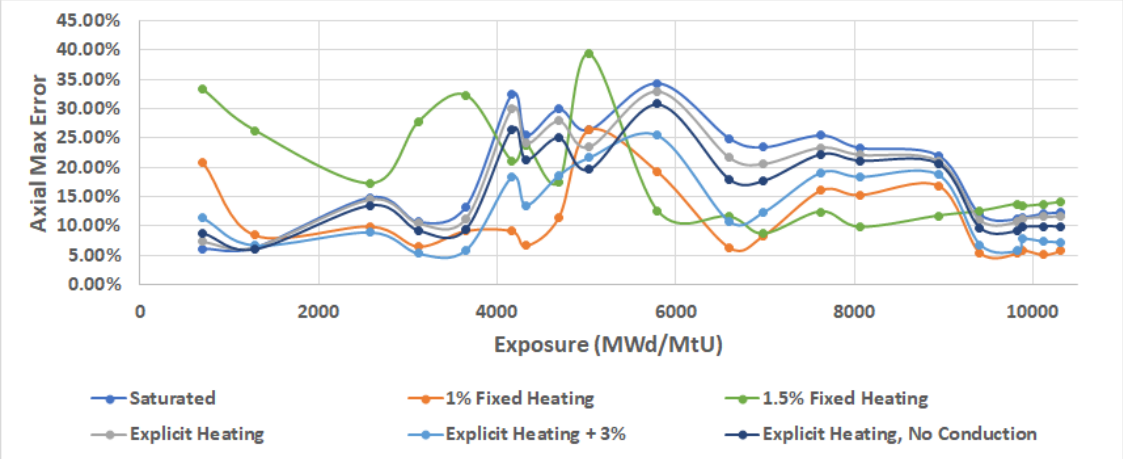

(b) Maximum

Figure 9. Axial TIP detector comparisons for various H1C1 bypass heating treatments

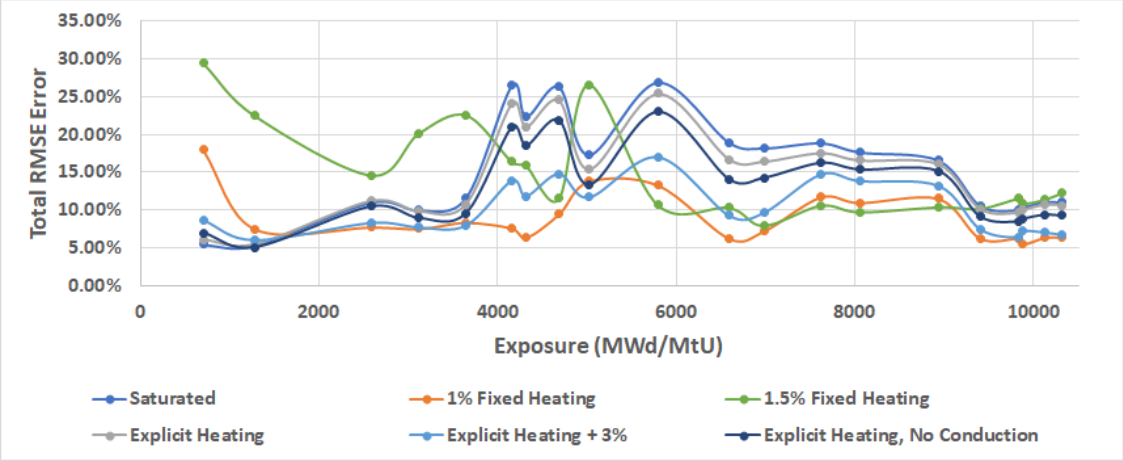

(a) root mean square error

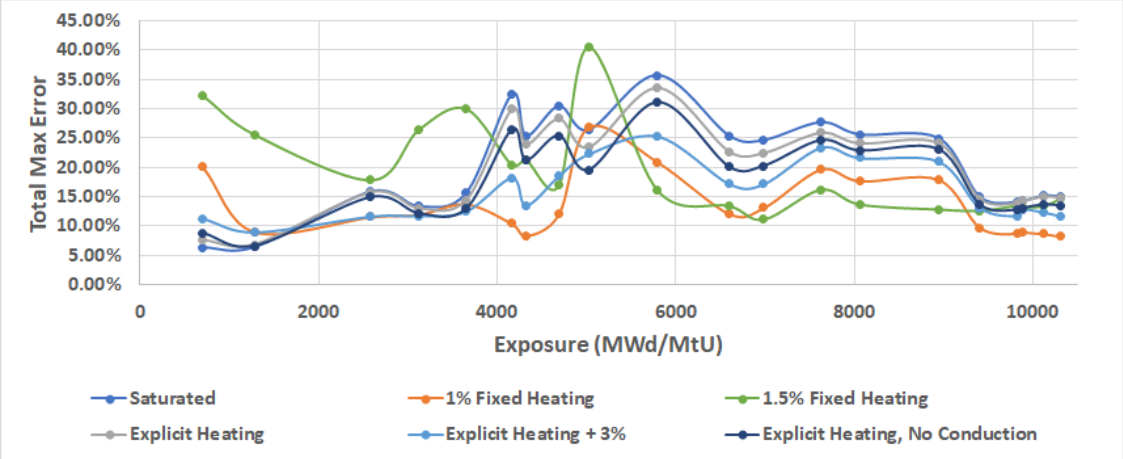

(b) Maximum

Figure 10. Total TIP detector comparisons for various H1C1 bypass heating treatments

2.4 HATCH TRANSIENT TEST

Before this work, VERA had never been used to run a transient simulation for a BWR model. Because of the anticipated focus on BWRs, it was important to assess this capability to see if any major obstacles were encountered. To test this, one of the end-of-cycle H1C1 models from the previous section was taken and restarted with a blade drop. The blade at 46-31 was set to drop from 20 notches removed, out of 48 total, to 48 notches removed in 1 second. The simulation was only run for a few hours, so the blade drop did not complete, but it made it through 0.3 seconds of the simulation, which is sufficient to assess if any major issues occurred. Figure 11 shows VERAView plots of the conditions at t=0 and t=0.3. In that time, the blade went from 213.36 cm to 152.4 cm; the axial slices in the figure are shown at 194.678 cm. It is clear from these images that the blade is moving at the expected rate and that the core power shape is responding to the movement of the blade. Follow-on work will perform more extensive testing, but this is sufficient to show that no major issues are present that must be addressed immediately specifically for BWRs transient calculations.

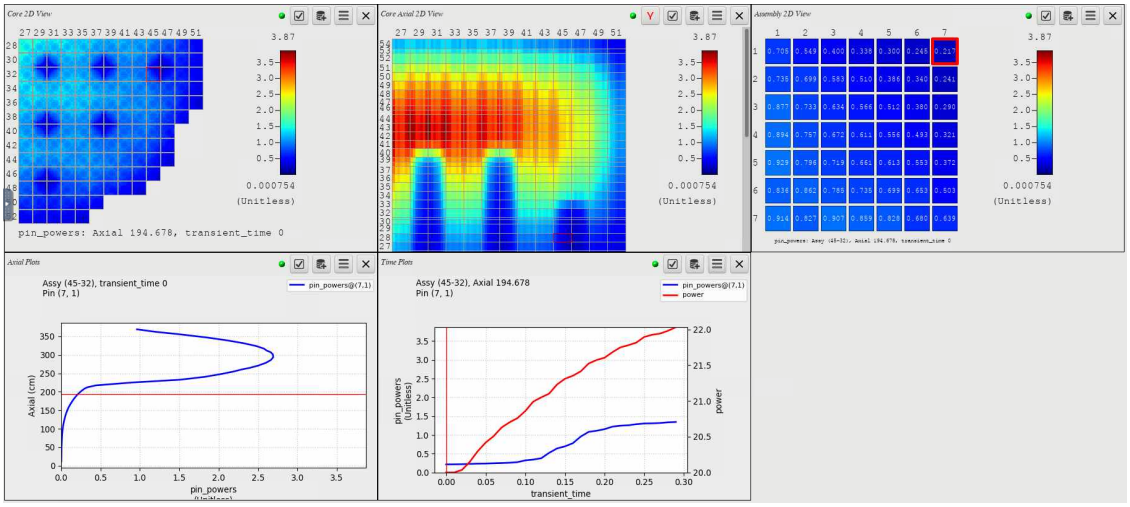

(a) t = 0.0 s

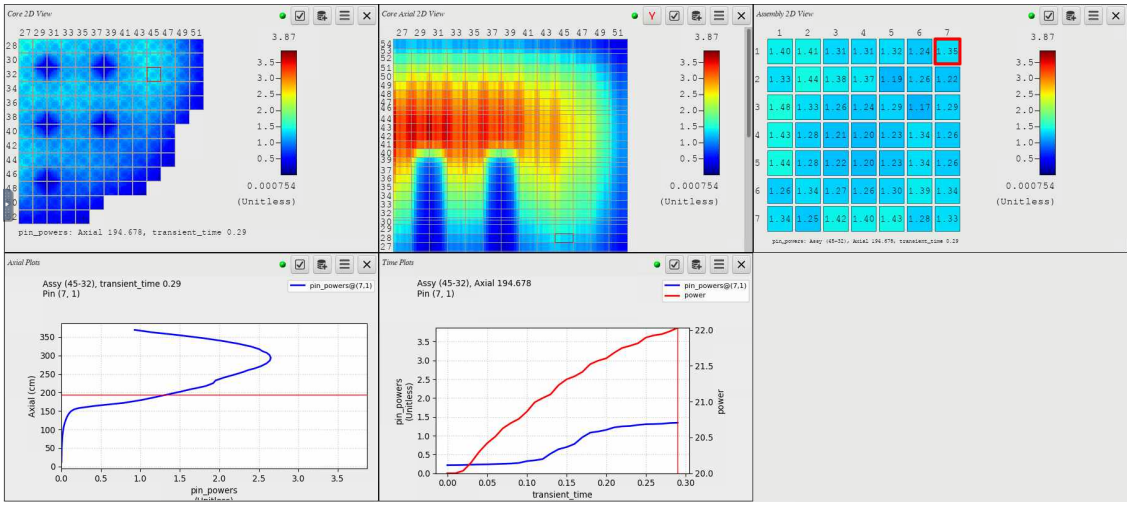

(b) t = 0.0 s

Figure 11. Results of Hatch unit 1 cycle 1 (H1C1) blade drop transient test

3. TRANSIENT IMPROVEMENTS

3.1 TRANSIENT XENON

One very important isotope for neutronics calculations is ^{135}Xe. This isotope has a large absorption cross section over one million barns in the thermal region. It has a half-life of approximately 9.2 hours. It is produced primarily from the β decay of ^{135}I (half-life 6.6 hours) with a small direct yield from fission. For depletion calculations, the depletion steps are much greater than the half-life of the two isotopes, so steady-state can be assumed to calculate the exact concentrations for a given power level. However, for transients, especially rapidly evolving superprompt transients, the rapid burning of the $_{135}$Xe during a power rise and subsequent build-up of $_{135}$I and $_{135}$ at high powers must be captured. This generally requires a special treatment of the differential equations describing these two isotopes.

The equations for the production of ^{135}Xe from ^{135}I are as follows:

$$\frac{dI}{dt} = \gamma_I \Sigma_F \phi(t) - \lambda_I I(t) \ , \tag{1}$$

$$\frac{dX}{dt} = \gamma_X \Sigma_F \phi(t) + \lambda_I I(t) - (\sigma_a \phi(t) + \lambda_X) X(t) \ , \tag{2}$$

where I and X refer to ^{135}I and ^{135}Xe, γ is the direct yield of an isotope from fission, Σ_f is the fission cross section, $\phi(t)$ is the scalar flux, λ is the decay constant, and σ_a is the microscopic absorption cross section for ^{135}Xe.

3.1.1 Equation Solution

First, the iodine equation can be solved using an integrating factor:

$$\frac{dI}{dt} + \lambda_I I(t) = \gamma_I \Sigma_F \phi(t)$$

$$\frac{d}{dt}\left[e^{\lambda_I t} I(t)\right] = e^{\lambda_I t} \gamma_I \Sigma_F \phi(t)$$

$$\int_{t_1}^{t_2} (\cdot) \Rightarrow e^{\lambda_I t_2} I(t_2) - e^{\lambda_I t_1} I(t_1) = \gamma_I \Sigma_F \int_{t_1}^{t_2} e^{\lambda_I t'} \phi(t')\, dt'$$

$$e^{\lambda_I t_2} I(t_2) = e^{\lambda_I t_1} I(t_1) + \gamma_I \Sigma_f \int_{t_1}^{t_2} e^{\lambda_I t'} \phi(t')\, dt'$$

$$I(t_2) = e^{\lambda_I (t_1 - t_2)} I(t_1) + \gamma_I \Sigma_f e^{-\lambda_I t_2} \int_{t_1}^{t_2} e^{\lambda_I t'} \phi(t')\, dt' \ .$$

It can be assumed that the flux is separable in space and time over short time steps:

$$\phi(t) \approx \overline{\phi} p(t) \ . \tag{3}$$

Substituting this expression into the equation for $I(t_2)$ gives the final solution for $I(t_2)$:

$$I(t_2) = e^{\lambda_I (t_1 - t_2)} I(t_1) + \gamma_I \Sigma_f \overline{\phi} e^{-\lambda_I t_2} \int_{t_1}^{t_2} e^{\lambda_I t'} p(t')\, dt' \ . \tag{4}$$

Everything in this expression is known except for the integral of the scalar flux shape function $p(t)$. This shape function is effectivley the shape of the core power change as a function of time. In the VERA iteration scheme, the next time step's power is calculated before feedback operations, and therefore the integral of $p(t)$ can be numerically evaluated at the start of the calculation. Thus, $I(t_2)$ is known if beginning-of-timestep values are used for everything else in Equation 4. Note that the approximation in Equation 3 is not necessarily a good one; this will be discussed in greater detail later.

Now a similar procedure can be applied to the xenon equation:

$$\frac{dX}{dt} + (\sigma_a \phi(t) + \lambda_X) X(t) = \gamma_X \Sigma_F \phi(t) + \lambda_I I(t) \ . \tag{5}$$

Since the iodine concentration does not depend on the xenon concentration, the RHS of the equation is already known. An integrating factor can be applied to the left hand side as was done with iodine, except that the integrating factor will be more complex in this case.

$$\begin{aligned} v(t) &= \int (\sigma_a \phi(t') + \lambda_X)\, dt' \\ &= \sigma_a \overline{\phi} \int p(t')\, dt' + \lambda_x t \\ &= \sigma_a \overline{\phi} \left(b_p t + \frac{m_p}{2} t^2 \right) + \lambda_X t \ , \end{aligned} \tag{6}$$

where we have assumed linear behavior for the power: $p(t) = p_0 + p_1 t$, $b_p = p(0)$, and $m_p = \frac{p(t_2)-p(t_1)}{t_2-t_1}$.

$$\begin{aligned} \frac{d}{dt}\left[e^{v(t)} X(t)\right] &= e^{v(t)} \gamma_X \Sigma_F \phi(t) + e^{v(t)} \lambda_I I(t) \\ \int_{t_1}^{t_2} (\cdot) \Rightarrow e^{v(t_2)} X(t_2) - e^{v(t_1)} X(t_1) &= \gamma_x \Sigma_f \overline{\phi} \int_{t_1}^{t_2} e^{v(t')} p(t')\, dt' + \lambda_I \int_{t_1}^{t_2} e^{v(t')} I(t')\, dt' \\ e^{v(t_2)} X(t_2) &= e^{v(t_1)} X(t_1) + \gamma_x \Sigma_f \overline{\phi} \int_{t_1}^{t_2} e^{v(t')} p(t')\, dt' + \lambda_I \int_{t_1}^{t_2} e^{v(t')} I(t')\, dt' \\ X(t_2) &= e^{v(t_1)-v(t_2)} X(t_1) + \gamma_X \Sigma_f \overline{\phi} e^{-v(t_2)} \int_{t_1}^{t_2} e^{v(t')} p(t')\, dt' + \lambda_I e^{-v(t_2)} \int_{t_1}^{t_2} e^{v(t')} I(t')\, dt' \\ t_1 = 0 \Rightarrow v(t_1) = 0 \Rightarrow X(t_2) &= e^{-v(t_2)} \left(X(0) + \gamma_X \Sigma_f \overline{\phi} \int_0^{t_2} e^{v(t')} p(t')\, dt' + \lambda_I \int_0^{t_2} e^{v(t')} I(t')\, dt' \right) \\ X(t_2) &= e^{-\left(\sigma_a \overline{\phi}\left(\frac{p(t_2)+p(0)}{2}\right)+\lambda_X\right)t} \left(X(0) + \gamma_X \Sigma_f \overline{\phi} \int_0^{t_2} e^{v(t')} p(t')\, dt' + \lambda_I \int_0^{t_2} e^{v(t')} I(t')\, dt' \right) \end{aligned} \tag{7}$$

3.1.2 Implementation Notes

Equation 4 can be solved directly if the power $p(t')$ is assumed to be linear. Applying integration by parts to the integral yields a solution that can be easily implemented in the code.

Equation 7 is less straightforward. Everything on the right-hand side of the equation is known, but the integrals have highly complicated solutions, even assuming that $p(t')$ and $I(t')$ are linear. To resolve this

issue, the integrals can be numerically integrated instead. This is accomplished simply by using the trapezoidal rule with small intervals to resolve the exponential shape. Many exponential evaluations are required because of the small integration intervals, so the exponential function was tabulated and the resulting table in-lined to optimize the calculation. This reduced the time required to evaluate the expression while introducing negligible error. The same table was also used in the evaluation of Equation 4 even though it had a much small impact there.

The most important approximations made in the derivation of Equations 4 and 7 is that the flux shape is constant and the power is linear. This is mitigated by two things. First, fast transients in which the flux changes rapidly are only solvable by VERA when using small timesteps due to numerical stability issues. Second, VERA's transient multilevel (TML) capability uses a point kinetics equations (PKE) solver to calculate the core power on a much finer timescale than what is used by the transport calculations. These two both work to reduce the timestep over which the ^{135}Xe calculation is performed, reducing the error caused by the linear approximations.

3.2 DEPLETION ENHANCEMENTS

As part of work on a previous milestone [10], VERA was modified to perform depletion calculations during transients. One side effect of the transient xenon implementation discussed in the previous section was that a defect in this depletion implementation was discovered. It was found that the depletion solver was using the same value for the current and previous power. While this likely did not have major effects on the results in the previous milestone, it was nonetheless incorrect. When power was increasing, the fuel was depleted too quickly; when the power was decreasing, the fuel was depleted too slowly. With the defect fixed, the average of the previous and current timesteps' power is used, which is a better approximation of the depletion. This fix will have small impacts on the decay heat source, which is important to accurately quantify for fuel performance and transient analysis. This defect was resolved and documented in a GitLab issue.

3.3 PARALLEL PERFORMANCE

Recently, issues have been identified in the parallel performance of the VERA neutron transport code MPACT due to poor parallel distribution of the workload. In some cases, certain parallel domains can have more than twice the work of others, resulting in significant runtime penalties overall. The partitioning for the first plane of VERA progression problem 9 [11], which is simply a model of the Watts Bar Unit 1 PWR cycle 1, shows the following partition data using the default partitioning scheme in VERA:

Partition Index	# Modules	# source regions
1	20	40500
2	22	44550
3	18	36450
4	14	28350
5	23	46575
6	10	20250
7	16	32400
8	14	28350
9	18	36450
10	10	20250
11	14	28350
12	14	28350
13	18	36450

14	10	20250
15	15	30375
16	21	42525

There are two major components to the MPACT solve: (1) the method of characteristics, which will scale approximately linearly in the number of source regions, and (2) coarse mesh finite difference, which scales approximately cubicly with the number of modules. As shown in the output, the number of modules ranges from 10 to 23, and the number of source regions is consistently about 2,000 regions per module. Thus, both portions of the solve are imbalanced with the current scheme. Similar results are obtained for other cores, including BWRs.

Work has begun on developing a new partitioning method based on an adjacency matrix. This uses information about how the modules are connected to partition the problem into partitions that are well-balanced in terms of the number of modules. This work was begun close to the due date for this report and is not yet complete. It is expected that it will be finished shortly after this report is submitted, making it available for FY24 and following needs.

3.4 MEMORY IMPROVEMENTS

While working on improving the BWR accuracy, it was discovered that the memory usage for BWR models was higher than anticipated; further testing revealed that this was also true to a lesser degree for PWR models. The neutron transport solver for VERA, MPACT [12], uses modular ray tracing to solve neutron transport problems with a high degree of accuracy. The purpose behind the modular ray tracing is to store one set of ray tracing data for multiple regions of the model that have the same mesh. These ray tracing modules can then be connected into the full-length ray, but the memory requirements are much lower because only one copy of the data is stored.

It was discovered that MPACT's checks used to determine whether two ray tracing modules are the same was far too strict: the material compositions were checked for consistency also, even though this sameness is not required because the ray tracing is performed independent of material properties. Thus, new checks were implemented that specifically checked only the data that impact the ray tracing. In the test cases reported on the GitLab Issue, up to 33% too much memory was consumed, which could cause the code to require more compute nodes than necessary. The new procedure for checking the modules enables the expected memory usage. This will significantly help with running coupled whole-core models, especially for BWRs.

3.5 DECAY HEAT COUPLING IMPROVEMENTS

3.5.1 Flux Normalization

In a previous NEAMS milestone report, work to couple decay heat calculations with TH calculations was performed to drive LOCA calculations using pin-by-pin decay heat [10]. One shortcoming identified during this work was that the flux normalization for coupling the neutron transport to other physics was based on total heat produced from fission. No mechanism existed in the code at that time to normalize based solely on prompt heat production so that the decay heat could be tracked separately.

To address this, a script was developed to scrape the various components of fission heat from the ENDF/B-VII.1 data [13] for the MT=458 channel, typically notated κ. The κs corresponding to heat produced from fission products, prompt neutrons, prompt gammas, and neutron captures were implemented in VERA as κ_{prompt} alongside the previous κ_{ER}, which contained all recoverable energy. This was done for every isotope in the cross section library.

Proper use of κ_{prompt} required some investigation. Ideally, it would be used everywhere when the decay heat calculations are enabled; unfortunately, the depletion and decay heat calculations themselves require

a normalized flux before obtaining a solution. Thus, those calculations must use κ_{ER} to normalize the flux to the total energy produced. After the depletion calculation is completed, κ_{prompt} and the decay heat calculations can then be used to separate the heating into the prompt and delayed components. These components are then used in the TH and transient portions of the calculations to obtain the proper feedback. The previous milestone report's results are improved and extended in a journal article that is expected to be published in the near future.

3.5.2 Transport and Depletion Isotope Mapping Improvements

Because decay heat has only recently become a focus for VERA applications, most VERA calculations have used a reduced depletion chain that consists of only a few hundred isotopes; in contrast, the full chain consists of approximately 2,200 isotopes. The few hundred isotopes were selected to encompass those important to reactivity and power distribution. Likewise, the transport library used for cross sections contains only a few hundred isotopes that are important for neutron transport effects, neglecting secondary effects such as decay heat. Additionally, many elements are grouped as a single isotope using their natural abundances instead of representing each constituent isotope separately.

For decay heat calculations, the full depletion chain must be used to obtain the correct decay heat. However, because decay heat was not a concern until recently, the mapping between the grouped isotopes in the cross section library and the individual isotopes in the depletion chain was never developed. Moreover, the cross section has multiple versions of certain isotopes whose cross sections are impacted by chemical bonds, such as ^{16}O bound to UO_2 or ^{12}C in graphite. These special cross sections were also not mapped correctly.

To resolve these issues, a new data structure was implemented in VERA to support the mapping between two sets of isotopes. This improved the depletion calculation by ensuring that the initial concentrations of each isotope correctly matched the VERA model, and it helped the rest of the VERA calculation by ensuring that the results of the depletion calculation were properly applied. The results of this improved mapping are documented in the GitLab issue and related merge requests. To summarize the results, the impacts on reactivity were nearly indiscernible except in the highest burnup test cases; even these cases saw less than a 10 pcm change in k_{eff}. However, tests that involved decay heat saw more than a 2% increase in the decay heat in some cases because of the increased concentrations of certain fission products. This could be important to capture for FFRD analysis and, at the very least, serves to eliminate a potential source of uncertainty in the calculations.

4. VERAONEWAY ASSESSMENT AND IMPROVEMENTS

4.1 CURRENT CAPABILITY ASSESSMENT

VERAOneWay strives to provide one-way coupling between the VERA core simulator and fuel performance codes, predominately BISON. The coupling can be accomplished across multiple cycles which may or may not include shuffling and/or depletion, as well as symmetric transients. However, the initial design was tailored to normal PWR operating conditions with transient processing being implemented late in the design. As a result, there are several lingering assumptions in the implementation which impose restrictions not found in VERA or BISON.

In the case of HBu transients, it is often desirable to run multiple quarter symmetry depletion cases with shuffling to establish a stable core design at the desired burnup. VERA then enables the user to utilize this stable core for a variety of transients, including asymmetric transients like rod ejection. Despite this being a valid VERA workflow, VERAOneWay does not handle going from full symmetry to quarter symmetry correctly and instead assumes that all feeder cycles match the symmetry of the final cycle. When this is not the case, undefined behavior may result.

Secondary to incorrect symmetry handling, VERAOneWay sometimes unexpectedly overwrites user supplied template values. The exact cause of this incorrect behavior is not well understood, nor are the issues well documented. Finally, there are quality of life issues with user input, and the performance on large cases does not meet expectations.

4.2 IMPROVEMENTS

An attempt at fixing the symmetry handling revealed that the constant symmetry assumption was widespread. Moreover, data encapsulation was nearly non-existent which resulted in significant ambiguity in the calculation of important quantities. This allowed for things like a constant symmetry assumption to be baked into specific routines with no realistic way of removing them after the fact.

A new unit test driven design has been proposed which respects separation of concerns and encapsulation. The main effort is separating the ingestion of VERA data from the translation of said data to VERAOneWay data. For clarity, VERAOneWay data could be a manipulation of VERA data, a view on VERA data, or something internally synthesized from available inputs. Such a separation will in a way provide a sort of VERA data application programming interface (API) that can hide a myriad of details making it easier to utilize the data correctly.

VERAOneWay, using the VERA data API, can deduce discharged assemblies utilizing the labels of actively modeled assemblies and the origin labels of assemblies shuffled into a given cycle. Once this deduction is complete VERAOneWay can construct entire histories and request required data from the API. In turn, the API can handle symmetry, among other details, internally and completely transparent to VERAOneWay.

Future efforts can focus on extending the API as needed and performance benchmark the overall process. Future work might also include reworking the methods used for constructing BISON inputs from the gathered data. There are active issues open in this area of concern, but they have not been evaluated in detail given a lack of open inputs demonstrating the error. Updating the VERA Hierarchical Data Format 5 (HDF5) output specification could be enabling for significant parts of this and other similar processes. All future efforts should include a substantial effort investment in unit testing, enabling expedited diagnosis and correction of semantic inconsistencies and ensuring accuracy of any future extensions.

4.3 SOFTWARE REQUIREMENTS DEVELOPMENT

A statistical down-selection of important fuel rods was investigated by Hirschhorn et al. [3]. Those findings demonstrated the necessity of down-selection because of the expense of repeatedly running the fuel performance calculations. The initial statistical approaches developed in that work will be implemented into VERAOneWay to streamline the analysis workflow by generating BISON inputs only for rods of interest based on the user-specified selection criteria. In light of this, a fresh set of software requirements was developed to clarify the exact set of capabilities that VERAOneWay should be equipped with. These requirements will guide future development and maintenance efforts for the package. They are documented in the GitLab issue. The current draft is reproduced below.

```
1. VERAOneWay (VOW) should generate fuel performance inputs based on VERA
power and burnup profiles for one or more VERA input/output files
   * The `only_cycle` input in the `[VERAONEWAY]` block indicates which
   cycle the inputs should be generated for.  If this is provided, then
   all rods that were present in that input will be generated.  If it's
   not present, then all rods in all inputs specified in the `[VERAONEWAY]`
   block will be generated.
   * The `only_assemblies` input in the `[VERAONEWAY]` block indicates which
   assembly positions each cycle should have inputs generated.  If not
   present, all rods will be generated.
   * The fuel performance input files will be generated for the code specified
   in the `fuel_code` of the `[VERAONEWAY]` block (currently `bison` or `fast`
   are accepted)
2. The files should follow a predictable naming convention:
`{casename}_C{cycle #}_A{assembly location}_P{pin X}-{pin Y}.bison.i`
   * `{cycle #}` is the number for the cycle from which the rod was discharged;
   if the cycle number cannot be determined from the `.xml` and `.h5` files,
   then the earliest cycle should be called cycle `1`.
   * `{assembly location}` is the `X-Y` location of the assembly using the
   `xlabel` and `ylabel` values from the final cycle
   * `{pin X}` is the X location of the pin in the assembly; this matches the
   indexing used in the VERA .h5 file for pin-wise datasets
   * `{pin Y}` is the Y location of the pin in the assembly; this matches the
   indexing used in the VERA .h5 file for pin-wise datasets
3. VERAOneWay (VOW) should track the rod histories in sequential order through
all VERA input/output files supplied by the user
4. The `-c` flag takes a "case name"; the file `{casename}.xml` must exist
5. The `{casename}.xml` file must have a fuel performance block in it (e.g.
`[BISON]`, `[FAST]`, etc.) with the following inputs:
   * `power_file` lists the VERA `.h5` files to be used, in order from oldest
   to newest
   * `cycle_xml` lists the VERA `.xml` files to be used; each one corresponds
   to an input in `power_file` and they should be listed in the same order
   * `shuffle_xml` lists the `.xml` files to use for shuffling between each
   pair of `.xml`; the number of inputs must be one less than the number of
   `cycle_xml` inputs; if no `shuffle_label` is given in one input value to the
   `cycle_xml` input, then it is assumed no shuffling took place at that point
   in time
```

   * ‘fuel_pin_input_file_template‘ gives the name of the fuel performance input file template that VOW should populate for each rod
6. VERAOneWay should support mapping between quarter and full symmetry calculations
   * When shuffling quarter->quarter, full->full, or full->quarter, nothing special needs to be done
   * When shuffle from quarter->full symmetry, the quarter core rod map should be unfolded to full symmetry, properly accounting for the ‘sym_bc‘ input in the quarter core ‘.xml‘ file
7. VERAOneWay must be able to support optional and required replacement of parameters in the fuel performance template files with values from the VERA inputs and outputs
8. VERAOneWay should have the capability to down-select rod input generation based on certain parameters and distributions
   * Must be able to look at ‘pin_*‘ data such as ‘pin_exposures‘, ‘pin_powers‘, ‘pin_decay_heat‘, etc.
   * Must be able to consider end-of-cycle, cycle-integrated, or cycle-maximum values
   * Must be able to consider rod-averaged, rod-maximum, or rod-minimum values
   * Should be able to explicitly include or exclude rods based on any of the above filters

5. CONCLUSIONS AND FUTURE WORK

Several important improvements have been documented in this report to support future fuel performance analysis. Defect fixes and multiphysics enhancements were implemented in VERA to improve BWR modeling capabilities; some validation was also done using the Hatch BWR cycle 1 model to determine the current status of VERA BWR capabilities. Several improvements have been made to VERA performance and transient capabilities. Finally, a substantial refactor was undertaken in VERAOneWay to resolve numerous issues with that code; effort was also made for planning future work on VERAOneWay since there is still much work that needs to be done for it to meet the future needs of the NEAMS program.

During this work, several future work items were also identified, which will be briefly summarized here and pursued in future work packages:

- Additional BWR improvements
  - Investigation of problems with explicit heat calculations
  - Implementation of fixed heating fraction for water rods
  - Extension of gamma transport to 3D
  - Coupling of gamma transport with thermal hydraulics calculations
  - Robustness improvements to BWR capabilities; certain combinations of bypass heating parameters were observed to cause the coupled calculation to become unstable
- Transient improvements
  - Automated calculation of time step sizes
  - Higher order shape functions in the transient xenon calculation
  - Completion of new partitioning scheme to improve performance
  - Other transient performance improvements, such as enabling multilevel coarse mesh finite difference (CMFD) for transient calculations
- VERAOneWay defect fixes and improvements
  - Resolve issues with junk values in BISON inputs
  - Resolve issues with overwriting user-specified BISON values
  - Improve input file generation performance
  - Develop new capabilities to selectively generate BISON input based on rod history statistics

6. ACKNOWLEDGMENTS

This research was funded by the Department of Energy Office of Nuclear Energy's Advanced Modeling and Simulation (NEAMS) program.

This research made use of the resources of the High Performance Computing Center at Idaho National Laboratory, which is supported by the Office of Nuclear Energy of the US Department of Energy and the Nuclear Science User Facilities under Contract No. DE-AC07-05ID14517.

REFERENCES

[1] Nathan Capps, Colby Jensen, Fabiola Cappia, Jason Harp, Kurt Terrani, Nicolas Woolstenhulme, and Daniel M. Wachs. A critical review of high burnup fuel fragmentation, relocation, and dispersal under loss-of-coolant accident conditions. *Journal of Nuclear Materials*, 546(1), 2021.

[2] Ian Greenquist, Aaron Wysocki, Jake Hirschhorn, and Nathan Capps. Multiphysics Analysis of Fuel Fragmentation, Relocation, and Dispersal Susceptibility–Part 1: Overview and Code Coupling Strategies, March 2023.

[3] Jake Hirschhorn, Ian Greenquist, Aaron Wysocki, and Nathan Capps. Multiphysics Analysis of Fuel Fragmentation, Relocation, and Dispersal Susceptibility–Part 2: High-Burnup Steady-State Operating and Fuel Performance Conditions, March 2023.

[4] Aaron Wysocki, Jake Hirschhorn, Ian Greenquist, and Nathan Capps. Multiphysics Analysis of Fuel Fragmentation, Relocation, and Dispersal Susceptibility–Part 3: Thermal Hydraulic Evaluation of Large Break Loca Under High-Burnup Conditions, March 2023.

[5] J. A. Turner, K. Clarno, M. Sieger, R. Bartlett, B. Collins, R. Pawlowski, R. Schmidt, and R. Summers. The Virtual Environment for Reactor Applications (VERA): Design and Architecture. *Journal of Computational Physics*, 326:544–568, 2016.

[6] S. M. Bajorek et al. Development, validation, and assessment of the TRACE thermal-hydraulics systems code. In *Proceedings of NURETH-16*, 2016.

[7] R. Williamson et al. Multidimensional Multiphysics Simulation of Nuclear Fuel Behavior. *Journal of Nuclear Materials*, 423(1–3):149–163, 2012.

[8] D. Gaston et al. MOOSE: A Parallel Computational Framework for Coupled Systems of Nonlinear Equations. *Nuclear Engineering and Design*, 239(10):1768–1778, 2009.

[9] Kyle A. Gamble, Larry K. Aagesen Jr., Antonio Martin Recuero, Jason D. Hales, Daniel J. Vanwasshenova, Nathan Capps, Ryan Sweet, Michael W. D. Cooper, Sudipta Biswas, and Wen Jiang. Advancements in modelin fuel pulverization and cladding behavior during a LOCA. Technical Report INL/EXT-21-64705-Rev001, Idaho National Laboratory, 2021.

[10] Aaron Graham and Andrew Godfrey. Coupled Decay Heat and Thermal Hydraulic Capability for Loss-of-Coolant Accident Simulations. Technical Report ORNL/TM-2023/2903, Oak Ridge National Laboratory, 2023.

[11] Andrew Godfrey. VERA Core Physics Benchmark Progression Problem Specifications. Technical Report CASL-U-2012-0131-004, Oak Ridge National Laboratory, 2014.

[12] B. Collins, S. Stimpson, B. W. Kelley, M. T. H. Young, B. Kochunas, A. Graham, E. W. Larsen, T. Downar, and A. Godfrey. Stability and Accuracy of Three-Dimensional Neutron Transport Simulations Using the 2D/1D Method in MPACT. *Journal of Computational Physics*, 326:612–628, 2016.

[13] M.B. Chadwick, M. Herman, P. Obložinský, M.E. Dunn, Y. Danon, A.C. Kahler, D.L. Smith, B. Pritychenko, G. Arbanas, R. Arcilla, R. Brewer, D.A. Brown, R. Capote, A.D. Carlson, Y.S. Cho, H. Derrien, K. Guber, G.M. Hale, S. Hoblit, S. Holloway, T.D. Johnson, T. Kawano, B.C. Kiedrowski, H. Kim, S. Kunieda, N.M. Larson, L. Leal, J.P. Lestone, R.C. Little, E.A. McCutchan, R.E. MacFarlane, M. MacInnes, C.M. Mattoon, R.D. McKnight, S.F. Mughabghab, G.P.A. Nobre, G. Palmiotti, A. Palumbo, M.T. Pigni, V.G. Pronyaev, R.O. Sayer, A.A. Sonzogni, N.C. Summers, P. Talou, I.J.

Thompson, A. Trkov, R.L. Vogt, S.C. van der Marck, A. Wallner, M.C. White, D. Wiarda, and P.G. Young. ENDF/B-VII.1 nuclear data for science and technology: Cross sections, covariances, fission product yields and decay data. *Nuclear Data Sheets*, 112(12):2887 – 2996, 2011. Special Issue on ENDF/B-VII.1 Library.