ORNL/TM-2023/2927

On-the-Fly Energy Condensation for Whole-Core Multiphysics Simulations

Aaron Graham
Kang Seog Kim

May 23, 2023

OAK RIDGE National Laboratory

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

ORNL/TM-2023/2927

Nuclear Energy and Fuel Cycle Division

ON-THE-FLY ENERGY CONDENSATION FOR WHOLE-CORE MULTIPHYSICS SIMULATIONS

Aaron Graham
Kang Seog Kim

Date Published: May 23, 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

ARI	all rods in
ARO	all rods out
BOC	beginning of cycle
BWR	boiling water reactor
C1	cycle 1
C2	cycle 2
CMFD	coarse mesh finite difference
CPM	collision probabilities method
CZP	cold zero power
EFPD	effective full-power days
EOC	end of cycle
FFRD	fuel fragmentation, relocation, and dispersion
gad	gadolinia-bearing
GMRES	generalized minimum residual
HBu	high-burnup
HFP	hot full power
HZP	hot zero power
IFBA	integral fuel burnable absorber
IRW	integral rod worth
ITC	isothermal temperature coefficient
LOCA	loss-of-coolant accident
LWR	light-water reactor
MC	Monte Carlo
MOC	method of characteristics
NEAMS	Nuclear Energy Advanced Modeling and Simulation
OTF	on-the-fly
PLR	part-length rods
PWR	pressurized water reactor
RIA	reactivity insertion accident
RMSE	root mean square error
SPH	SuPerHomogéneísation
TH	thermal hydraulic
VERA	Virtual Environment for Reactor Applications
WABA	wet annular burnable absorber
WB1	Watts Bar Unit 1
WSA	Wigner-Seitz approximation

1. INTRODUCTION

As the nuclear energy industry considers ways to achieve improved economics in the current fleet of light-water reactors (LWRs), one possible approach is to operate each cycle longer. This causes a greater portion of the fuel to be burned and reduces the frequency of outages, which ultimately reduces the cost to operate the reactor. However, this also leads to higher burnup fuels than have traditionally been allowed in these reactors. There are concerns about integrity of high-burnup (HBu) fuel, especially during accident conditions such as loss-of-coolant accidents (LOCAs), as shown by Capps et al. [1]. To investigate these concerns, advanced modeling and simulation capabilities are under development to determine the susceptibility of HBu fuel to fuel fragmentation, relocation, and dispersion (FFRD). Improvements have previously been made to fuel performance capabilities to model these phenomena more accurately [2]; multiphysics simulations have also been conducted to determine the power and burnup histories of the HBu fuel [3], which are needed as inputs to the fuel performance calculations. Most recently, new statistical approaches have been developed to identify a subset of fuel rods that have greater FFRD susceptibility, reducing the total number of fuel performance simulations required [4–6].

To properly assess FFRD susceptibility for a fuel rod, accurate power and burnup histories are required. This makes the Virtual Environment for Reactor Applications (VERA) code suite an ideal tool for core simulation because it generates rod-by-rod axial power and burnup distributions using state-of-the-art high-fidelity physics methods. However, the detailed nature of VERA calculations also means that it is slow compared to industry methods. Thus, efforts are currently being made to speed up VERA calculations without sacrificing their accuracy and fidelity.

This report focuses on the implementation of on-the-fly (OTF) energy group condensation in MPACT, the neutron transport module of VERA. Typical MPACT calculations use 51 energy groups [7]; there is also a 252-group library that can be used for calculations when the 51-group library may not be sufficient. The 51-group cross section library has been generated for LWR calculations and is extensively validated against a wide range of LWR problems. The OTF energy condensation presented in this report aims to reduce the number of groups significantly without sacrificing the accuracy of the 51-group library. This is done by generating the few-group cross sections from the 51-group cross sections at runtime so that the exact problem of interest is used for the cross section generation. This eliminates a typical source of error in cross section generation in which representative problems are used that have some minor differences compared with the actual problems in which the cross sections are used. Because MPACT is ultimately coupled to other physics codes as part of VERA, the cross sections have to be updated frequently as the conditions of the problem change. Thus, performance of the OTF methodology itself is crucially important to ensure that it does not cancel out performance improvements in the method of characteristics (MOC) and coarse mesh finite difference (CMFD) solvers.

2. ON-THE-FLY ENERGY CONDENSATION METHODOLOGY

2.1 VERA ITERATION SCHEME

MPACT uses the MOC, a geometry-agnostic method useful for resolving the complex heterogeneity of tens of thousands of fuel rods in LWRs, to solve the neutron transport equations. This method is applied in the radial direction and is coupled to a faster running, lower-order transport method in the axial direction—an approach called 2D/1D. The 2D/1D iteration scheme is then accelerated by the CMFD acceleration method. The MOC calculations are the most expensive part of the calculation because they require sweeping the length of a very large number of rays drawn across the geometry. The CMFD calculations are close behind because they involve solving a non-linear system of equations; the number of rows in the matrix for this system is equal to the number of coarse mesh cells multiplied by the number of energy groups. The VERA iteration scheme is illustrated in Figure 1.

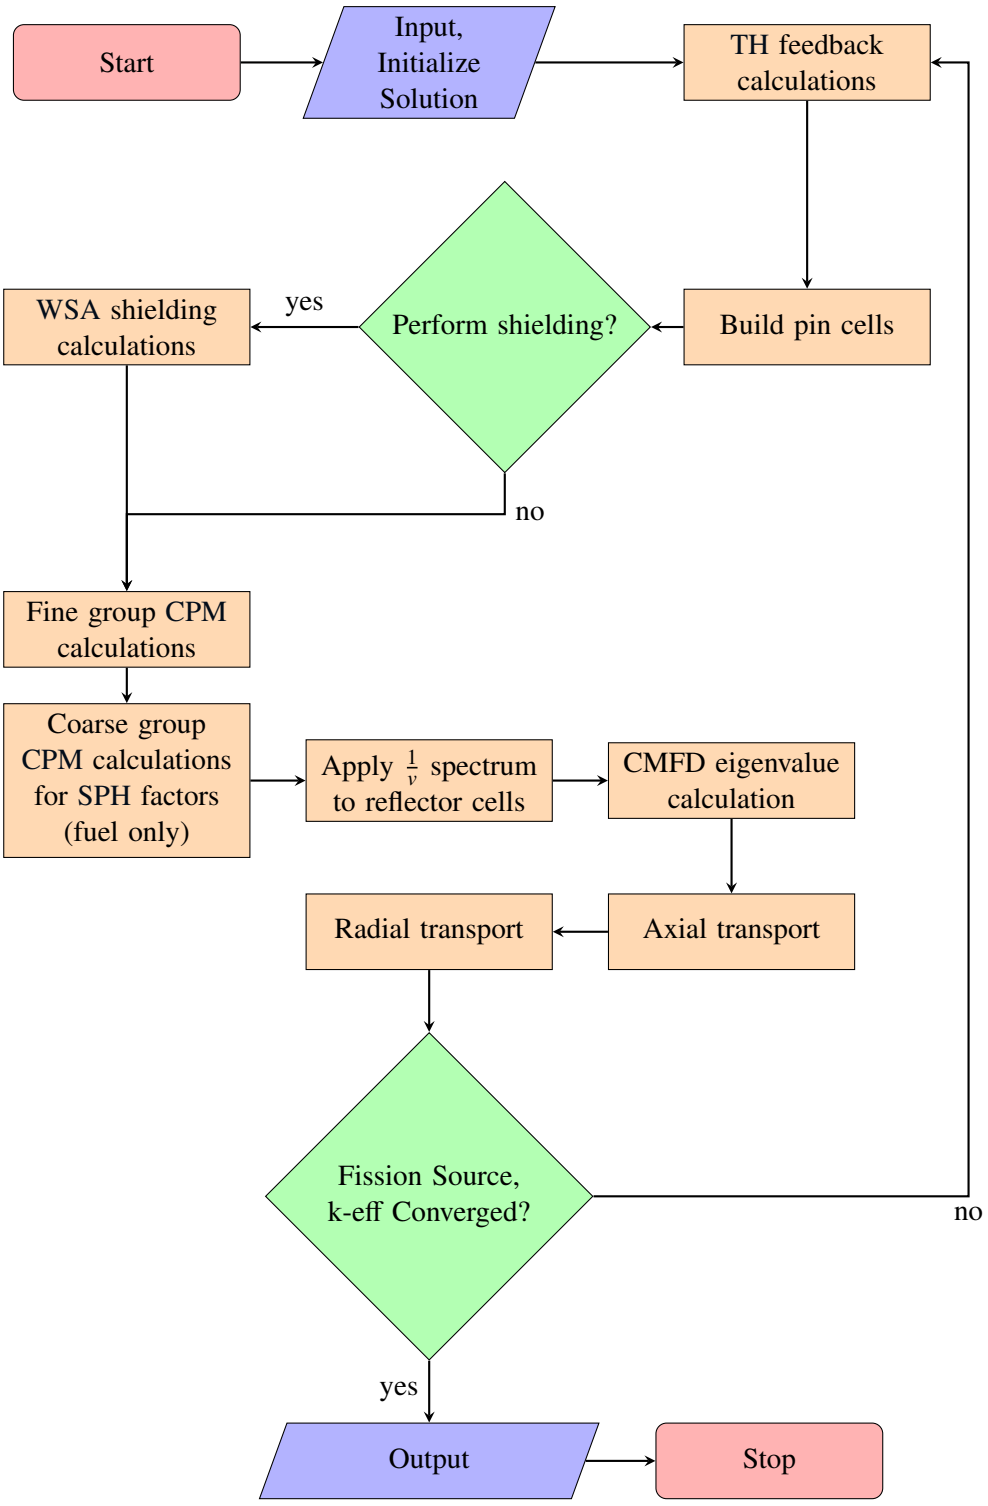

Figure 1. Illustration of VERA iteration scheme.

2.2 CONDENSATION PROCEDURE

To accurately condense cross sections, it is important that an accurate energy spectrum be used. Since this spectrum is not known prior to solving the problem, it must be accurately approximated. To do so, the 3D reactor is decomposed into constituent pin cells that can be solved efficiently. A fine-group spectrum is approximated for each region of each pin cell and is used to collapse the cross sections for each region. These collapsed cross sections are then used for the whole-core transport calculations.

2.3 PIN CELL IDENTIFICATION

Several categories of pin cells can be identified:

- fuel pins,
- gadolinia-bearing (gad) fuel pins,
- guide tubes,
- control rods,
- discrete burnable poison inserts,
- control blades
- cells comprising the channel box, bypass flow, and control blade in boiling water reactors (BWRs), and
- everything else (including axial and radial reflector cells).

Cells are grouped into these categories because some of them require special treatments. These treatments are discussed in the sections in which they are applicable.

2.4 WIGNER-SEITZ APPROXIMATION FOR SELF-SHIELDING CALCULATIONS

Before performing the spectrum calculation, spatial and energy self-shielding calculations must be performed using the fine-group structure. Otherwise, the shielding effects will not be treated in the spectrum calculation, resulting in a poor approximation of the fine-group spectrum. MPACT relies on the subgroup approach to shielding calculations. Historically, it performed subgroup calculations for each category of isotopes in each resonance group. Recently, the Wigner-Seitz approximation (WSA) approximation was implemented to accelerate the subgroup calculations [8]. In this approach, the core is decomposed into pin cells, as described in the previous section. For each pin cell category, a one-group transport calculation is performed to obtain a Dancoff factor for each pin cell and each subgroup isotope category. The collision probabilities method (CPM) is then used on each pin cell to determine a Dancoff factor for the isolated pin cell. The outer radius of the pin cell is then adjusted and the CPM calculation repeats until the isolated pin cell Dancoff factor equals that obtained from the whole-core calculation. Once this iteration scheme is complete, the new Dancoff-equivalent cell can be used to calculate the spatial and energy self-shielding effects accurately for each region in the pin cell.

This WSA approach has two major benefits. First, it is more computationally efficient than performing whole-core shielding calculations. The 51-group library has 22 resonance groups that have subgroup levels and isotope categories for each resonance group. For a typical pressurized water reactor (PWR) calculation, this results in 268 “pseudogroups” that must be solved using MOC. For the WSA approach, this number is reduced to fewer than 20 for most problems. When using the 252-group library, there are 133 resonance groups that lead to more than 1,000 pseudogroups to solve. The CPM calculations require non-trivial amounts of time, but the overall shielding calculation is still significantly faster than the whole-core approach since the number of MOC calculations is no longer a function of the number of resonance groups. Second, the constituent pin cells required for the OTF energy condensation are already constructed for the WSA approach to shielding, allowing them to be re-used to generate an energy spectrum as well.

2.5 PIN CELL SPECTRUM CALCULATION

For the pin cell calculations, fine-group cross sections are calculated using the results of the WSA shielding calculations. These cross sections are used to set up a 1D CPM calculation. The CPM calculations are one-group calculations, so the groups must be swept starting with the fastest group and ending with the most thermal group. Additionally, the presence of up-scatter and more than one region in the pin cell requires iterating over all the groups to resolve the up-scattering source as well as the spatial distribution of the fission source. For each iteration over the groups, the total fission source in the pin cell is fixed to a constant value. This ensures that the calculation eventually converges even though the pin cell will likely not be perfectly critical. It should be noted that the spectrum calculation should use the original radius of the pin cell that preserves volume, not the Dancoff-equivalent radius. The results of the WSA calculations are incorporated into the fine-group cross sections, so there is no need to use the Dancoff-equivalent radius for the spectrum calculation too.

2.6 CROSS SECTION GENERATION

Once the fine-group spectrum ϕ_g is obtained for each region, the spectrum and the cross sections can be collapsed:

$$\phi_G = \sum_{g \in G} \phi_g \,, \tag{1a}$$

$$\Sigma_{x,G} = \frac{\sum_{g \in G} \phi_g \Sigma_{x,g}}{\phi_G} \,, \tag{1b}$$

$$\Sigma_{s,G \to G'} = \sum_{g' \in G'} \sum_{g \in G} \frac{\phi_g \Sigma_{s,g \to g'}}{\phi_G} \,, \tag{1c}$$

$$\chi_G = \sum_{g \in G} \chi_g \,, \tag{1d}$$

where G indicates a coarse group quantity, and $g \in G$ refers to all fine groups g that are inside the coarse group G.

2.7 SPH FACTORS

Energy condensation always results in some error in the collapsed cross sections if angle-dependent cross sections are not used [9, 10]. To mitigate these effects, SuPerHomogéneísation (SPH) factors are used. To calculate these factors, a second CPM calculation is set up for each pin cell, this time using the coarse group structure and cross sections. The converged source from the fine group CPM calculation is condensed to the coarse group structure and used to drive the coarse group CPM calculation. After the calculation is converged, an SPH factor is calculated:

$$f^k_{i,G} = \frac{\phi_{fine,i,G}}{\phi^k_{coarse,i,G}} \,, \tag{2}$$

where $\phi_{fine,i,G}$ is the condensed flux from the fine group CPM calculation in region i and coarse group G, and $\phi^k_{coarse,i,G}$ is the flux from the coarse group CPM calculation. The cross sections are updated using this factor:

$$\Sigma^k_{x,i,G} = f^k_{i,G} \Sigma^0_{x,i,G} \,, \tag{3}$$

where $\Sigma^0_{x,i,G}$ is calculated with Eqs. (1b) and (1c) and $f^0_{i,G} = 1$.

2.8 GEOMETRY CONSIDERATIONS

The 1D cylindrical CPM calculations fit naturally with LWR fuel rods. However, there are several regions that do not fit well in this type of calculation. First, the baffle and radial reflector regions are not well-suited for cylindrical calculations. For this reason, a $\frac{1}{\nu}$ spectrum is used in those regions instead of using the CPM calculations. The cylindrical CPM method was tested for them, but gave worse results than a simple predefined spectrum. Several other more complicated hard-coded spectra were attempted but did not make a significant difference.

Second, the channel box and bypass regions of BWRs do not fit well in cylindrical geometry. These regions are still solved using the cylindrical CPM calculations, but this represents an area for improvements. It is likely that a 1D slab CPM calculation would be a better approximation, but some other approach may be valid as well.

Third, the axial reflector regions do not necessarily fit well into the cylindrical approach. It is expected that this is less important than the radial reflectors. For the results in this report, the axial reflectors are treated with the same spectrum as the radial reflectors.

3. OTF ENERGY CONDENSATION VALIDATION

All results in this report are based on the VERA progression problems [11] and the BWR progression problems [12]. The VERA progression problems are well-known in the reactor modeling and simulation community. They present a set of problems ranging from a single 2D pin cell up to 3D multiphysics multi-cycle calculations. The results are broken up into several different logical sets for the purpose of testing the new OTF capability. Each set is discussed in its own section immediately preceding the results for those problems. This section focuses on the neutronics-only problems, in which the only difference between the old and new VERA results is the use of the OTF capability. Section 4 presents results for multiphysics problems, which also includes the use of the pin census capability.

The BWR progression problems were inspired by the VERA progression problems and follow a similar progression of increasing size and complexity for BWR calculations. Each of these problems is a neutronics-only calculation, with reference solutions available via Serpent2 [13]. Comparisons for the OTF capability are also presented for these BWR progression problems after the VERA progression problems.

3.1 51→19 OTF

The first set of results presented is based on the 51-energy group transport library, which is the default for VERA. The calculations were collapsed to 19 groups using a modified version of the 18-group library. The energy group structures are illustrated in Figure 2.

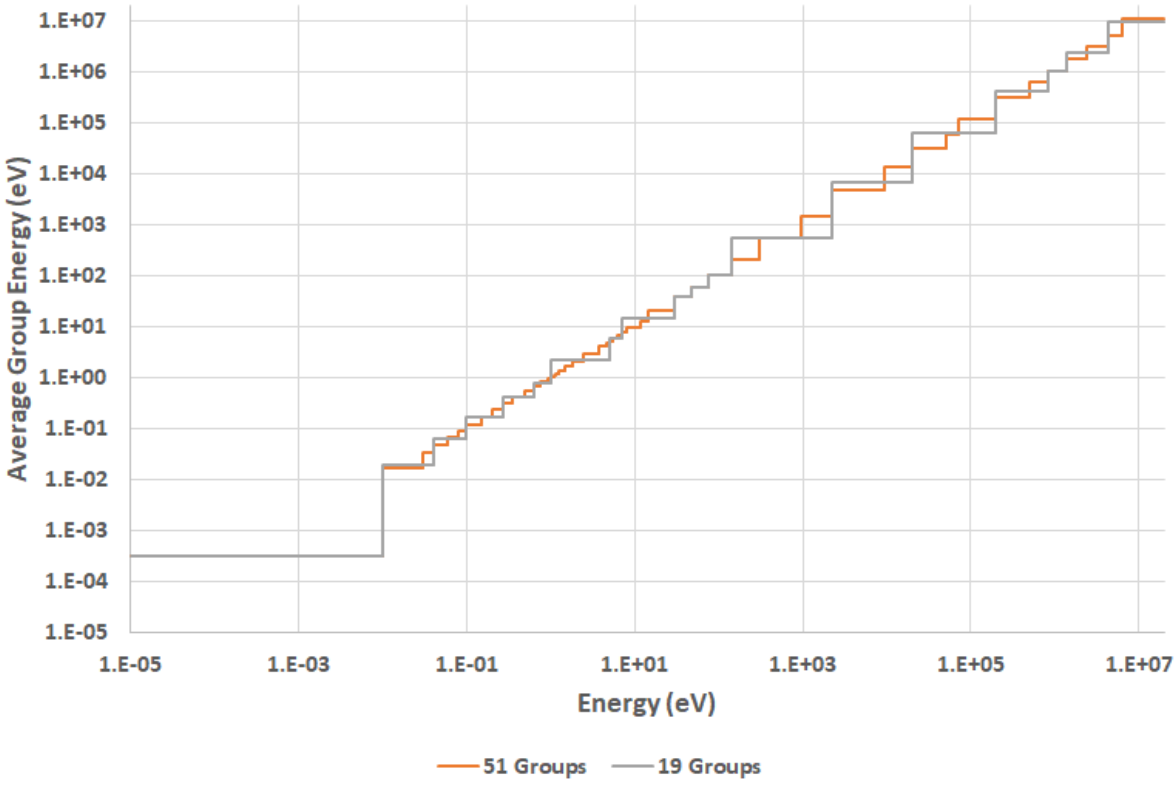

Figure 2. Average energy and energy boundaries of each group for the 51 and 19 group libraries.

3.1.1 2D PWR RESULTS

The first set of results covers problems 1a–1e, problems 2a–2q, problems 4a-2d–4c-2d, and problems 5a-2d–5c-2d. Most codes with OTF-like capabilities tend to solve 2D lattices and other similar 2D calculations to collapse cross sections for a nodal code, so these results are the first step in testing the OTF capability in MPACT, as they are well understood from other results. Additionally, all these problems have Monte Carlo (MC) results generated by the KENO-VI [14], so we can compare the OTF results not only to the fine-group transport but also to MC. These problems are briefly described in Table 1. The results are summarized in Table 2. Each set of problems is shown in more detail in the following sections.

Table 1. Summary of the 2D VERA progression problems

Problem	Variations	Description
1	5, a–e	Pin cells with various temperature conditions
2	17, a–q	Lattices at various conditions, combinations of burnable poisons, and control rods
3	2, a–b	3D assembly without (3a) and with (3b) pyrex burnable poisons
4-2D	3, a–c	3×3 array of assemblies with burnable poisons and 3 different control rod configurations (none in 4a-2d, AIC rods in 4b-2d, B_4C rods in 4c-2d)
5-2D	3, a–c	Variations of control rod configurations

Table 2. Summary of the VERA progression problem 51→19 OTF results compared to KENO-VI; all values are the average of |VERA − KENO-VI| over the set of problems

Problem	51 Groups			51→19 OTF		
	k_{eff} (pcm)	Power Distribution		k_{eff} (pcm)	Power Distribution	
		RMSE	Max.		RMSE	Max.
1	53 ± 97	—	—	140 ± 69	—	—
2	157 ± 128	0.001 ± 0.000	0.003 ± 0.001	217 ± 208	0.002 ± 0.001	0.004 ± 0.001
3	68 ± 56	0.002 ± 0.000	0.005 ± 0.001	103 ± 144	0.003 ± 0.000	0.008 ± 0.001
4-2D	86 ± 7	0.003 ± 0.001	0.006 ± 0.002	119 ± 18	0.003 ± 0.001	0.008 ± 0.001
4-3D	69 ± 4	0.006 ± 0.003	0.037 ± 0.023	113 ± 10	0.006 ± 0.002	0.038 ± 0.021
5-2D	75 ± 3	0.003 ± 0.001	0.011 ± 0.003	115 ± 7	0.003 ± 0.001	0.010 ± 0.002
5-3D	56 ± 0	0.005 ± 0.000	0.029 ± 0.000	82 ± 0	0.004 ± 0.000	0.028 ± 0.000
All 2D	122 ± 126	0.001 ± 0.001	0.004 ± 0.003	182 ± 207	0.002 ± 0.001	0.004 ± 0.003
All 3D	60 ± 21	0.005 ± 0.003	0.032 ± 0.023	102 ± 28	0.005 ± 0.003	0.033 ± 0.022

3.1.1.1 Problem 1 – Pin Cells

Figure 3 shows the k_{eff} comparisons for problem 1. There is a clear negative bias in the OTF results of about −100 pcm for most of the cases. The two capabilities compare much closer for 1e, which is an integral fuel burnable absorber (IFBA) pin cell. This bias is likely due to the lack of angle-dependent cross sections in the energy condensation along with some impact from the cylindricization of the cell. The results here were generated with SPH factors, which reduced the errors slightly but did not eliminate them. It is unclear whether there is an effective way to reduce these errors any further.

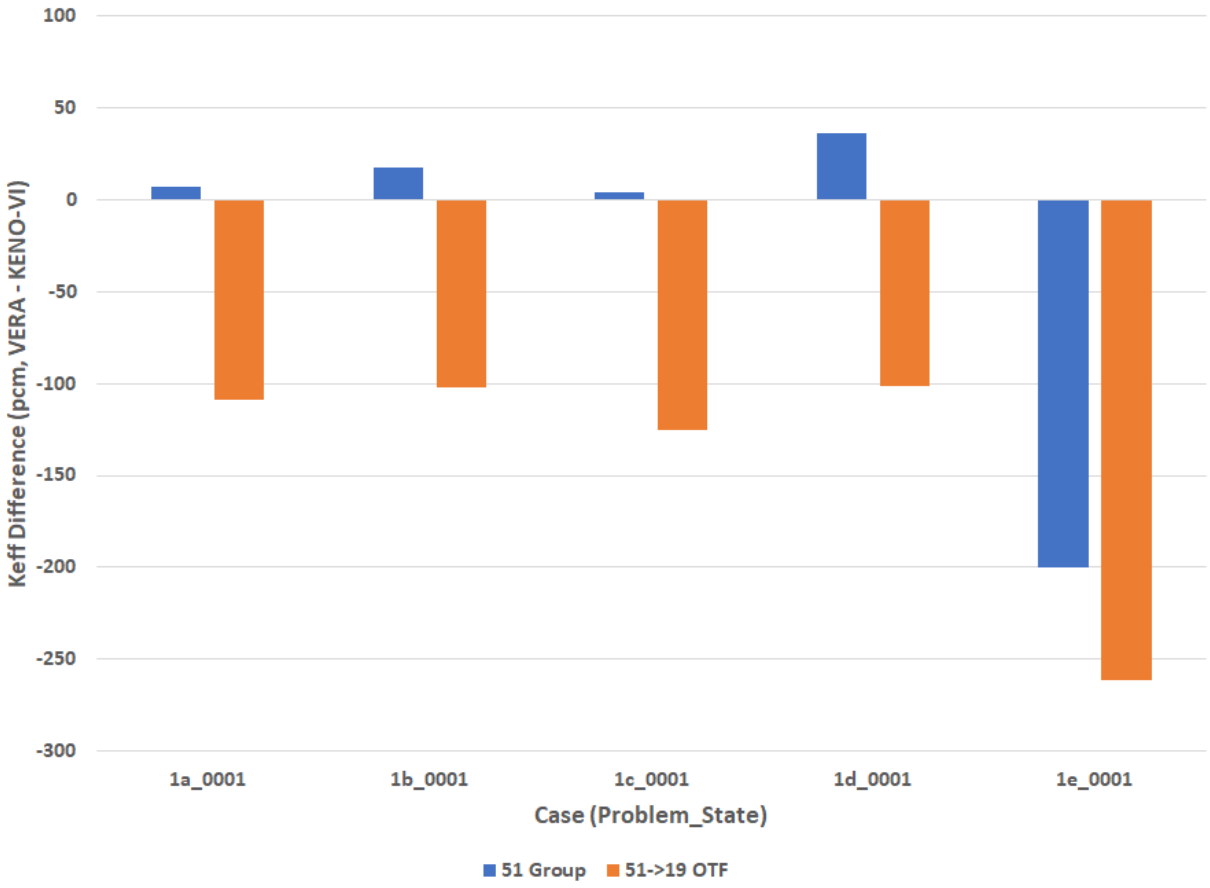

Figure 3. VERA problem 1 results for 51→19 OTF compared to KENO-VI.

3.1.1.2 Problem 2 – Lattices

Figure 4 shows the k_{eff} comparisons for problem 2. There is a slight negative bias for cases 2a–2d, 2i, and 2q, which are all lattices without strong absorbers; this bias is smaller than that for the pin cells. Small cases being highly sensitive compared to larger cases is typical and holds true in this instance. Most of the other lattices contain either burnable poisons or control rods; these lattices frequently show positive errors instead of negative. The errors tend to be around 100 pcm or less except for 2h (which has the strongly absorbing B_4C control rods) and 2p (which has 24 gad rods).

Figure 5 shows the pin power distribution comparisons. Every OTF case sees larger errors than the corresponding fine group calculation; however, the differences for many cases are around 0.001 or less, which is acceptably small. Both OTF and fine-group calculations show larger errors in the presence of strong absorbers: 2g, 2h, 2l, 2m, 2n, 2o, and 2p. The exceptions to this are 2e, 2f, 2j, and 2k—all of which contain pyrex absorbers but show maximum differences of 0.003 or less.

The largest increases in error compared to KENO-VI occurred in 2m, 2n, and 2p. To examine the differences, OTF can be compared directly with the fine-group calculations. Figure 6 shows the power distributions (OTF - 51 groups) for these three cases along with 2a and 2h, two lattices that did not show large differences between OTF and the fine group even though the 2h case shows large errors compared to KENO-VI. Problem 2a shows very good agreement, the largest difference being less than 0.002. Problem 2h contains 24 B_4C control rods, which are very strong thermal absorbers. The differences between OTF

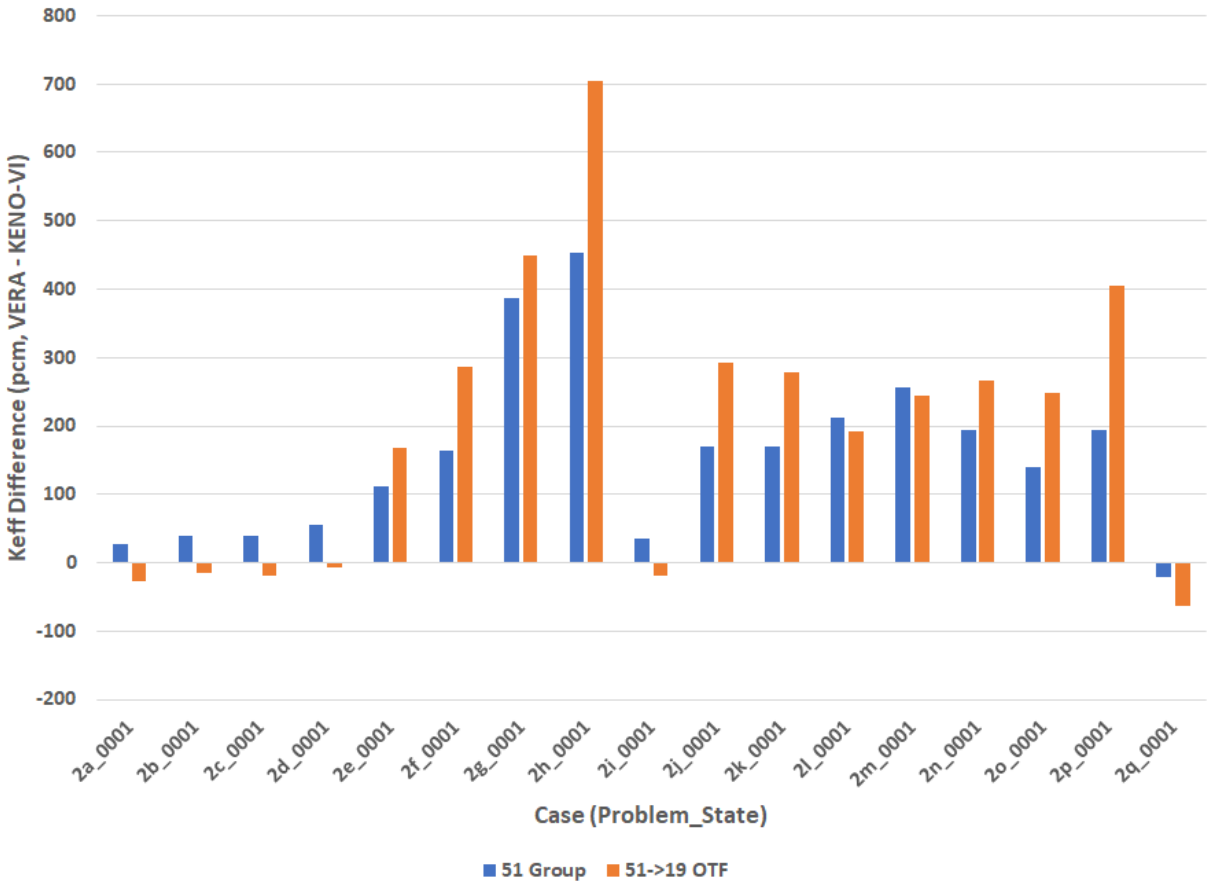

Figure 4. Problem 2 k_{eff} results for 51→19 OTF compared to KENO-VI.

and the fine group calculations for 2h are actually quite small, around 0.0015 or less everywhere; however, both the fine group and OTF calculations show the largest errors of any lattice compared to KENO-VI. Problem 2m contains 128 IFBA rods. Without exception, the IFBA rods are underpredicted and the non-IFBA rods are overpredicted. Problem 2n contains 104 IFBA and 20 wet annular burnable absorber (WABA) and is the only lattice that contains WABA. In 2n, the pattern is the same as in 2m, with underprediction in IFBA rods exclusively. The only difference is that the errors are not quite as large in the IFBA rods that neighbor WABAs; the largest errors are seen in the corner and center IFBA rods that are far from the WABAs. Finally, 2p contains 24 gad rods with 1.8% gad loading. Gadolinium has two isotopes that are very strong thermal absorbers. All 24 gad rods are overpredicted by nearly 0.005, with the remaining rods all being underpredicted.

The differences between OTF and fine group for individual lattices provides two major takeaways. First, for most rods and lattices, OTF does an excellent job of predicting a sufficiently accurate spectrum for individual cells to obtain an accurate global solution quickly. This is demonstrated in Figure 6 by 2a, a simple lattice with fuel and guide tubes, and 2h, a lattice that contains strongly absorbing control rods. Even though the errors compared to KENO-VI are larger for 2h, the differences between fine group and OTF are quite small. Second, there is clearly some weakness in the OTF capability when it comes to integral poisons such as IFBA and gad rods. The IFBA rods are underpredicted, whereas the gad rods are significantly overpredicted. It is likely that this result is related to the angle dependence of the spectrum collapse in the location of the poison—gad mixed into the fuel versus IFBA around the outside of the

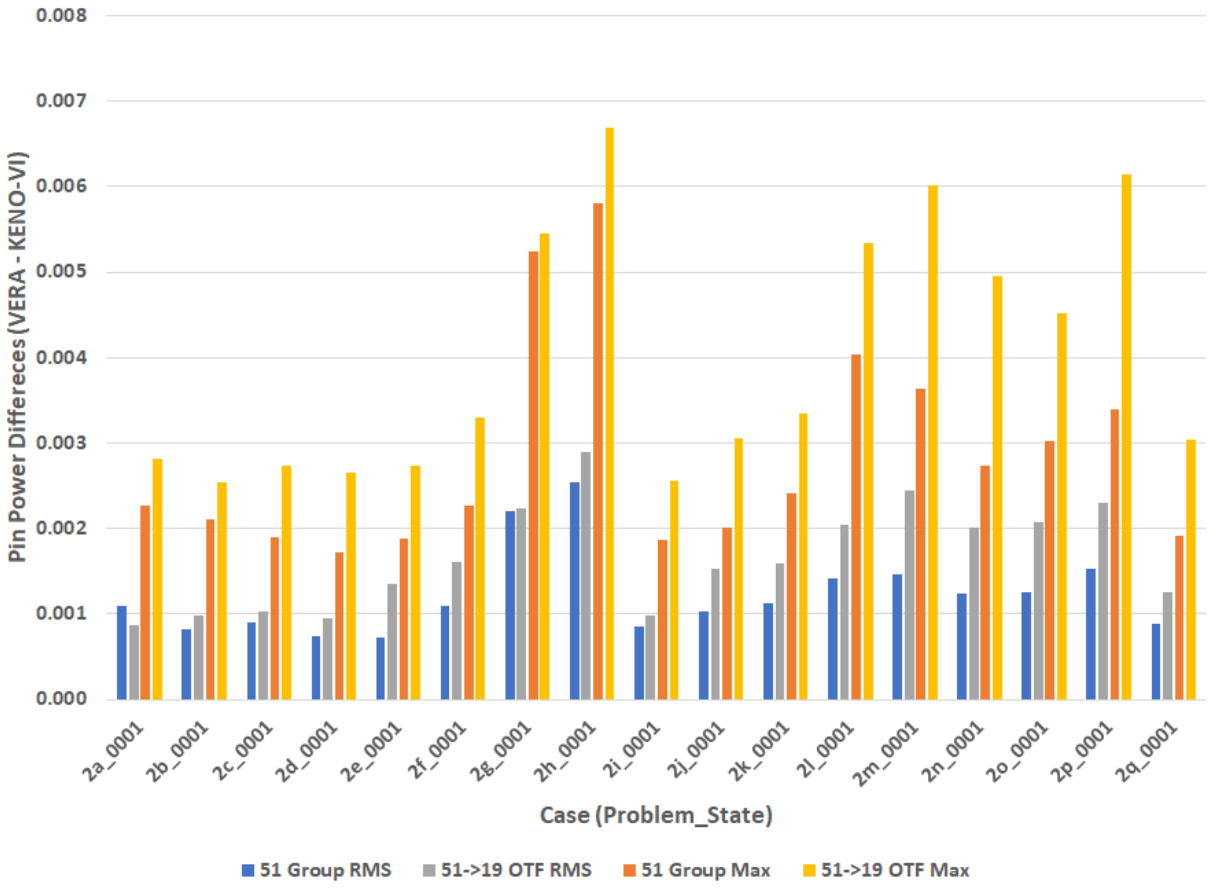

Figure 5. VERA problem 2 power distribution results for 51→19 OTF compared to KENO-VI

pellet—but further investigation is warranted to determine whether the spectra for these rods can be improved.

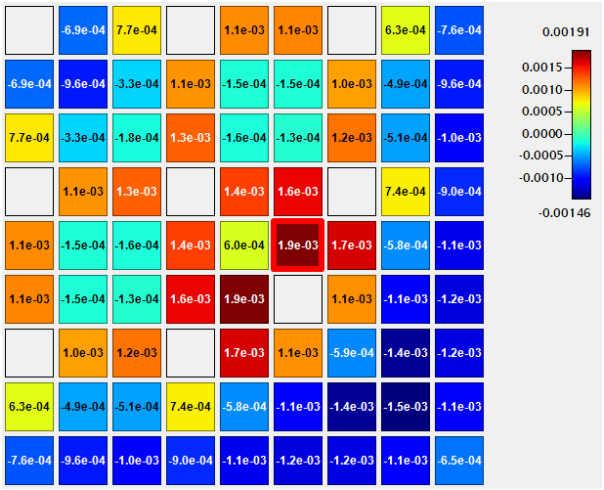

(a) Problem 2a – no poisons

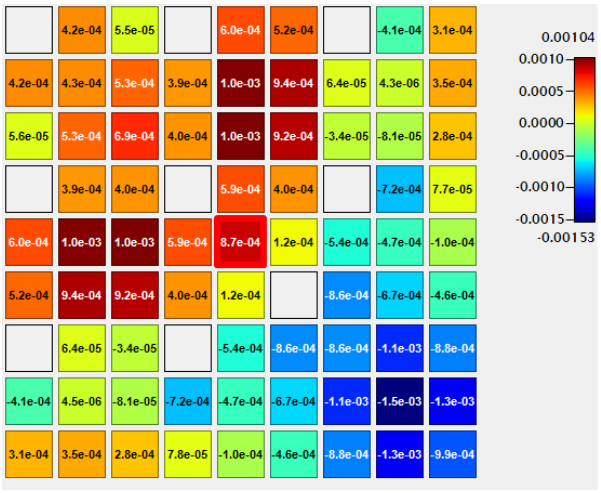

(b) Problem 2h – 24 B_4C control rods

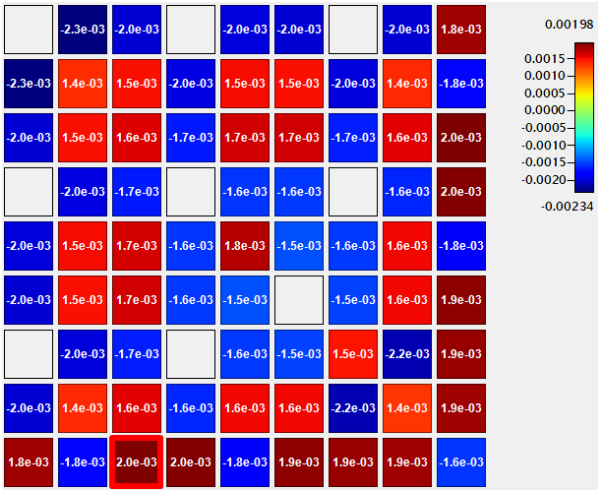

(c) Problem 2m – 128 IFBA rods

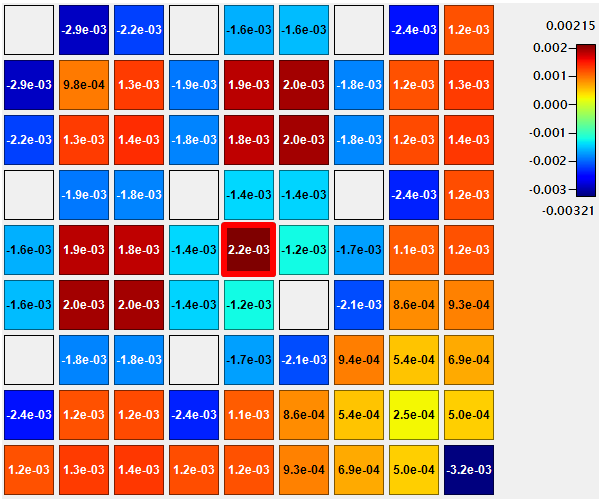

(d) Problem 2n – 104 IFBA rods, 20 WABAs

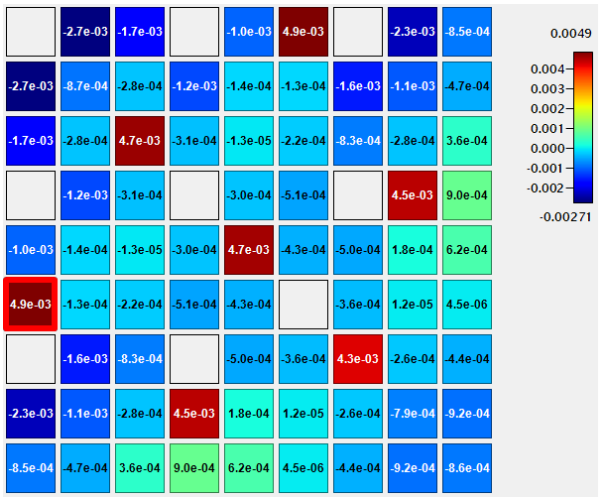

(e) Problem 2p – 24 gad rods

Figure 6. (OTF - fine group transport) for several of the VERA Progression Problem 2 lattices.

3.1.1.3 Problems 4 and 5 – 2D Multi-lattice

With the larger cases, the k_{eff} differences between the OTF and fine group calculations diminish significantly, as shown in Figure 7. The largest changes from the fine group calculation to OTF are less than 50 pcm, and the single largest error by any OTF calculation is less than 140 pcm. These results all represent excellent agreement compared to Monte Carlo.

The power distributions are compared in Figure 8. The errors are larger than those for the pin cell and lattice calculations, which is expected given the greater heterogeneity in these models. The OTF fine group spectra are calculated using isolated pin cells with white boundary conditions. This causes the pin cell to have a different energy spectrum for the boundary condition than what is seen by that cell in the larger problem. For this reason, it is expected to see larger errors the more heterogeneous a problem is. The increase in error from fine group transport to OTF is small for all cases; in the case of 5b-2D and 5c-2D, the maximum errors actually decrease when using OTF. The root mean square error (RMSE) also decreases with OTF for 5b-2D. This is discussed in greater detail later in this section. Overall, the single greatest increase in power distribution error when moving from fine group transport to OTF occurs in 4a-2D with an increase of about 0.002. This is a perfectly acceptable increase given the computational performance improvements.

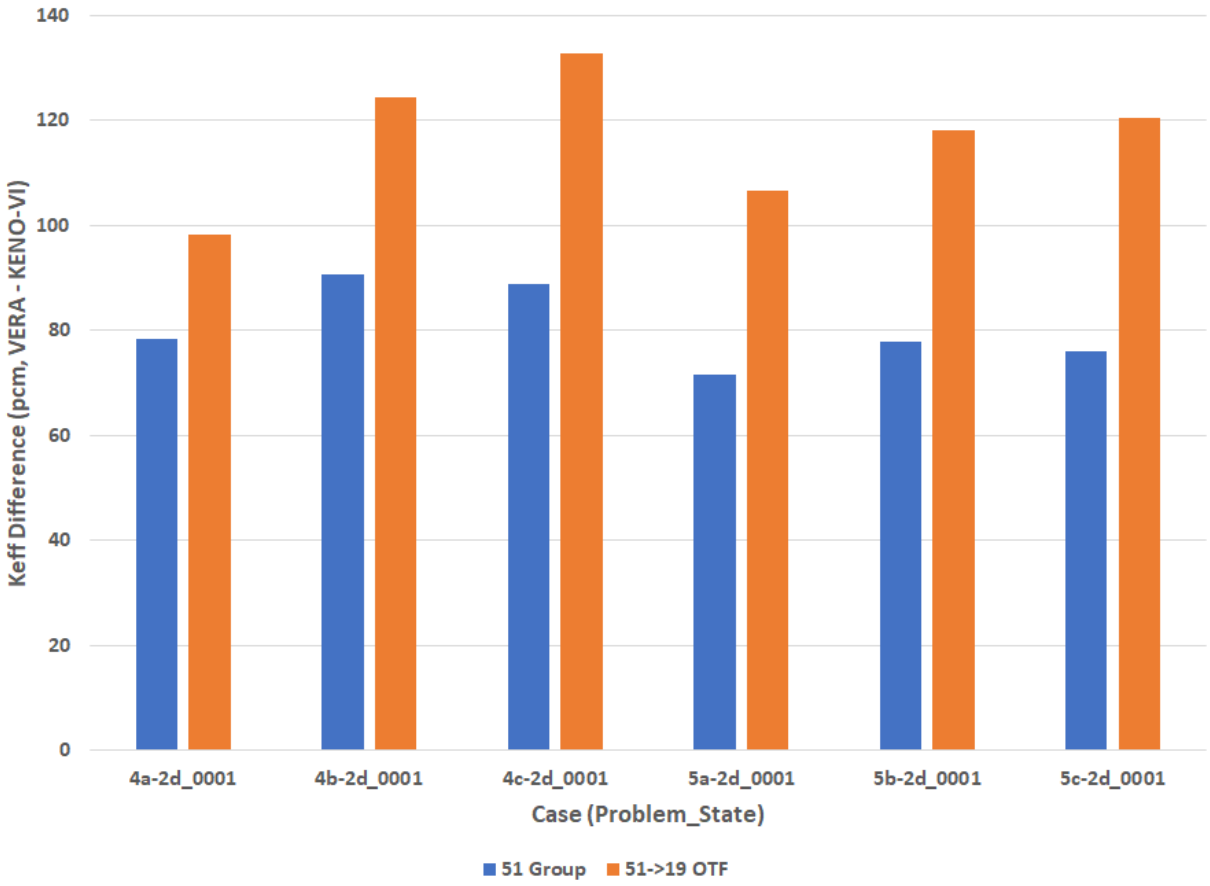

Figure 7. VERA problems 4 and 5 2D k_{eff} results for 51→19 OTF compared to KENO-VI.

Figure 9 shows the differences between OTF and the fine group calculations for each of the 2D variations of problem 4. Problem 4a-2D shows good agreement overall. The largest overpredictions are seen around

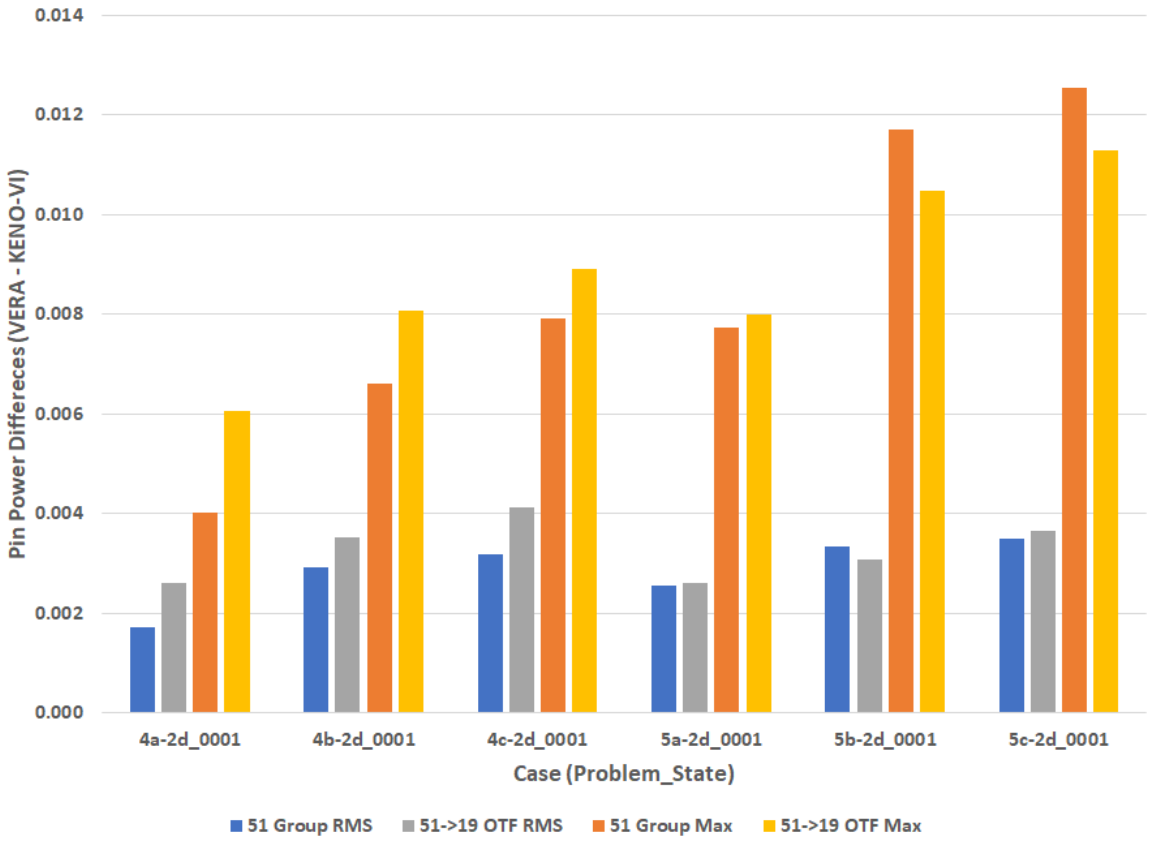

Figure 8. VERA problems 4 and 5 power distribution results for 51→19 OTF compared to KENO-VI.

the empty guide tubes in assemblies with pyrex; the largest underpredictions are on the southeast periphery of the problem. While the magnitude of these differences is acceptably small, it again indicates a skewing of the power shape in the presence of strong absorbers. Problem 4b-2D adds AIC control rods in the center assembly. Overall, the comparison is similar to 4a-2D: the magnitudes of the differences are slightly larger, but the overall shape of the differences is similar. Problem 4c-2D continues to show the largest differences in the same locations, but there is a significant change in the center assembly due to the presence of the B_4C control rods. The entire assembly was underpredicted for 4a-2D and 4b-2D but was overpredicted for 4c-2D. This is similar to the trend observed in problem 2h—but for multiple assemblies instead of a single lattice.

Figure 10 shows the differences between OTF and the fine group calculations for each of the 2D variations of problem 5. The variations are the same as those for problem 4: 5a-2D has no control rods while 5b-2D and 5c-2D have bank D inserted in locations H-8, D-8, H-12, and D-12. Problem 5b-2D uses AIC rods; problem 5c-2D uses B_4C rods. It can be seen that in the rodded locations, 5b-2D has the smallest differences and 5a-2D has the largest differences. This trend in the control rod, pyrex, and IFBA behaviors is consistent with the smaller problems. The most obvious feature of all 3 variations is the reflector boundary. The overprediction of power in the core interior results in an underprediction on the core periphery, except in the outermost row of pins. This row neighbors the baffle, which is treated by a very simple $\frac{1}{v}$ spectrum instead of a transport-based spectrum. This obviously results in a poor approximation of the reflector cross sections and causes significant errors in the pins neighboring the baffle. This is a clear area for improvement in the current OTF methodology.

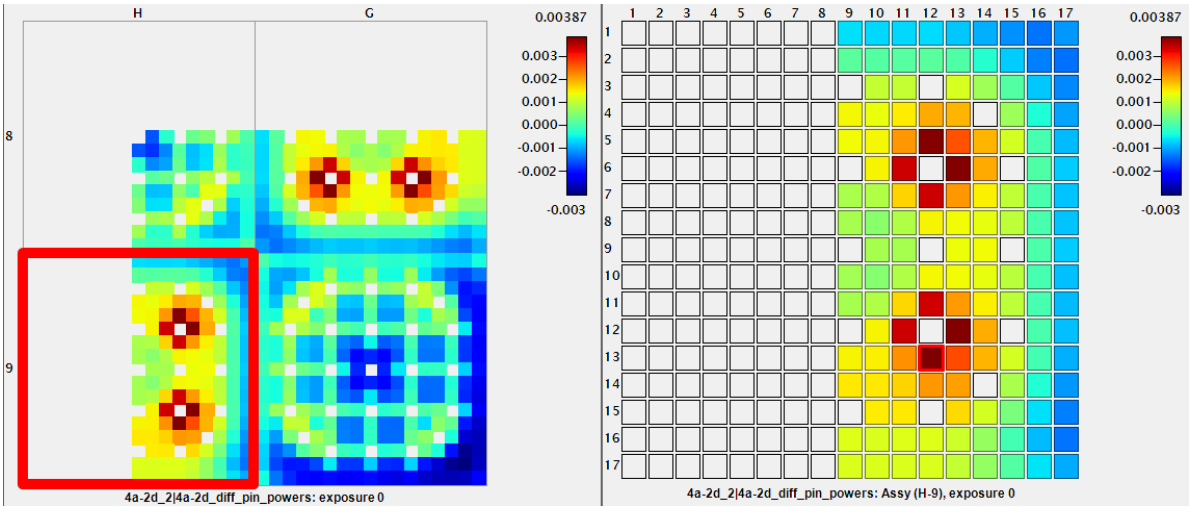

(a) Problem 4a-2d – no control rods

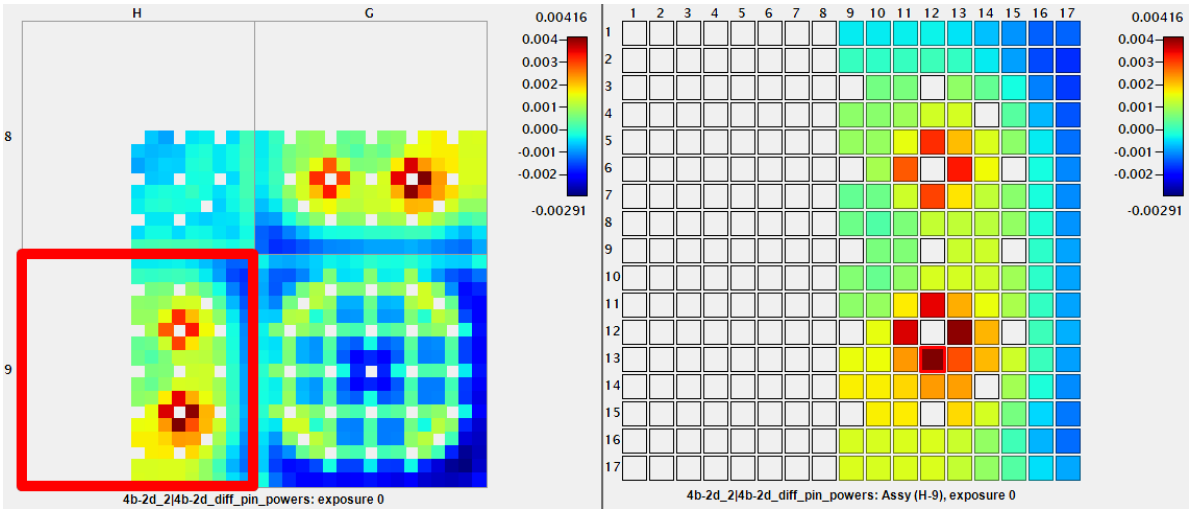

(b) Problem 4b-2d – AIC control rods

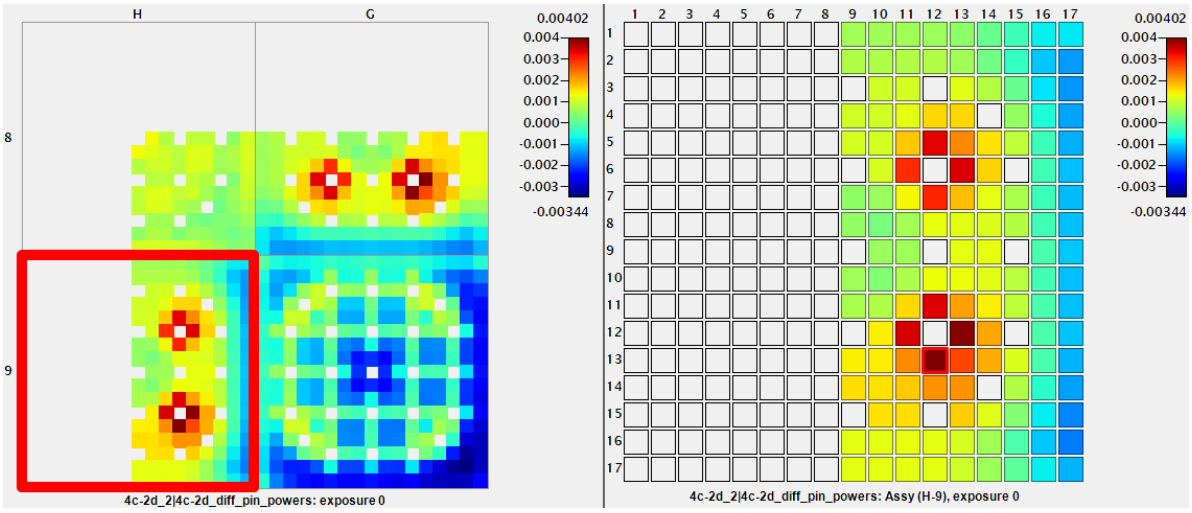

(c) Problem 4c-2d – B_4C control rods

Figure 9. (OTF - fine group transport) for 2D variations of VERA Progression Problem 4

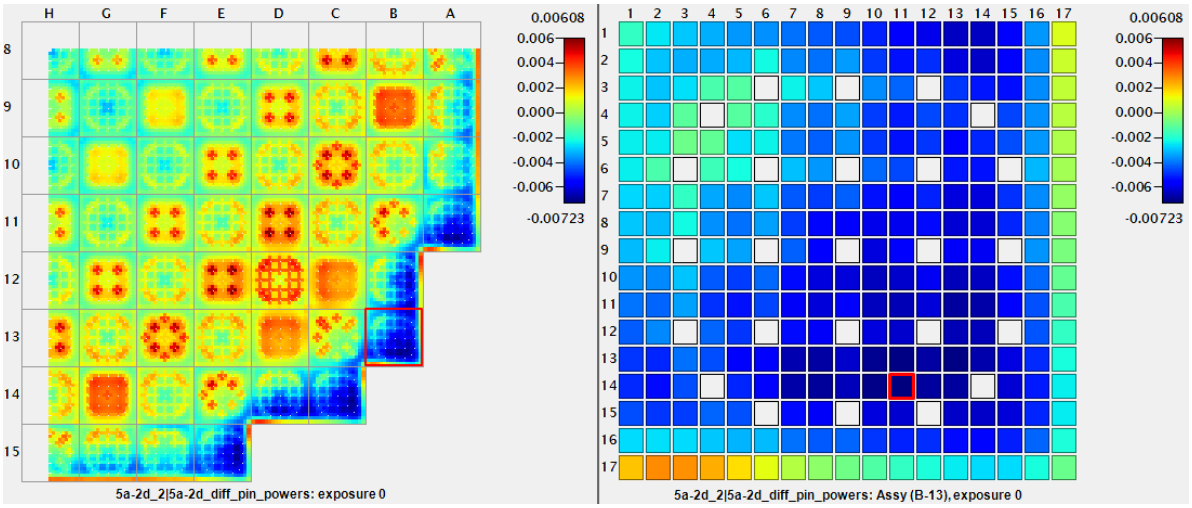

(a) Problem 5a-2d – no control rods

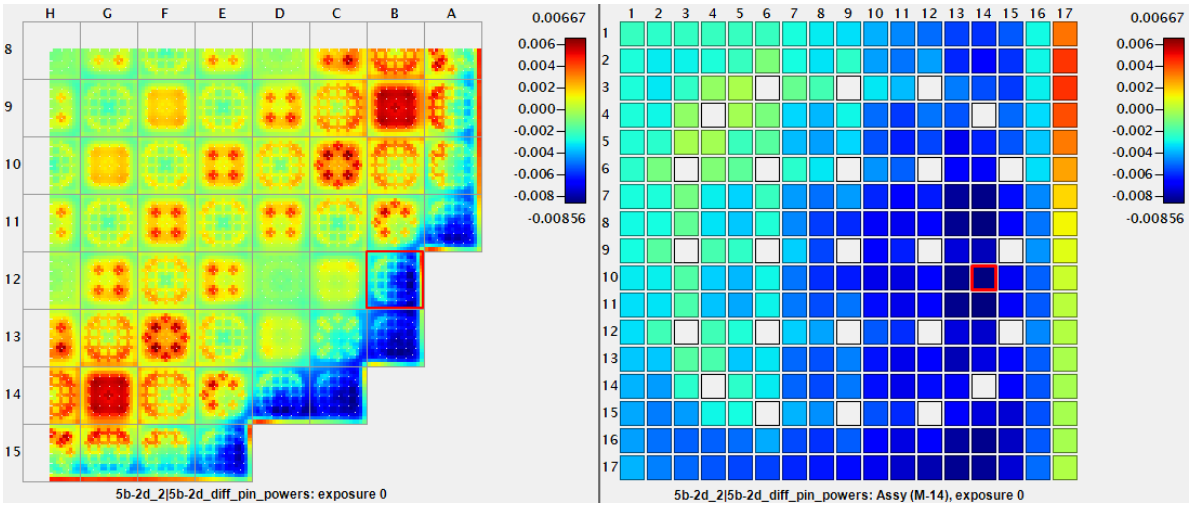

(b) Problem 5b-2d – AIC control rods

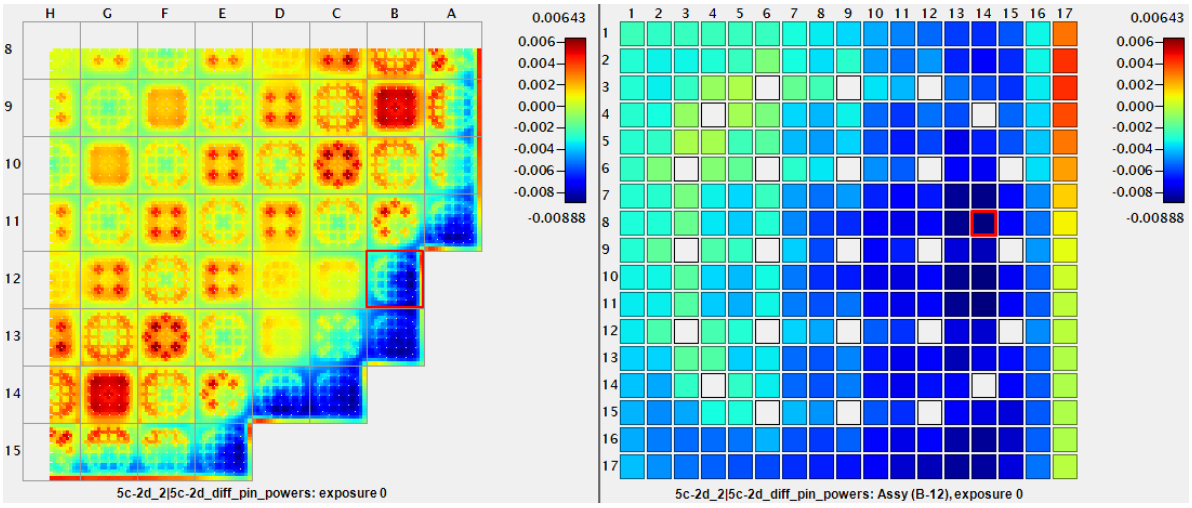

(c) Problem 5c-2d – B_4C control rods

Figure 10. (OTF - fine group transport) for 2D variations of VERA Progression Problem 5.

3.1.2 3D PWR RESULTS

The second set of results covers the extension of the OTF capability from 2D to 3D, which is less common amongst other codes. This will be tested on problem 3, 3D versions of problem 4, and 3D versions of problem 5. These problems are briefly summarized in Table 3. As in the previous section, KENO-VI results are also available for these problems.

Table 3. Summary of the 3D VERA Progression Problems

Problem	Variations	Description
3	2	hot zero power (HZP) isothermal assembly with and without discrete burnable poisons
4	11 rod positions	3×3 array of assemblies with discrete burnable poisons and control rods
5	32	Fresh HZP Watts Bar Unit 1 (WB1) core; 10 criticals, 9 bank worths, 2 isothermal temperature coefficient (ITC), and 11 integral bank D worth calculations

3.1.2.1 Problem 3

Table 4 shows the differences between fine group transport and OTF compared to KENO-VI for problem 3. Problem 3a is a 3D assembly with no burnable poisons and shows a decrease in k_{eff} when using OTF; problem 3b contains pyrex poisons and shows an increase in k_{eff} when using OTF. This trend is consistent with that observed in problem 2a–2f. The power distributions compare well. The RMSE is almost the same between the fine group transport and the OTF calculations. The maximum differences increase some but remain under 1% and are, therefore, not of significant concern.

Table 4. Summary of results for VERA Progression Problem 3

Problem	51 Groups			51→19 OTF		
	k_{eff} (pcm)	Power Distribution		k_{eff} (pcm)	Power Distribution	
		RMSE	Max.		RMSE	Max.
3a	28	0.002	0.005	1	0.003	0.009
3b	107	0.002	0.006	205	0.003	0.008

Figure 11 shows the power distribution differences between OTF and fine group for problem 3. Both problems show similar agreement, with the highest overpredictions on the interior and underprediction closer to the problem boundary radially and axially. In 3b, the power is pulled toward the interior more than in 3a, again showing a tendency of overprediction around the pyrex. The power is also moved toward the center of the assembly axially, thought 3a also shows a change in axial power shape. This is likely affected by the treatment of the axial reflectors, which receive the same $\frac{1}{v}$ treatment that the baffle receives as part of the radial reflector.

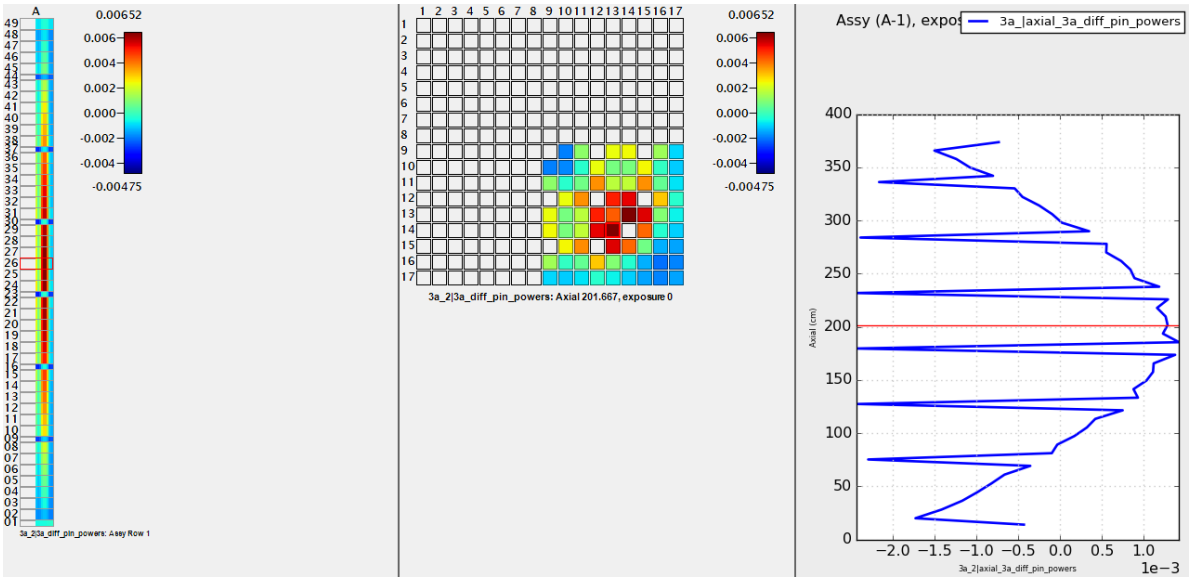

(a) Problem 3a – no poisons

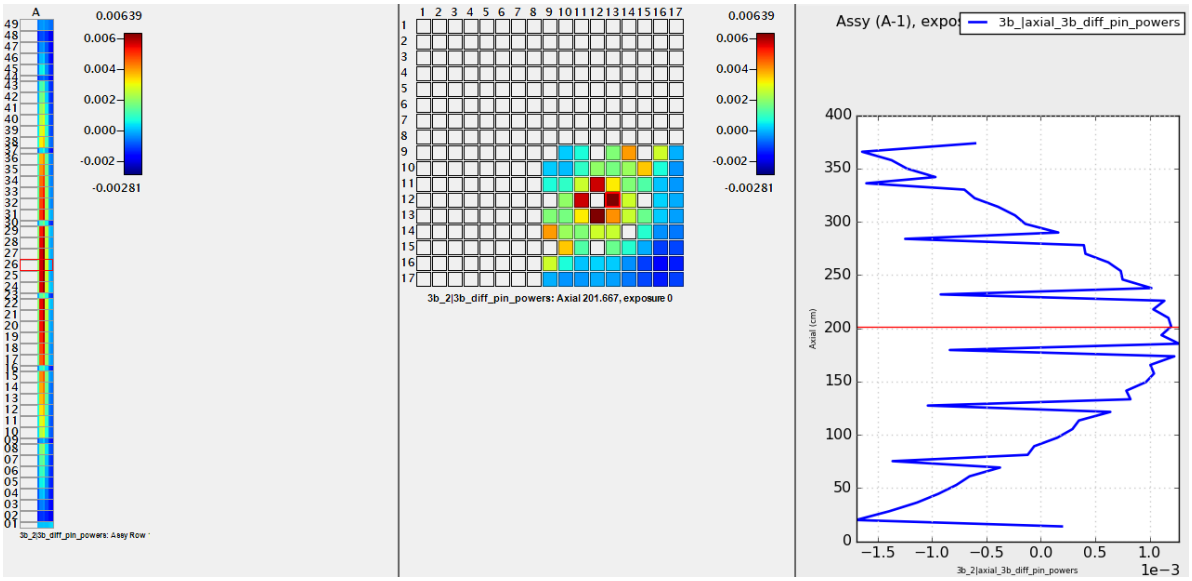

(b) Problem 3b – pyrex poisons

Figure 11. (OTF – fine group transport) for VERA Progression Problem 3.

3.1.2.2 Problem 4

Figure 12 shows the problem 4 k_{eff} comparisons, and Figure 13 shows the power distribution differences. The k_{eff} changes for problem 4 are consistent, ranging from about 60 pcm high in the first state, all rods in (ARI), down to about 30 pcm high in the last state, all rods out (ARO). Again, this is consistent with the difference between rodded and unrodded lattices—problems 2a–2d versus 2g–2h—as well as with the differences between a 3D assembly without and with pyrex absorbers—problem 3a versus problem 3b.

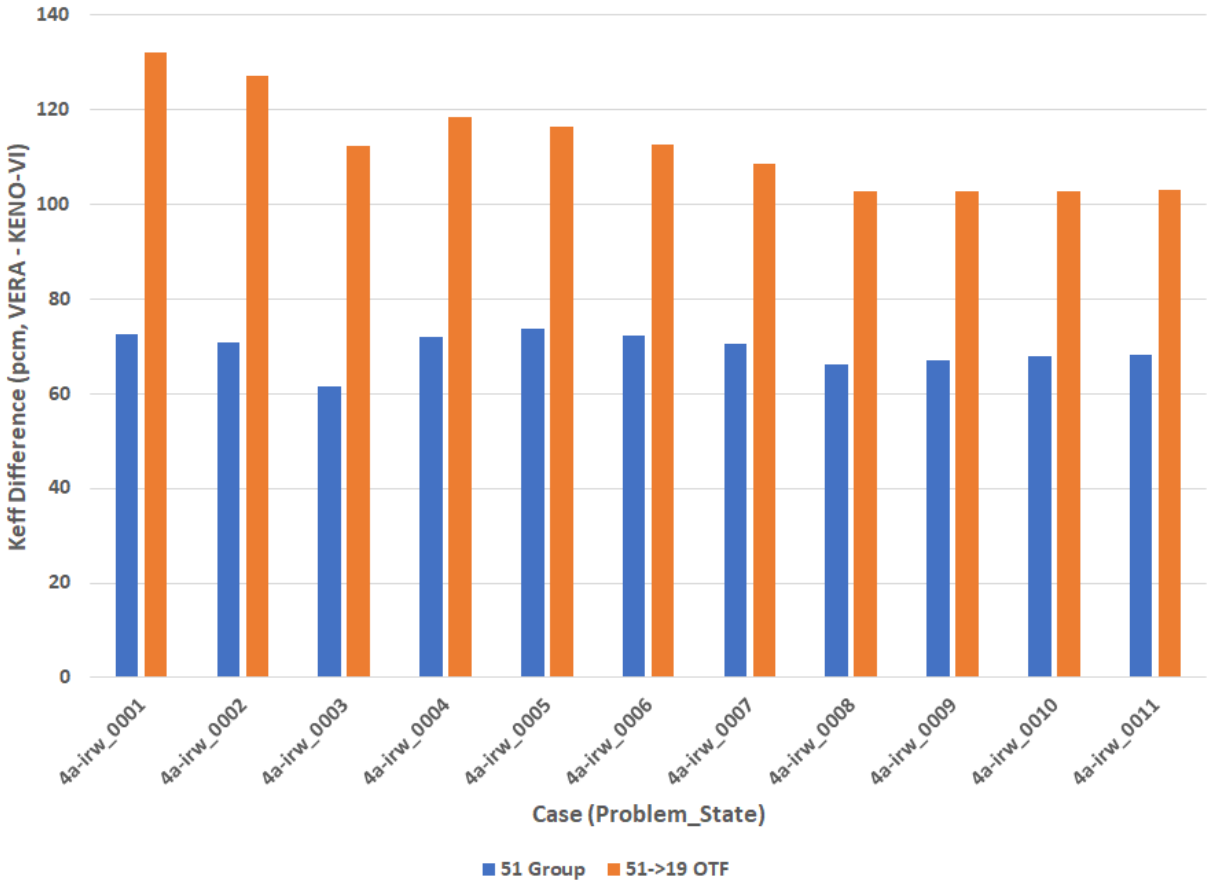

Figure 12. VERA problem 4 3D k_{eff} results for 51→19 OTF compared to KENO-VI.

Figure 14 shows the problem 4 integral rod worth (IRW) power distribution differences between OTF and fine group transport at ARI, the rod position of maximum error, and ARO. The agreement for ARI and ARO is less than 0.01 everywhere, which is good for a highly heterogeneous 3D calculation. For ARI, the largest difference is 0.009; for ARO, the largest difference is 0.007. In both cases, the overall shape of the differences is similar to that exhibited by problems 3a and 3b. The largest difference over all 11 state points occurs when the rod is withdrawn 69 steps in the fourth state point. This state point has a severe axial power shape that tends to stress transport calculations to begin with. Add to that the isolation effects of the OTF capability, and it is not surprising that the largest errors are seen just below the tip of the partially inserted rod.

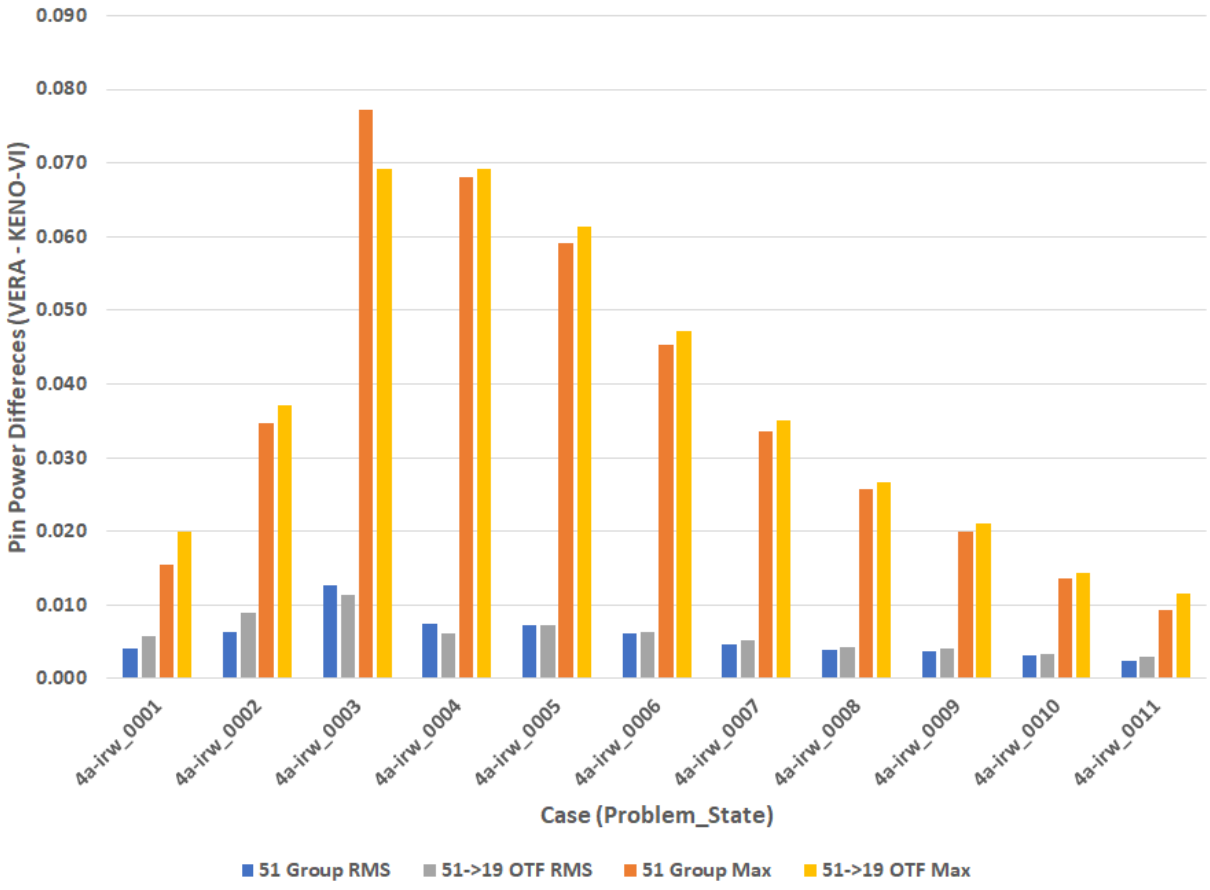

Figure 13. VERA problem 4 3D power distribution results for 51→19 OTF compared to KENO-VI.

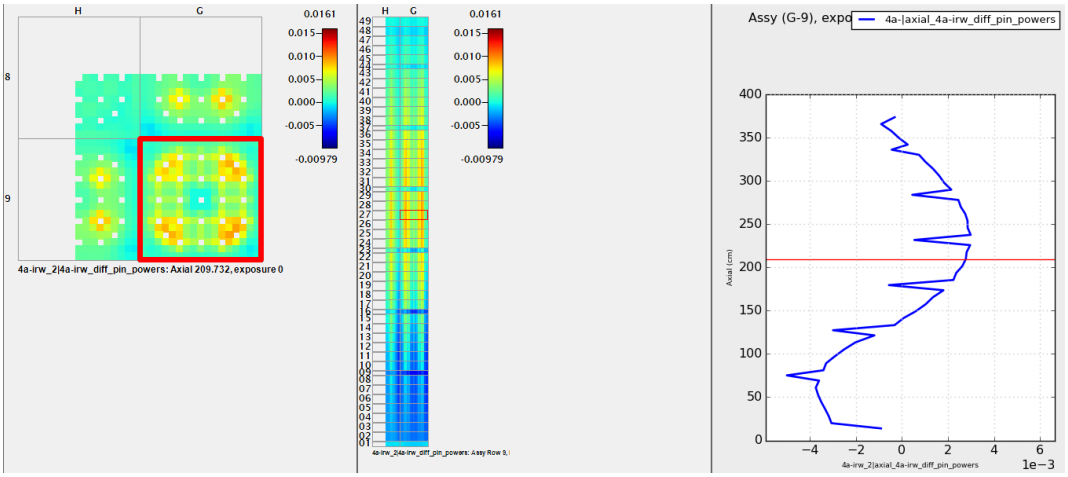

(a) Problem 4 – all rods in

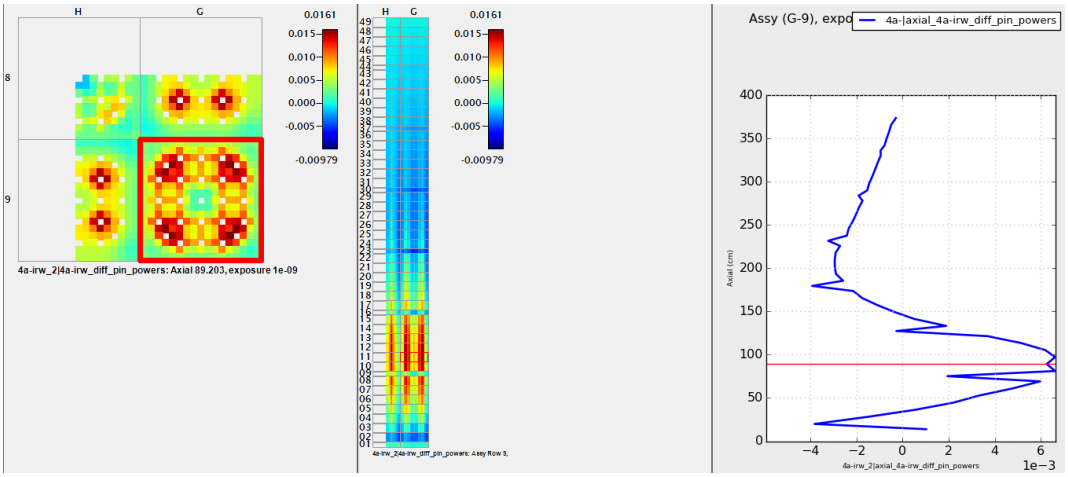

(b) Problem 4 – maximum error

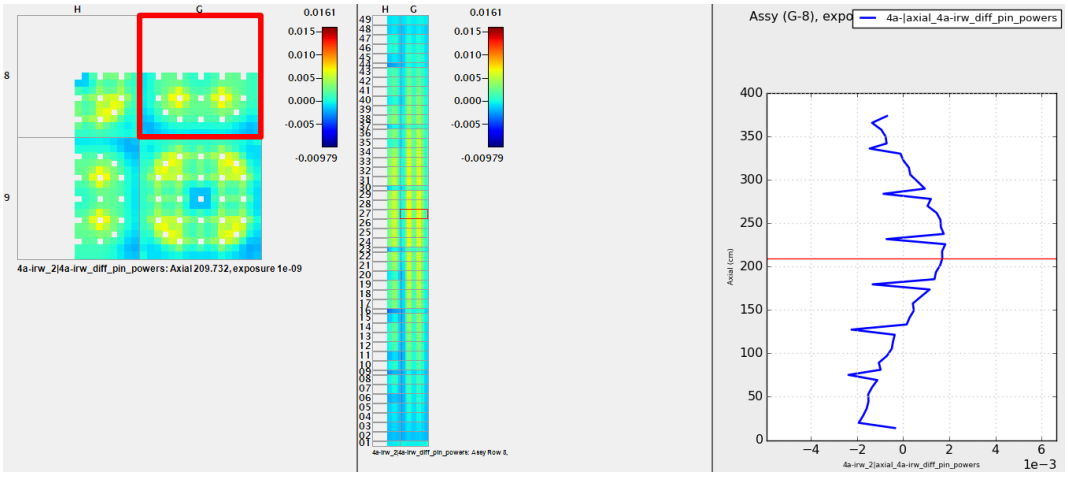

(c) Problem 4 – all rods out

Figure 14. (OTF – fine group transport) for VERA Progression Problem 4 IRW calculation.

3.1.2.3 Problem 5

Problem 5 consists of 32 state points used to calculate bank worths, ITC, integral bank D worth, and critical configurations. The first critical is the only one with available KENO-VI power distribution results because of the difficulty of achieving acceptably low uncertainties in a problem of this size; the remaining 31 state points have k_{eff} results from KENO-VI. The power distribution compared to KENO-VI is shown in Table 2. The k_{eff} error was 82 pcm, the RMSE power distribution error was 0.004, and the maximum power distribution error was 0.028.

Table 5 shows the k_{eff} differences between VERA and KENO-VI. The fine group calculations are consistently 50–60 pcm off, whereas the OTF calculations are around 100 pcm off. Overall, this is good agreement for both sets of calculations. Table 6 shows the OTF calculations compared to fine group transport for all 32 state points of VERA problem 5. The power distributions show agreement around 0.003 for the RMSE power distributions differences and 0.02 or less for the maximum power distribution differences. For a 3D calculation, this is good agreement between the two, especially considering the simplistic treatment of the reflector. Furthermore, because the maximum error compared to KENO-VI for the first critical was 0.028 for OTF and 0.029 for fine group transport, the 0.02 differences between fine group and OTF are not of great concern.

Table 5. Summary of VERA problem 5 OTF results compared KENO-VI

Problem	51 Groups k_{eff} (pcm)	51 →19 OTF k_{eff} (pcm)
Criticals	53 ± 35	97 ± 34
Rod Worths	57 ± 2	110 ± 4
ITC	59 ± 1	103 ± 0
Bank D IRW	52 ± 1	102 ± 5
Total	54 ± 20	103 ± 19

Table 6. Summary of VERA problem 5 OTF results compared to fine group transport

Problem	k_{eff} (pcm)	Power Distribution	
		RMSE	Max.
Criticals	50 ± 20	0.003 ± 0.000	0.021 ± 0.004
Rod Worths	54 ± 4	0.003 ± 0.000	0.017 ± 0.003
ITC	48 ± 4	0.003 ± 0.000	0.014 ± 0.001
Bank D IRW	50 ± 4	0.003 ± 0.000	0.017 ± 0.002
Total	51 ± 12	0.003 ± 0.000	0.018 ± 0.002

The power distribution slices at the maximum difference location for problem 5 are shown in Figure 15. Not surprisingly, these differences look like a combination between the 5-2D results and the 4-3D results. The largest errors occur below the tip of the partially inserted rod at D-12. There is also a stripe of larger errors along the baffle all the way around the core. There are no features in this comparison that are unique compared to the previously discussed cases.

One feature that is noticeable in all the 3D calculation differences is the spacer grids. The OTF power prediction in spacer grid locations is always lower than in the planes immediately above and below when compared with the fine group transport. VERA uses a semi-explicit spacer grid treatment in which there is a thin strip of spacer grid around the outside of each pin cell in the locations with spacer grids. The

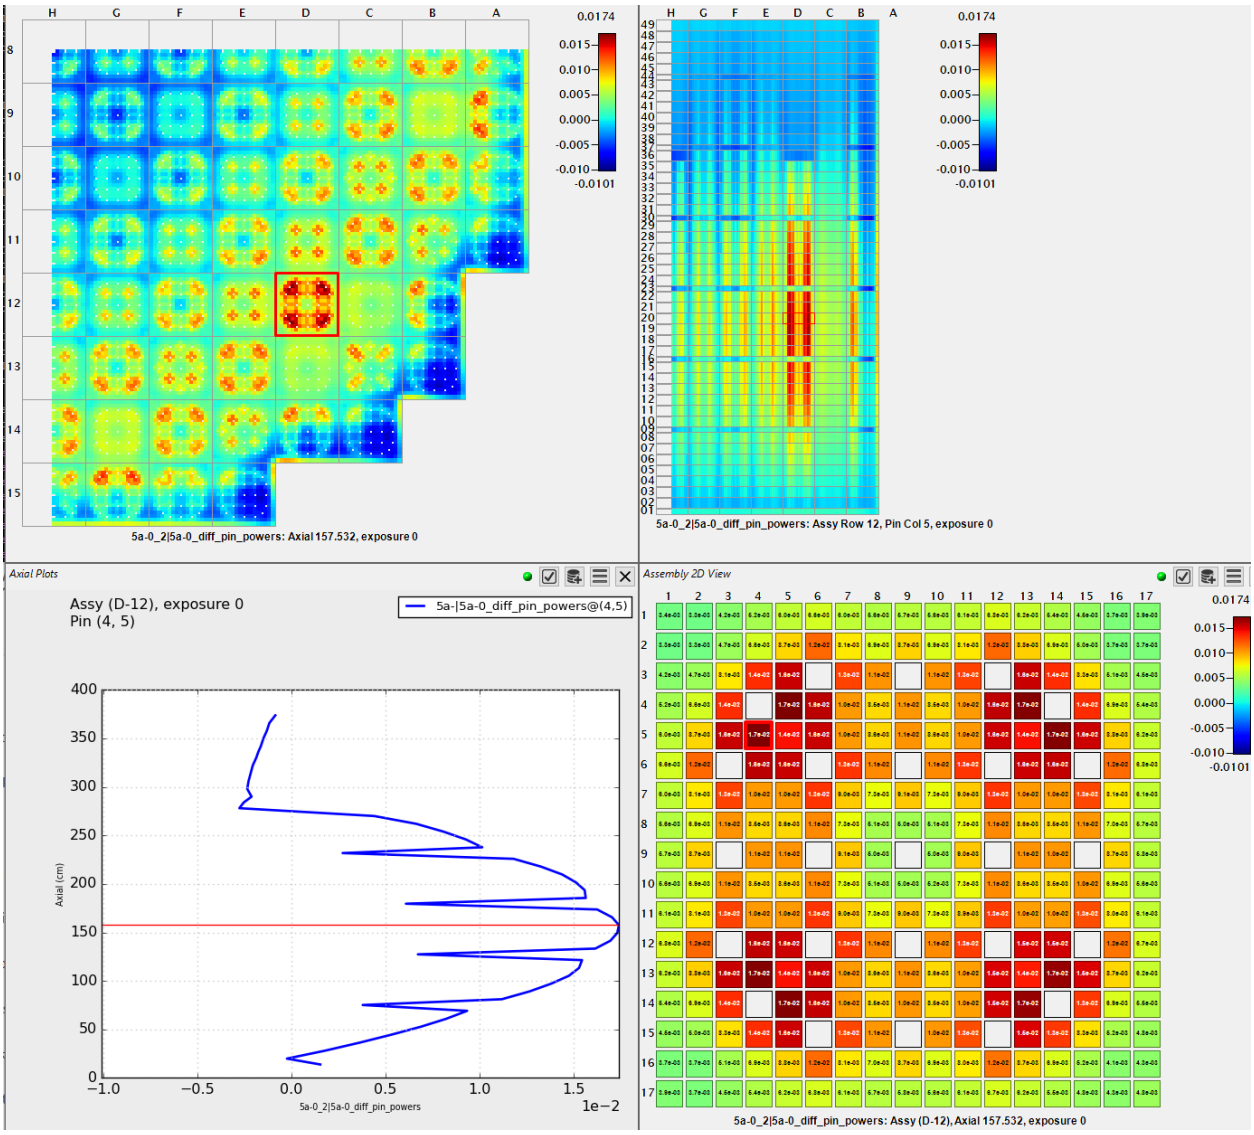

Figure 15. (OTF – fine group transport) for VERA Progression Problem 5.

explanation for the larger-than-expected power decreases at the spacer grids is not clear yet, but something about the CPM calculations with the spacer grid in the outermost ring causes a shift in the spectrum; this effect should be investigated further.

3.1.3 2D BWR RESULTS

The 2D BWR cases consist of a matrix of GE-9, and GE-14 bundle types; 0%, 40%, and 80% void; nominal steady-state, steady-state with control blade, and 127-state depletion without control blade. There are also Peach-6 and part-length rods (PLR) bundles, controlled and uncontrolled, at 0%, 40%, and 80% void. Finally, there are 2D mini-cores consisting of a controlled or uncontrolled 4×4 array of GE-14 bundles at 0%, 20%, and 60% void. This gives a total of 387 cases each for GE-9 and GE-14 and 6 cases each for Peach-6, PLR, and 4×4 GE-14. Table 7 gives the summary results for these calculations.

All calculations were run using P_2 scattering, which is typical for BWR calculations since they have greater anisotropy than PWRs. The fine group calculation shows maximum power errors around 1.5% and k_{eff}

errors below 200 pcm; the ge-14 80% void depletion case has larger k_{eff} errors around 300 pcm. The OTF calculations show worse agreement. The k_{eff} errors are up to 100 pcm higher for GE-9 and GE-14 bundles and closer to 300 pcm higher for the Peach-6 and PLR lattices. The RMSE power distribution differences are higher, ranging from 0.004 to 0.014, and the maximum power differences are worse by about 0.01 for many cases. The PLR lattices show the worst differences, with maximum errors around 0.044.

These errors are not particularly surprising. The approximation of the channel box and bypass cells as cylindrical, discussed in Section 2.8, is probably a poor one. These would probably be better approximated with slab geometry, which may improve the prediction of the edge pins, which tend to have the largest errors regardless.

Table 7. Summary of 2D BWR results for 51→19 OTF

Problem	51 Groups			51→19 OTF		
	k_{eff} (pcm)	Power Distribution		k_{eff} (pcm)	Power Distribution	
		RMSE	Max.		RMSE	Max.
GE-9	87 ± 112	0.003 ± 0.001	0.008 ± 0.002	101 ± 167	0.004 ± 0.002	0.011 ± 0.006
GE-14	82 ± 18	0.003 ± 0.001	0.009 ± 0.002	172 ± 30	0.007 ± 0.003	0.018 ± 0.009
Peach-6	60 ± 28	0.004 ± 0.001	0.014 ± 0.005	352 ± 81	0.009 ± 0.000	0.024 ± 0.003
PLR	60 ± 59	0.004 ± 0.002	0.013 ± 0.005	352 ± 151	0.014 ± 0.003	0.044 ± 0.014
4×4 GE-14	17 ± 109	0.004 ± 0.001	0.016 ± 0.002	149 ± 193	0.012 ± 0.000	0.036 ± 0.003
Total	84 ± 92	0.003 ± 0.001	0.009 ± 0.003	140 ± 126	0.006 ± 0.003	0.015 ± 0.009

3.1.4 3D BWR RESULTS

For 3D BWR validation, there are three sets of problems. First is a 3D GE-14 bundle at cold zero power (CZP), HZP, and hot full power (HFP). Second is a 3D bundle blade worth calculation, with the blade being pulled from 0 to 48 steps in increments of four. Finally, there are four 3D mini-cores with various rodded configurations, as shown in Table 8.

Overall, the 3D results are similar to the 2D results but worse. The higher order scattering results with 51 groups are high to begin with, with errors between 6% and 8% for many of the cases. Therefore, maximum errors over 0.10 for the OTF are unsurprising at this time given the current treatment of the bypass flow regions. Improvements to geometry should bring those error down significantly.

Table 8. Summary of 3D BWR results for 51→19 OTF

Problem	51 Groups			51→19 OTF		
	k_{eff} (pcm)	Power Distribution		k_{eff} (pcm)	Power Distribution	
		RMSE	Max.		RMSE	Max.
3D GE-14	27 ± 40	0.017 ± 0.004	0.067 ± 0.023	119 ± 122	0.034 ± 0.008	0.126 ± 0.018
3D Blade Worth	23 ± 24	0.013 ± 0.003	0.069 ± 0.031	169 ± 68	0.030 ± 0.004	0.146 ± 0.072
3D Mini-core	24 ± 33	0.012 ± 0.002	0.080 ± 0.042	217 ± 48	0.040 ± 0.009	0.211 ± 0.103
Total	23 ± 27	0.013 ± 0.003	0.071 ± 0.031	179 ± 85	0.032 ± 0.006	0.159 ± 0.076

3.2 252→51 OTF

The next set of results presented is based on the 252 energy group transport library. The calculations were collapsed to the default 51 group structure. The energy group structures are illustrated in Figure 16. The results are summarized in Table 9. Each set of problems is shown in more detail in the following sections.

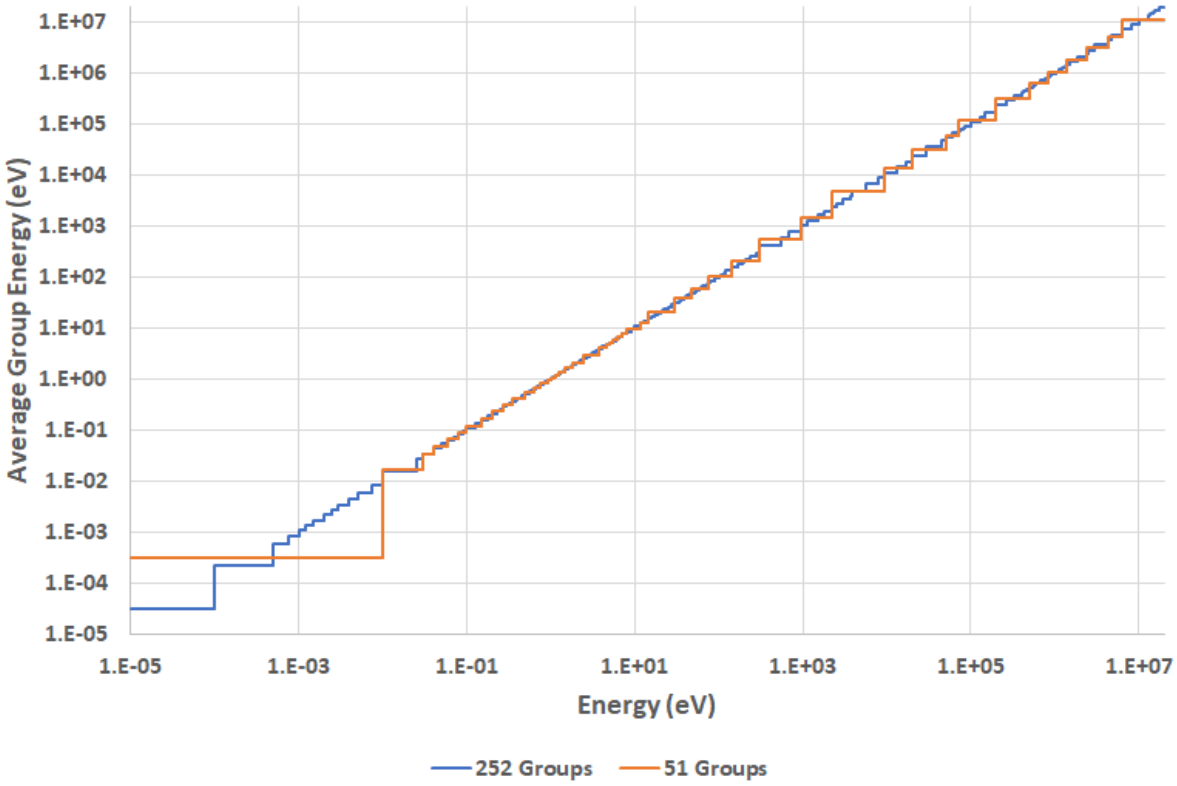

Figure 16. Average energy and energy boundaries of each group for the 252- and 51-group libraries.

Table 9. Summary of the VERA progression problem 252→51 OTF results compared to KENO-VI; all values are the average of |VERA − KENO-VI| over the set of problems; DNC means the calculation did not converge

Problem	252 Groups			252→51 OTF		
	k_{eff} (pcm)	Power Distribution		k_{eff} (pcm)	Power Distribution	
		RMSE	Max.		RMSE	Max.
1	172 ± 115	—	—	482 ± 82	—	—
2	90 ± 115	0.001 ± 0.000	0.003 ± 0.000	203 ± 153	0.001 ± 0.001	0.003 ± 0.001
3	DNC	DNC	DNC	263 ± 82	0.002 ± 0.000	0.005 ± 0.001
4-2D	61 ± 4	0.002 ± 0.001	0.006 ± 0.001	182 ± 17	0.002 ± 0.000	0.005 ± 0.001
4-3D	DNC	DNC	DNC	211 ± 8	0.005 ± 0.002	0.036 ± 0.023
5-2D	72 ± 1	0.003 ± 0.001	0.011 ± 0.001	177 ± 13	0.009 ± 0.001	0.020 ± 0.002
5-3D	DNC	DNC	DNC	210 ± 0	0.010 ± 0.000	0.090 ± 0.000
All 2D	99 ± 114	0.001 ± 0.001	0.004 ± 0.003	247 ± 175	0.002 ± 0.003	0.005 ± 0.006
All 3D	DNC	DNC	DNC	212 ± 69	0.005 ± 0.003	0.036 ± 0.028

3.2.1 2D RESULTS

3.2.1.1 Problem 1 – Pin Cells

Figure 17 shows the k_{eff} comparisons for 252-group transport and 252→51 OTF calculations. The most noticeable difference between these results and those for the 51→19 OTF calculations is that the negative reactivity bias is much larger in this case, about 300 pcm more negative than the fine group transport calculation. As mentioned previously, the cylindricization and lack of angle-dependent cross sections are expected to introduce at least some error. However, this error is larger than expected, especially since SPH factors are employed. Thus, these discrepancies warrant greater investigation.

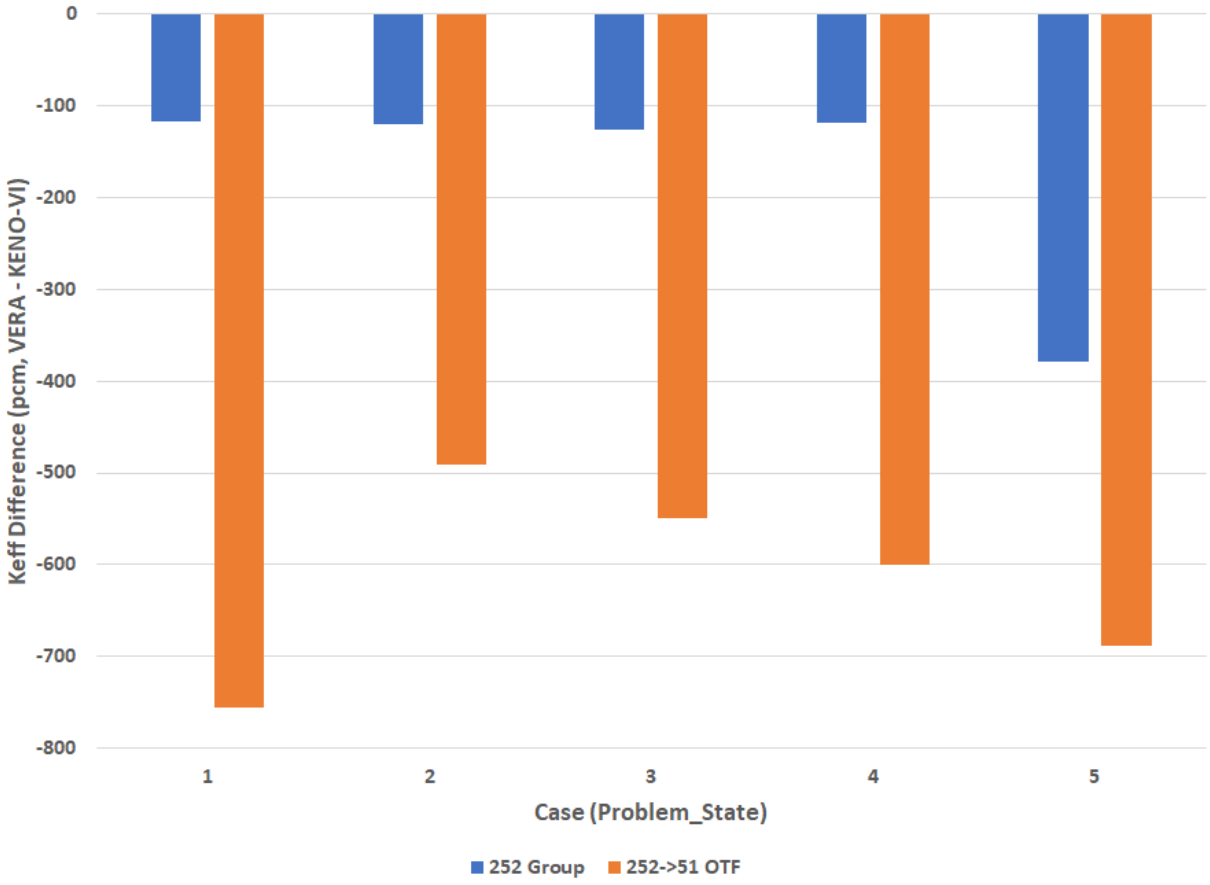

Figure 17. VERA problem 1 results for 252→51 OTF compared to KENO-VI

3.2.1.2 Problem 2 – Lattices

Figure 18 shows the k_{eff} differences for the 252-group calculations. Again, a larger negative reactivity bias is observed when collapsing from 252 groups than from 51 groups. However, in the case of the lattices, this has the effect of bringing some of the results much closer to the fine group transport. Where the 51→19 OTF was significantly overpredicting with strong absorbers, the 252→51 OTF is frequently around 100 pcm or less different from the fine group transport. However, this is surely cancellation of error.

Figure 19 shows the power distribution differences for the lattices. The trends in these results are very similar to those in Section 3.1.1.2. The largest errors are around 0.006 in lattices with strong thermal absorbers. Most lattices of a maximum error around 0.003 or less.

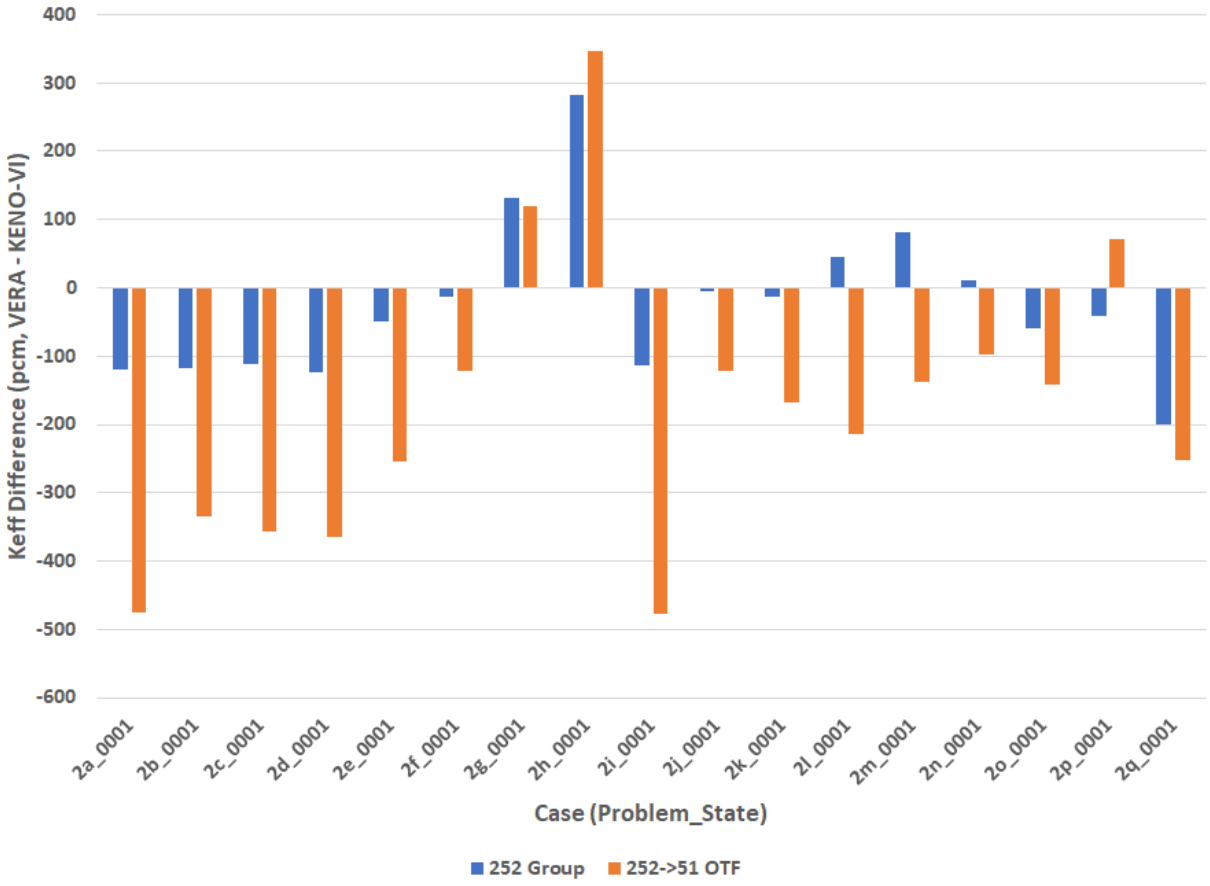

Figure 18. VERA problem 2 k_{eff} results for 252→51 OTF compared to KENO-VI.

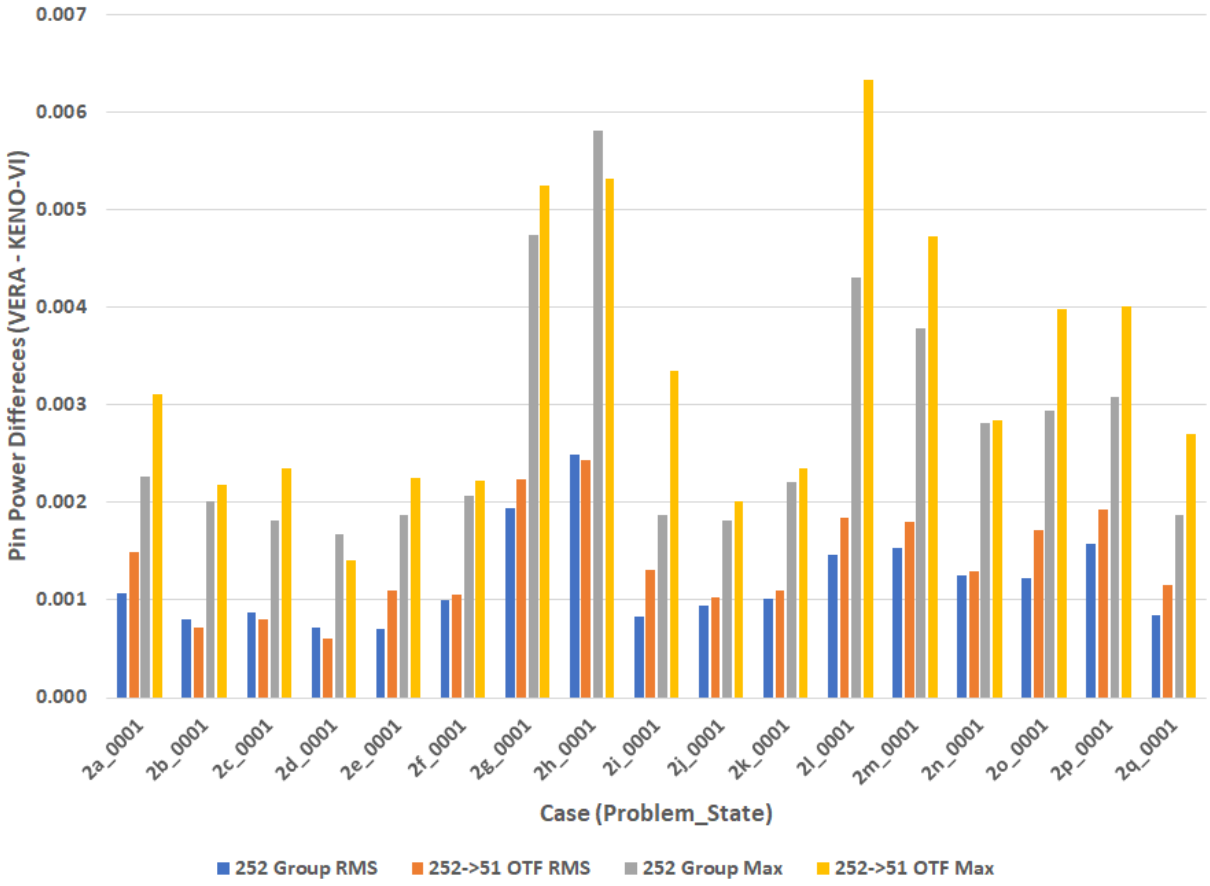

Figure 19. VERA problem 2 power distribution results for 252→51 OTF compared to KENO-VI.

3.2.1.3 Problems 4 and 5 – 2D Multi-lattice

Figure 20 shows the k_{eff} differences for the 2D variations of problems 4 and 5. The negative bias continues, though for these larger problems it is around 150 pcm instead of 300 pcm. Figure 21 shows the power distribution differences. The results for problem 4 are similar to those in Section 3.1.1.3. The results for problem 5 are similar for the fine group transport but significantly worse for OTF, with differences around 0.02. Figure 22 shows the difference between the fine group and OTF calculations for problem 5c-2D. For the 51 group calculations, the structures related to the burnable poisons, control rods, and radial reflector could be clearly seen in the differences. For the 252-group calculations, those structures are barely discernible; instead, we see a significant in→out tilt in the power distribution. This result indicates that the $\frac{1}{v}$ approximation in the radial reflector region causes greater problems for the 252-group library than for the 51-group library.

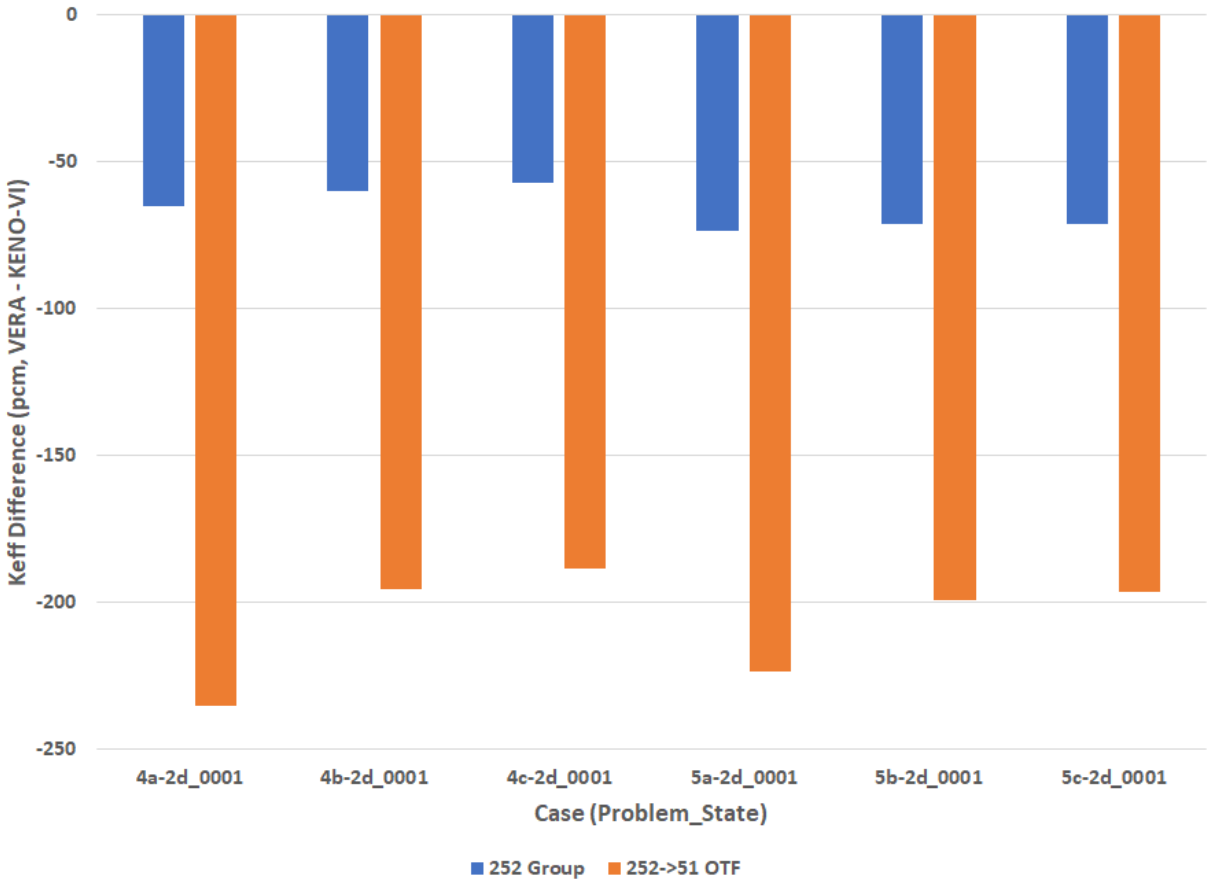

Figure 20. VERA problems 4 and 5 2D k_{eff} results for 252→51 OTF compared to KENO-VI.

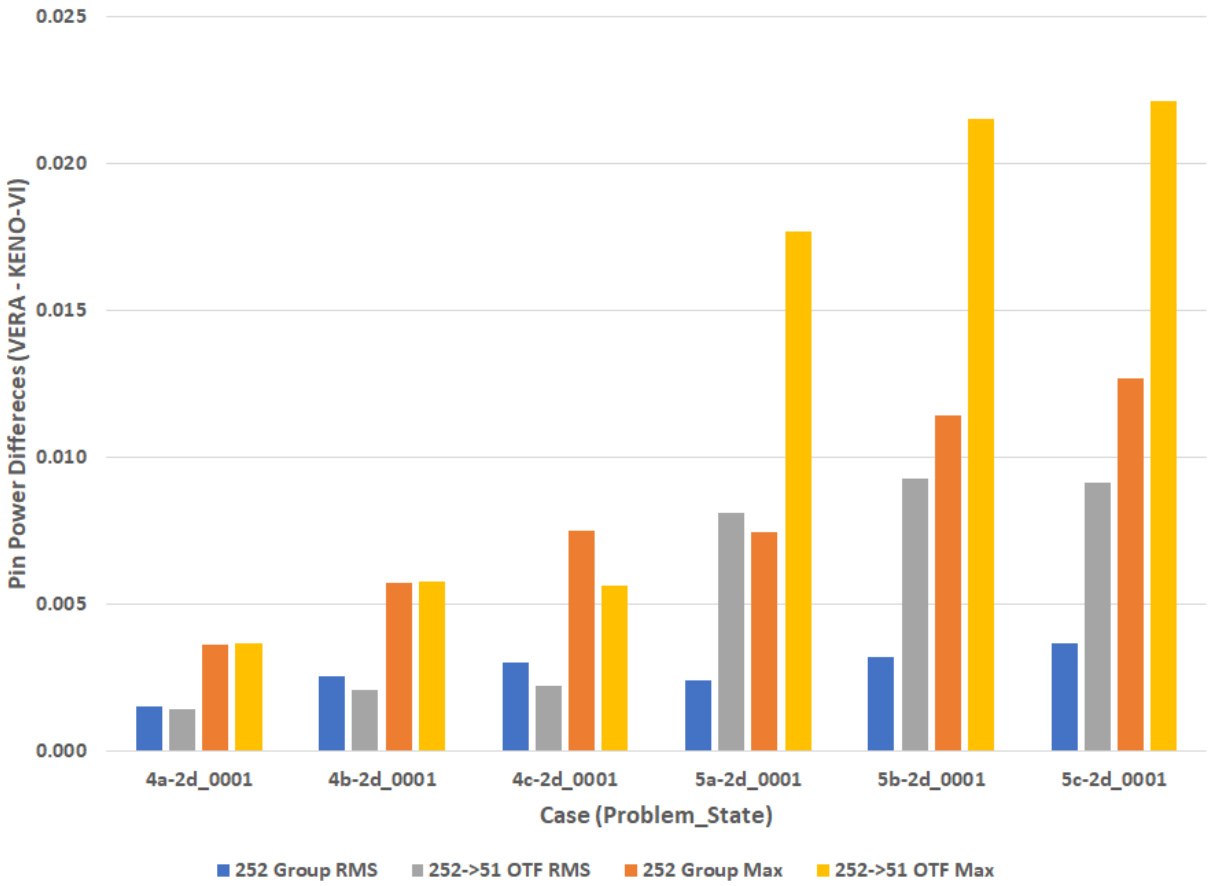

Figure 21. VERA problems 4 and 5 power distribution results for 252→51 OTF compared to KENO-VI.

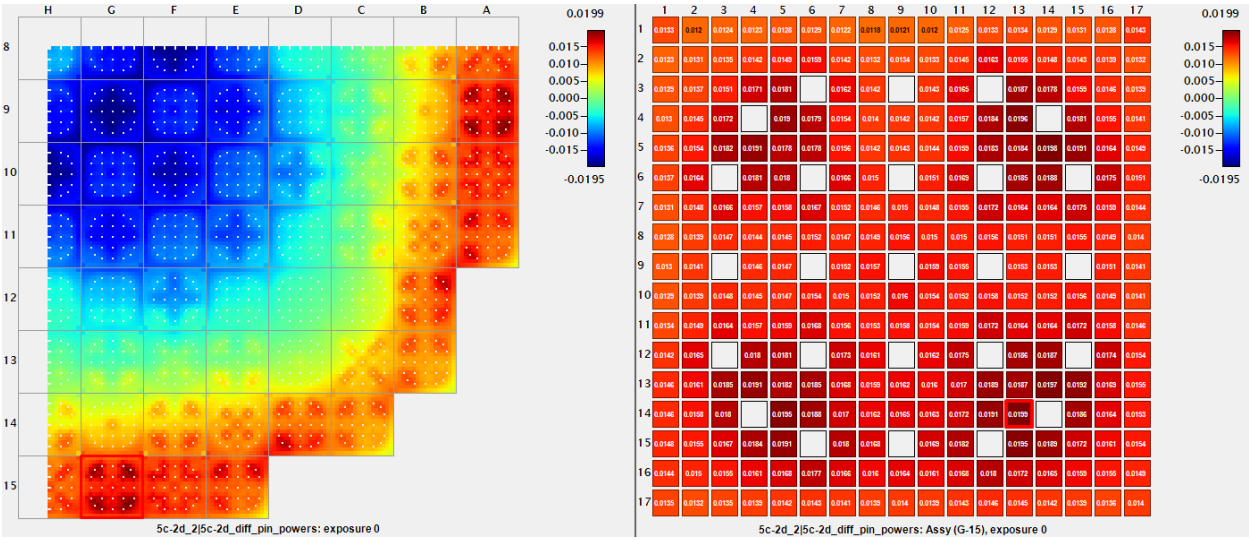

Figure 22. (OTF - fine group transport) for VERA Progression Problem 5c-2D.

3.2.2 3D RESULTS

3.2.2.1 Problem 3

The 252→51 OTF calculations for problem 3 are summarized in 10. Aside from the larger k_{eff} error, the results are very similar to those in Section 3.1.2.1, with slightly better power distribution comparisons. No detailed figures are shown for either case, as the results are very similar to others previously shown.

Table 10. Summary of results for VERA Progression Problem 3

Problem	252→51 OTF		
	k_{eff} (pcm)	Power Distribution	
		RMSE	Max.
3a	-321	0.001	0.004
3b	-206	0.002	0.006

3.2.2.2 Problem 4

Figures 23 and 24 show the problem 4 k_{eff} and power distribution differences, respectively. The maximum power distribution errors seem to show a slightly improvement over the 51-group errors, though the difference is quite small. The overall trends are the same as those shown Section 3.1.2.2, except that the reactivity is once again much more negative by about 300 pcm.

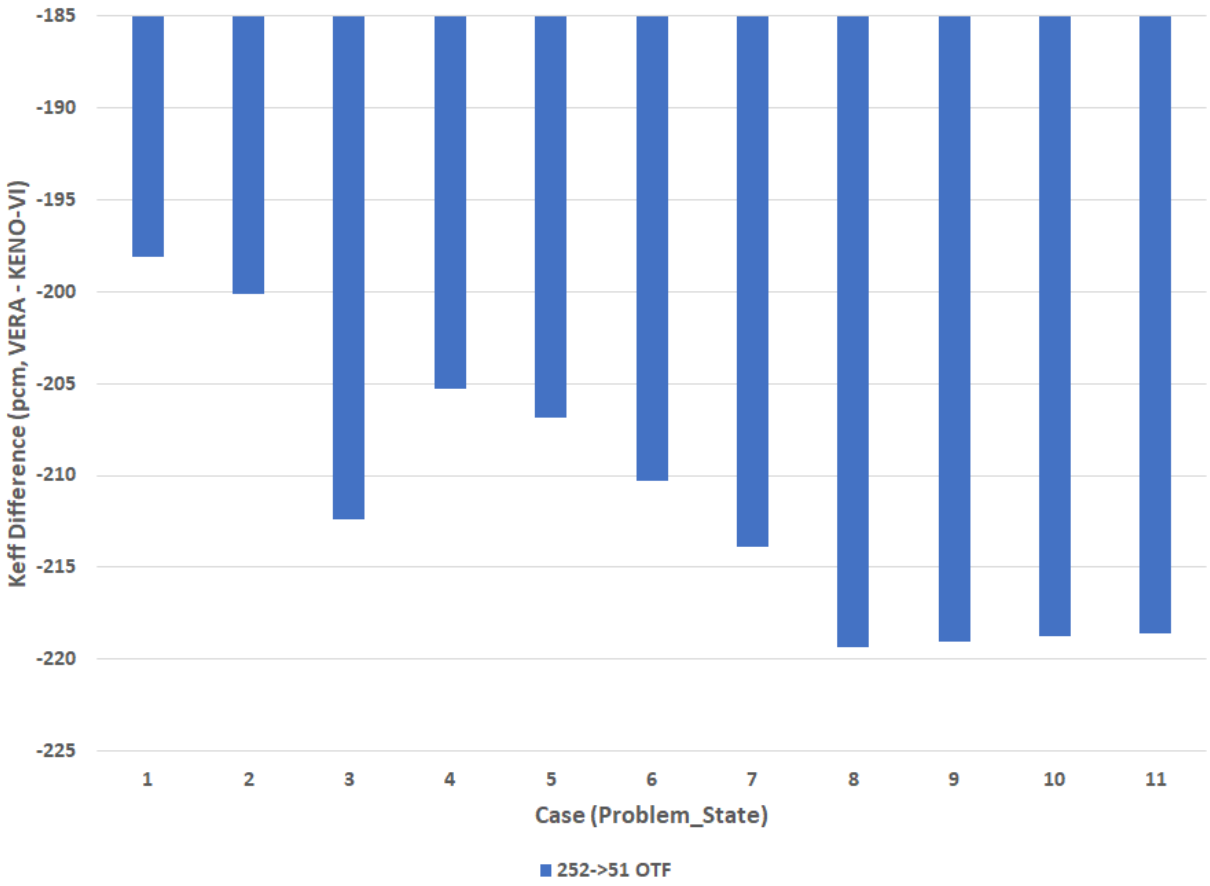

Figure 23. VERA problem 4 3D k_{eff} results for 252→51 OTF compared to KENO-VI.

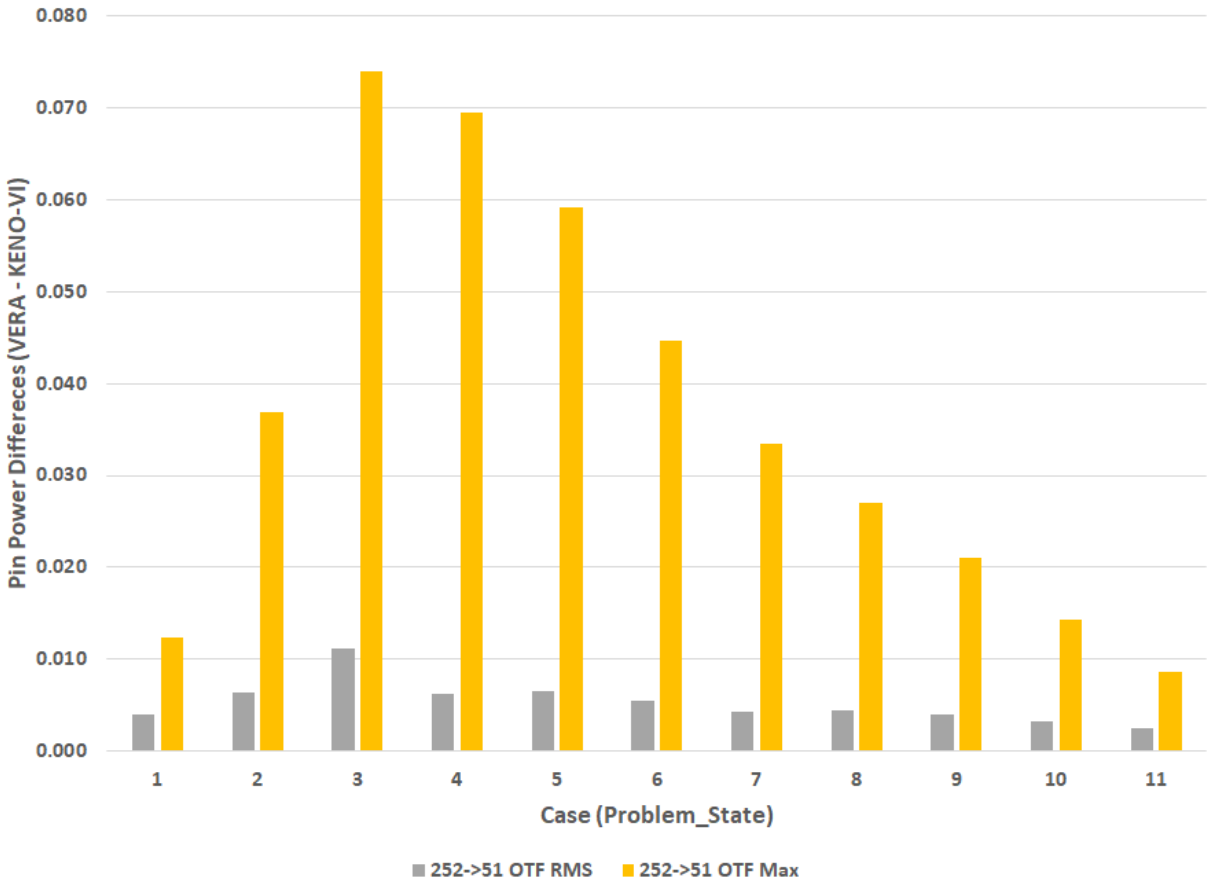

Figure 24. VERA problem 4 3D power distribution results for 252→51 OTF compared to KENO-VI.

3.2.2.3 Problem 5

Table 11 Shows the k_{eff} differences between KENO-VI and 252→51 OTF. The results are very consistent, with about 200 pcm difference between the two. Overall, this is good agreement. The agreement is worse than that for the 51 group comparisons but better than for 252→51 OTF for the smaller problems. This is in keeping with the the pattern discussed in previous sections. The power distribution differences could look better, but this is likely due to the reflector treatment, as seen in the 2D versions of problem 5. Overall, there are no major concerns, but there is room for improvement as well. No power distribution comparisons can be made for 31 of the 32 state points because this problem fails to converge for 252-group fine group transport, and there are no KENO-VI solutions available for those problems.

Table 11. Summary of VERA problem 5 OTF results compared KENO-VI

Problem	**252 →51 OTF k_{eff} (pcm)**
Criticals	216 ± 143
Rod Worths	206 ± 9
ITC	206 ± 0
Bank D IRW	208 ± 7
Total	210 ± 79

3.2.3 BWR RESULTS

Because the errors in the 51→19 OTF calculations were quite high for the BWR progression problems, it was deemed unhelpful to spend the computing time to run the 252→51 OTF calculations when the results will certainly not be very good. These calculations can be run in the future once accuracy improvements have been made to the BWR OTF capability.

4. 3D MULTIPHYSICS RESULTS

With validation of the OTF capability complete for cases involving only neutron transport, we can turn our attention to the extension of the OTF capability to multiphysics problems. This involves a far greater number of cross section condensation calculations. Additionally, multiphysics problems can be more challenging to converge; prior to this work, it was unknown whether OTF energy condensation in the neutron transport would impact the multiphysics convergence or not. This capability was tested on problems 6, 8, 9, and 10, which are summarized in Table 12; problems 7 was skipped because it is a beginning of life HFP calculation, which is already covered by problems 8 and 9. Instead, OTF calculations are compared to fine group transport calculations, and both sets of calculations are compared to plant measurements where available.

Table 12. Summary of the 3D Multiphysics VERA Progression Problems

Problem	Description
6	3D assembly at full power with thermal hydraulic (TH) coupling
8	Startup physics for HFP beginning of cycle (BOC) core; 49 total state points
9	Coupled depletion of WB1 cycle 1 (C1); 32 total state points
10	Core shuffle and refuel of WB1 from C1 to cycle 2 (C2) followed by HZP BOC simulation of C2

4.1 51→19 OTF

4.1.1 PROBLEM 6

The first of the multiphysics is VERA problem 6. This problem is a 3D coupled assembly, similar to problem 3a but with TH feedback. Comparing the fine group calculation with the OTF calculation gave a pin power difference of −56 pcm, RMSE power difference of 0.002, and a maximum power difference of 0.004. This is very similar to what was observed with problem 3, except that the power distributions actually compare better. This is actually a helpful side effect of the TH coupling. The TH feedback causes the power in the hottest rods to be suppressed to a greater extent than the power in the colder rods due to Doppler feedback. No such feedback exists without the TH coupling, which results in sharp flux gradients between pin cells. The TH feedback has the effect of smoothing the power some, especially between neighboring rods that vary in enrichment or burnable poison loading. Thus, the OTF capability is actually closer to the fine group transport when TH is enabled. Figure 25 shows the power distribution differences. The primary difference between this figure and the problem 3 results in Figure 11 is that the TH coupling has pushed the power downward in the core, causing the axial power difference not to have the cosine-like shape. However, the magnitude and features of the errors are similar in all other ways.

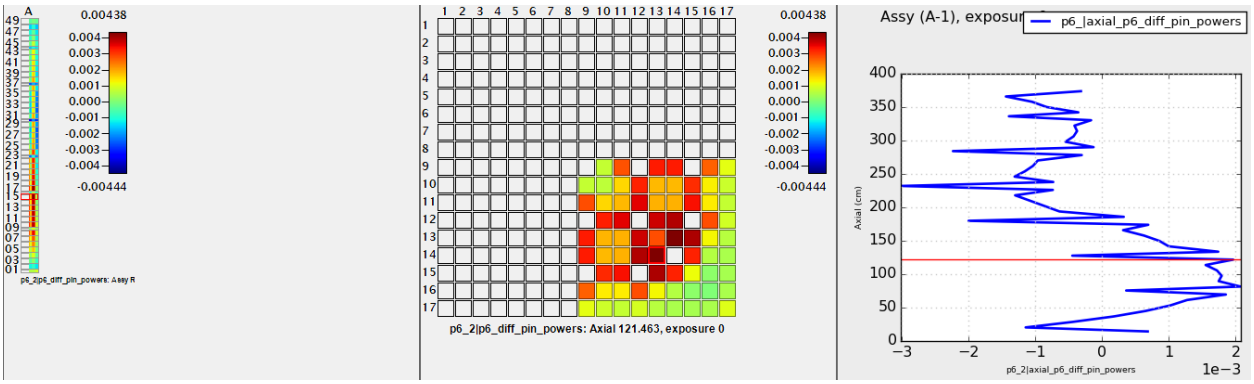

Figure 25. VERA problem 6 power distribution differences for 51→19 OTF compared to KENO-VI.

4.1.2 PROBLEM 8

Problem 8 consists of a series of startup physics tests performed at the WB1 plant. The tests were conducted over the course of 48 h as the control rods were withdrawn and power was gradually brought from 0% to 100%. This is modeled in VERA with 49 state points: an initial HZP calculation followed by forty-eight 1 h time steps. At each state point, the critical boron concentration is calculation, so this problem is essentially a series of 49 multiphysics criticals.

Figure 26 shows the results of the calculations, comparing 51→19 OTF to 51-group transport. The critical boron concentration difference starts at just under 7 ppm and decreases to around 4 ppm by the end of the calculation. This is comparable to a 70 to 40 pcm change in k_{eff} which is small and well within acceptable ranges. The power distribution is the worst near the beginning of the power ascent, with maximum errors just above 0.02. However, as the reactor approaches full power and operating temperatures, the TH feedback flattens the power distribution out and enables the OTF capability to more accurately approximate the fine group flux. The maximum power error at the end of the startup is only 0.01, half of what it was at the start. The RMSE power differences are around 0.0025–0.003 for each state point. Overall, this is excellent agreement and shows that OTF ought to be effective at the full range of conditions from HZP to HFP.

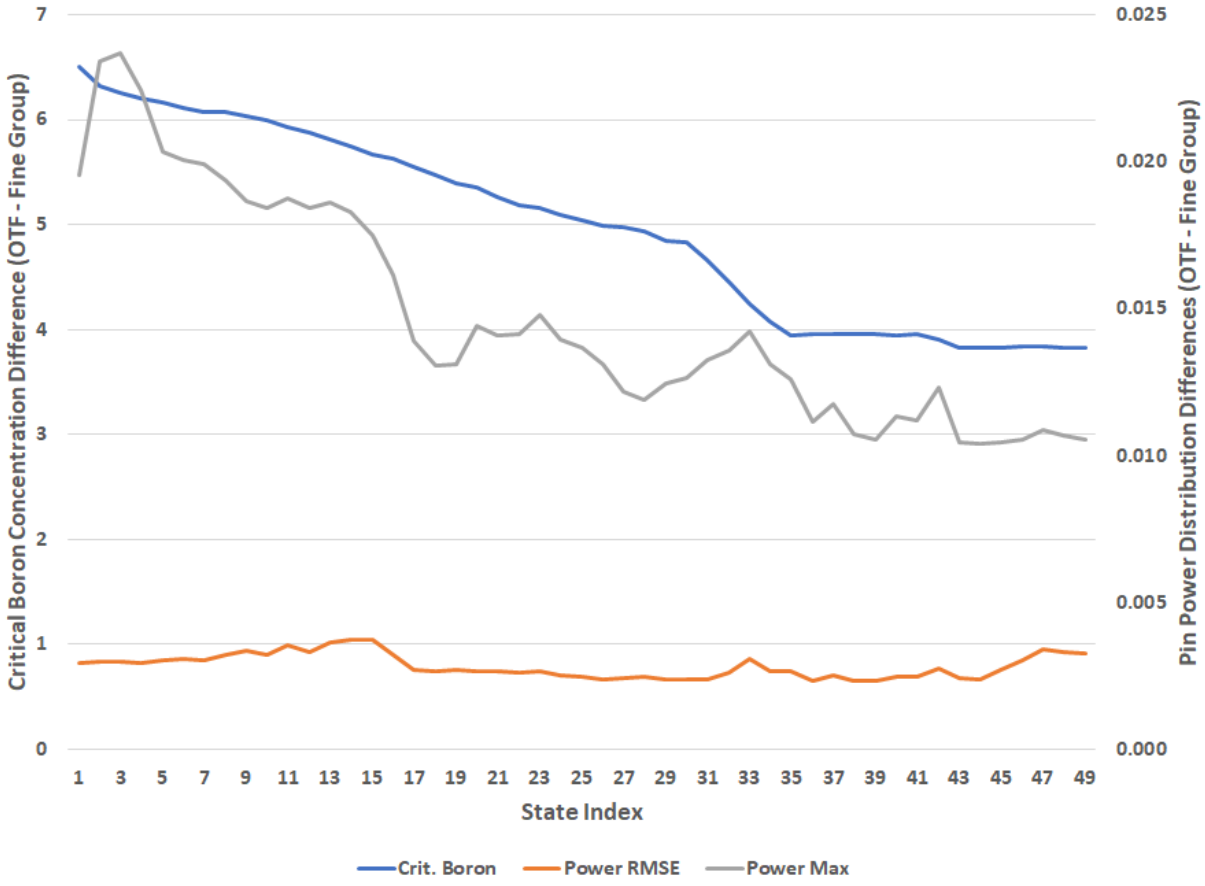

Figure 26. VERA problem 8 k_{eff} and power distribution differences for 51→19 OTF compared to fine group transport.

4.1.3 PROBLEM 9

Problem 9 is one of the most important for validating the OTF capability. It is a quarter-core cycle depletion with TH coupling and 32 state points. The TH was already incorporated in problem 6, but problem 9 also incorporates isotopic depletion. This will have a significant effect on the spectrum in the fuel. Furthermore, cycle depletion calculations are quite long and computationally expensive, so they could benefit greatly from the OTF capability if it performs well.

Figure 27 shows the boron letdown curves for the WB1 measured data, VERA fine group calculation, and VERA OTF calculation. Differences between fine group and OTF are also shown, as well as each of the two VERA calculations compared to the measurements. Although there is a consistent error between VERA and the measurement of −30 to −50 ppm, the difference between OTF and the fine group calculation are much smaller, ranging from about +4 ppm down to about −2.5 ppm. Assuming a typical value of about 10 pcm per ppm, this equates to less than 50 pcm differences throughout the cycle, which is excellent agreement between the two calculations.

Across all 32 state points of problem 9, the RMSE power distribution difference is 0.002, and the maximum averages 0.01. The maximum difference anywhere in any state point is 0.014 in the very first state. The end of the cycle also shows larger errors than the middle of the cycle, with errors around 0.012. Figure 28 shows power distribution differences between OTF and fine group transport at various points in

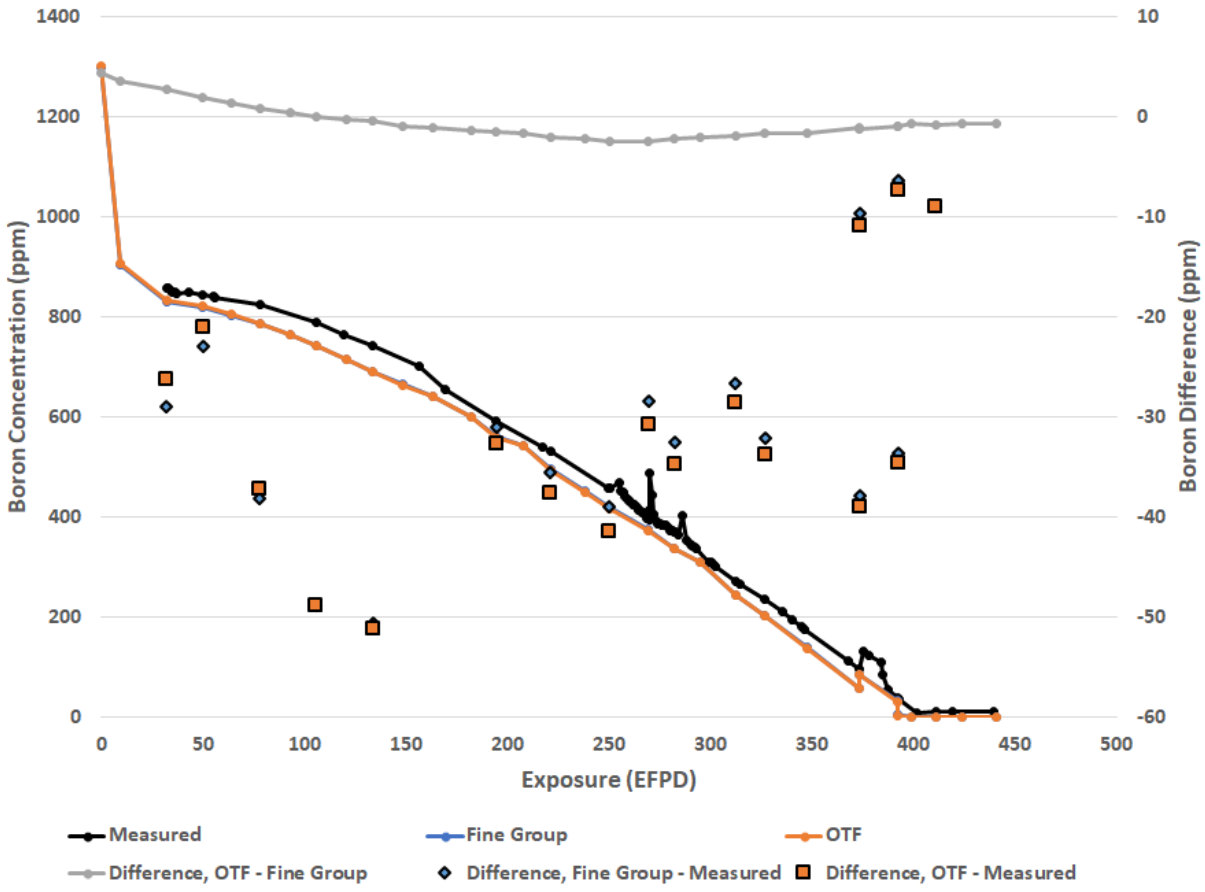

Figure 27. VERA problem 9 boron letdown curves.

the cycle. The largest error is in the first state index, which is 0% power. By the third state, those errors have reduced significantly. This is smaller because of the TH coupling, which flattens the power distribution by suppressing power in the hottest pins. Thus, the third state is more representative of BOC results for a coupled core. At that point, there is still a noticeable overprediction of power in the so-called ring of fire portion of the core, in which the periphery and interior are both underpredicted. By 200 effective full-power days (EFPD), the over-depletion of overpredicted rods has flattened out the power distribution differences similarly to how TH coupling did at BOC. At this point, the dominant errors are caused by the reflector treatments. The axial power is overpredicted at the bottom of the core, and the radial power has an out→in tilt compared to the fine group transport. At end of cycle (EOC), things look similar, the major difference being that the control rods are fully withdrawn such that rod-related effects are seen in the power distribution. Overall, the comparisons are good, especially regarding consistency throughout the cycle. The most notable issue continues to be the reflector.

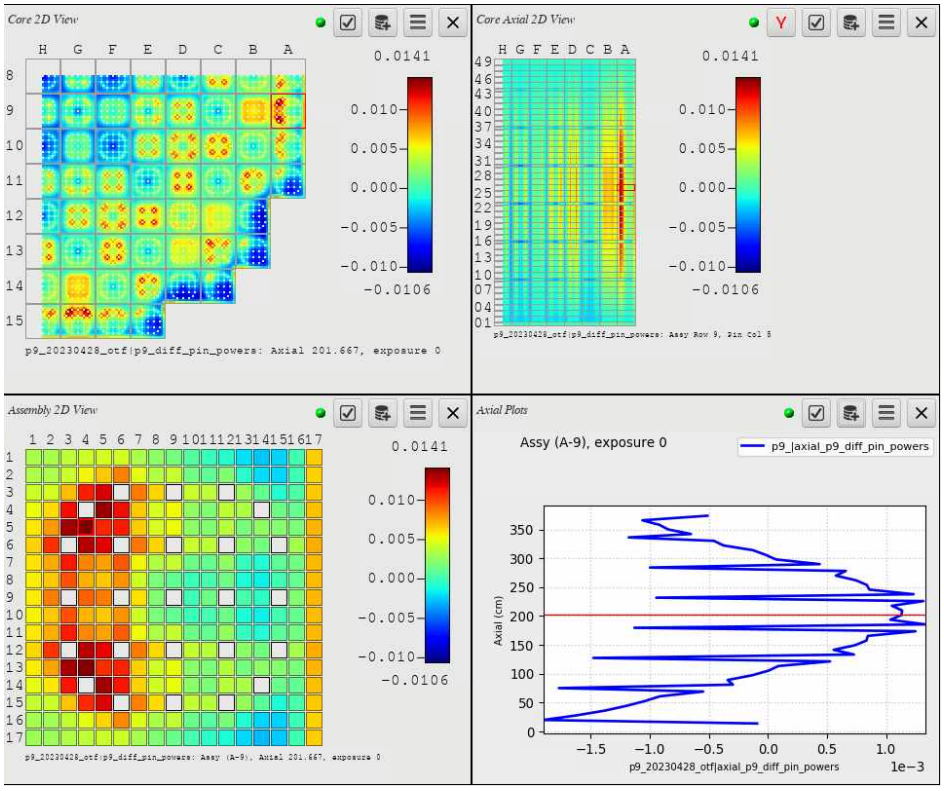

(a) State 1, 0 EFPD, 0% Power

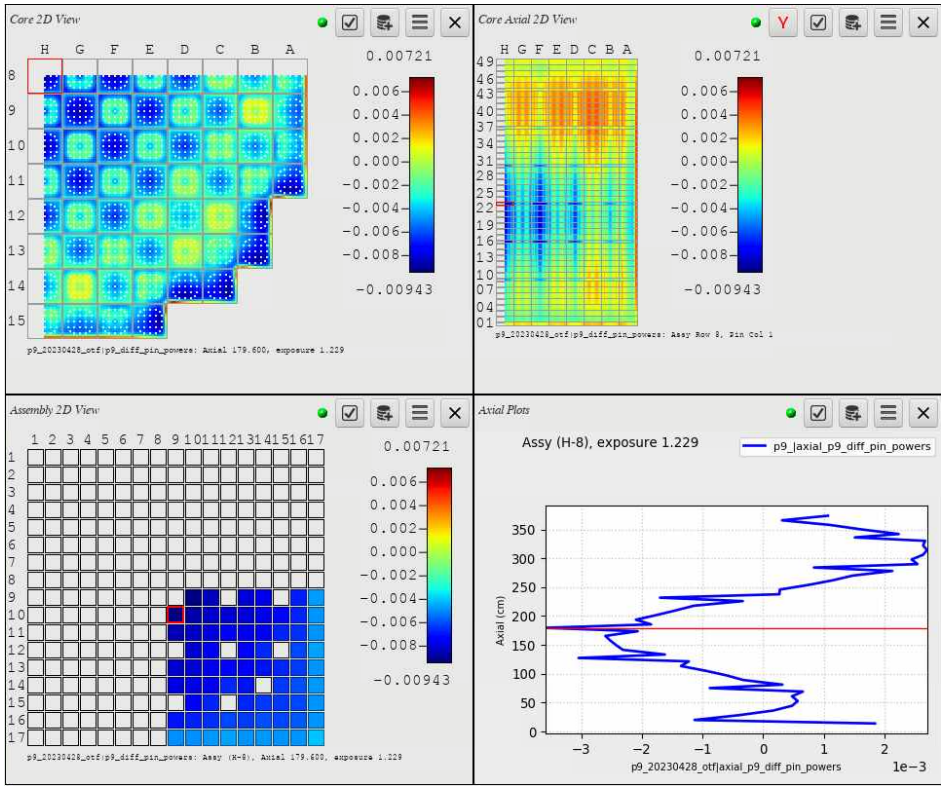

(b) State 3, 32 EFPD, 99.7% power

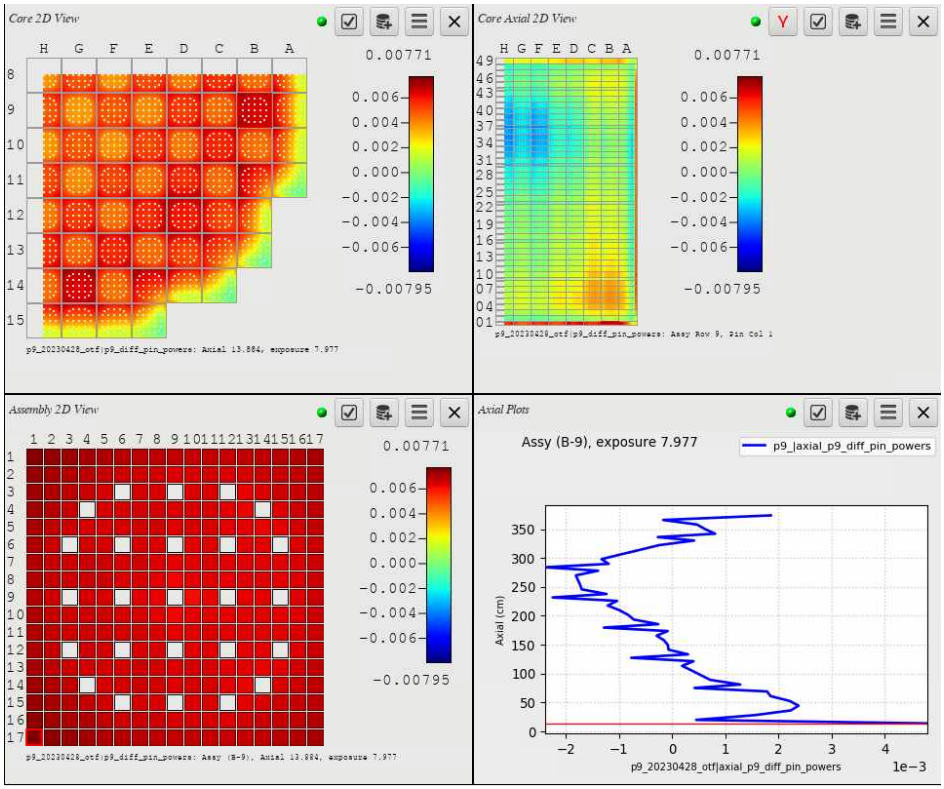

(c) State 15, 207.7 EFPD, 93.9% power

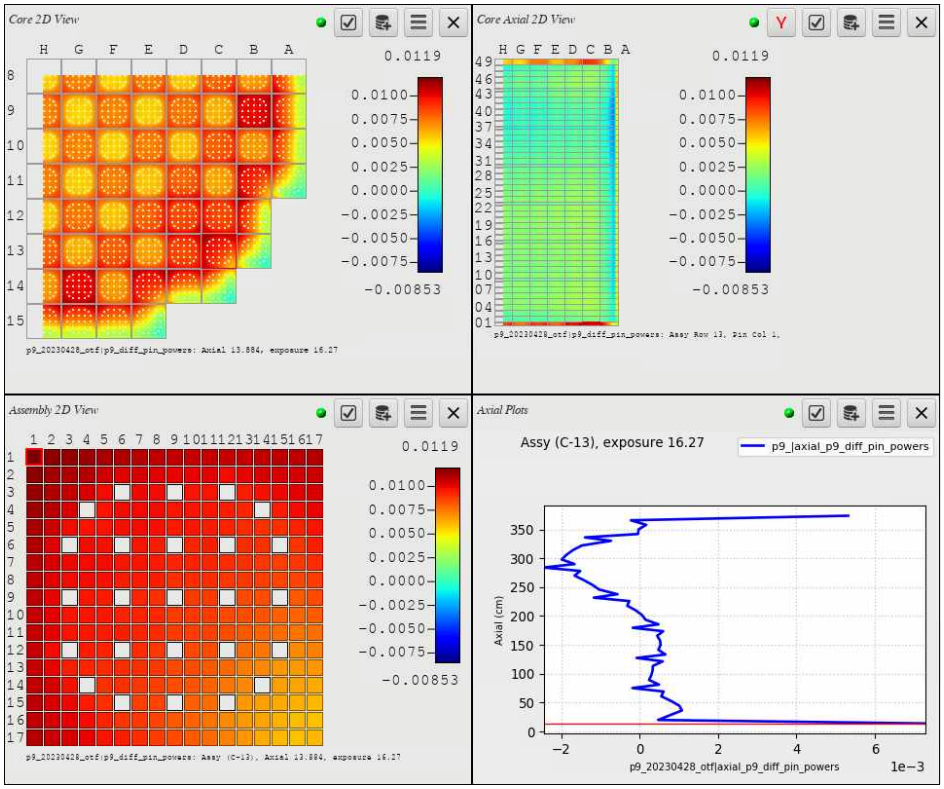

(d) State 31, 423.6 EFPD, 78.8% power

Figure 28. VERA problem 9 power distribution differences, sliced at largest differences.

4.1.4 PROBLEM 10

Problem 10 is simply BOC conditions for cycle 2 of WB1, having shuffled the burned fuel from cycle 1. Comparing OTF to fine group transport, the k_{eff} difference is less than 1 pcm. The power distribution differences are 0.007 RMSE and 0.035 maximum. This is larger than for problem 9. Figure 29 shows slices of the power distribution differences and exposure distribution at the maximum and minimum power difference locations. The maximum is the most notable difference; it occurs near the top of the core in a fresh bundle, B-11, specifically in the northwest quadrant of that bundle that neighbors burned fuel, C-10. Additionally, it is just below the top of some WABA inserts in C-10 that do not extend through the entire active fuel region, and there are IFBA in the fresh fuel at C-11. These features will introduce more severe flux gradients than would be observed at the start of problem 9, which was a fresh core. Additionally, there has been a clear trend of increased errors around burnable poisons; thus, a difference of 0.035 is not shocking. Once again, this shows the weakness of the isolated pin cell treatments.

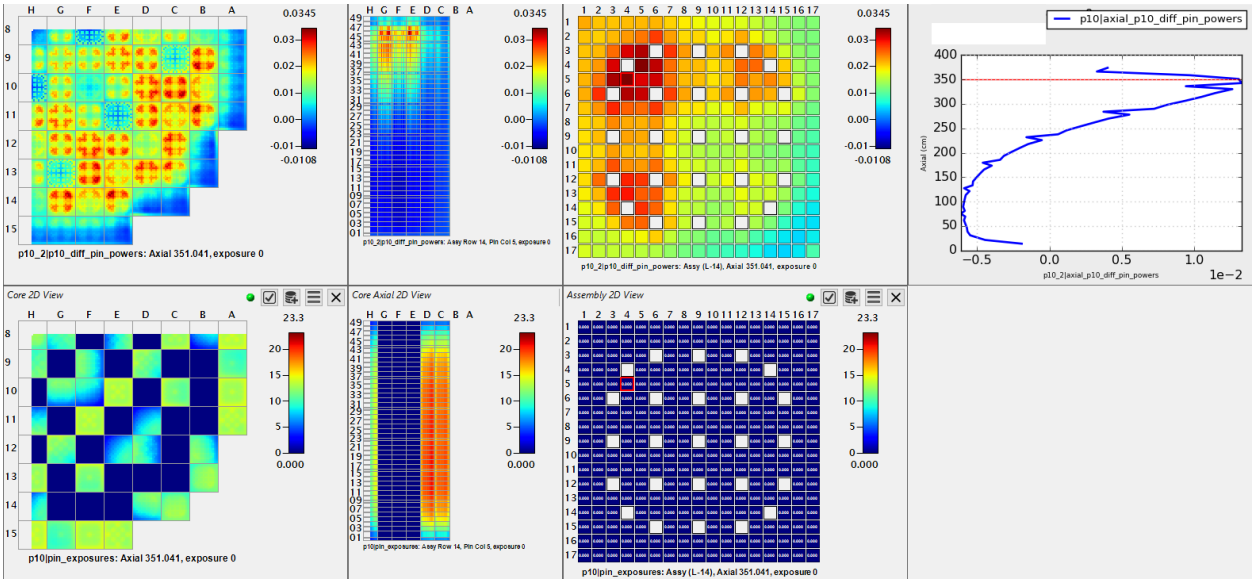

(a) Slice at maximum difference

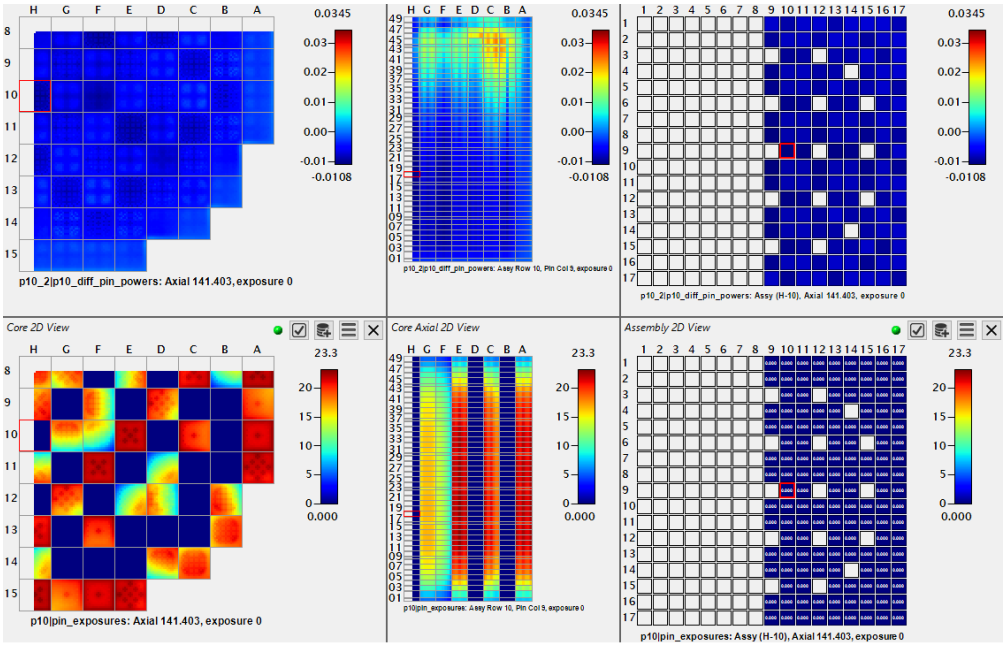

(b) Slice at minimum difference

Figure 29. Comparison of fine group and OTF VERA problem 10 calculations.

4.2 252→51 OTF

The 252 group calculations could only converge these 3D calculations if OTF was used; for the sake of time and computing resources, problem 9 is the only multiphysics calculation shown for the 252→51 OTF capability. This calculation was very slow, completing only 9 of 32 state points. The boron concentration comparisons for these 9 state points are shown in Figure 30. This comparison is rather limited, but without the fine group calculation, little more can be shown. For the 51-group comparisons in Figure 27, the difference between OTF and the measurements ranged between about −30 and −50 ppm. For this comparison, those results are more negative by 10–15 ppm. This continues a trend of negative reactivity bias in 252-group OTF calculations.

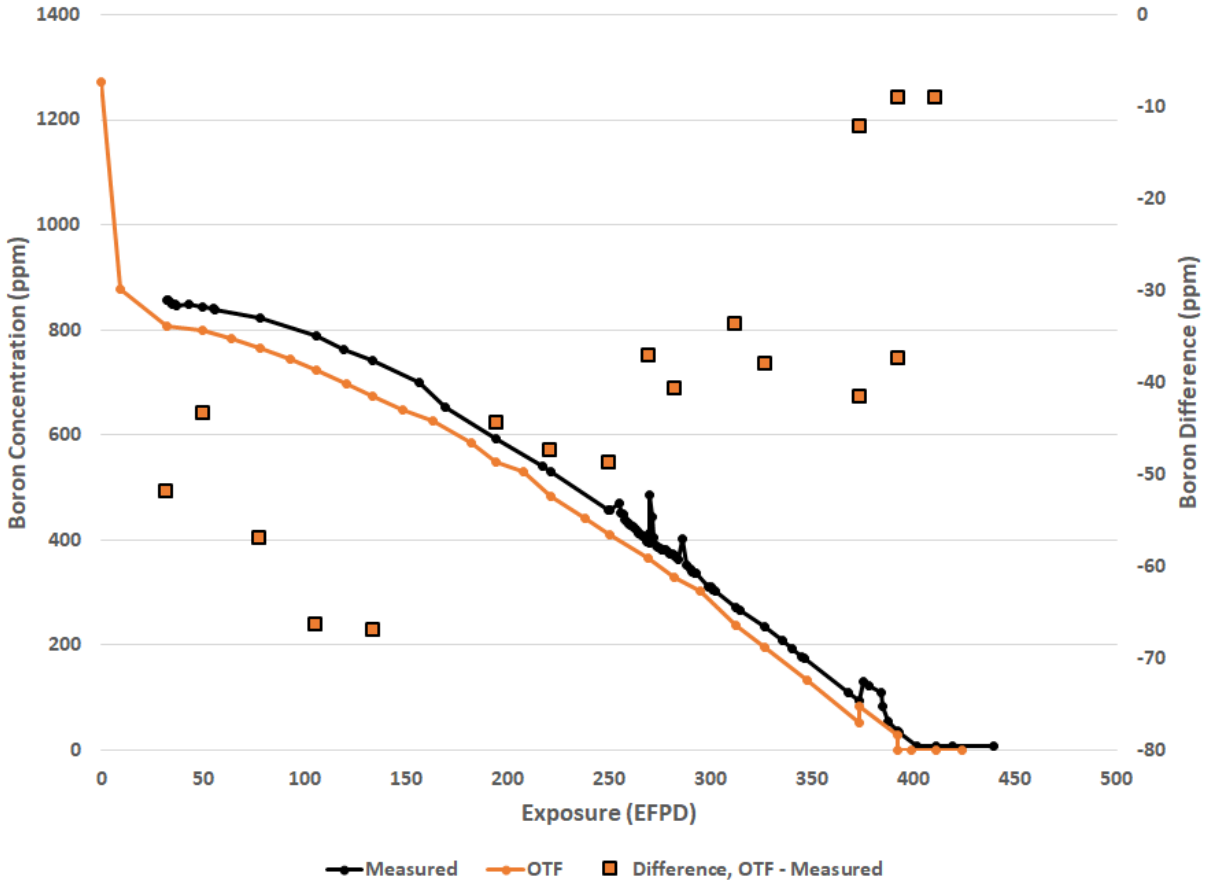

Figure 30. VERA problem 9 boron letdown curves.

5. TRANSIENT MULTIPHYSICS RESULTS

5.1 NULL TRANSIENT

The first calculation was to simply perform a null transient to identify obvious issues with the code. This transient was done with a pin cell, and it did in fact expose some issues with the code, specifically related to the neutron velocities. Previously the velocities were stored only in the cross section library. However, when using OTF, the velocities are needed in the coarse group structure. Furthermore, collapsing the velocities should be done using the appropriate spectrum. This causes the neutron velocities to become region-dependent since the spectrum is different in each region. The code was modified to perform an OTF collapse of the velocities from the fine group library structure to the coarse group structure. With this complete, the code could successfully run in a null transient in which there was no significant change in power.

5.2 PROBLEM 4-MINI ROD EJECTION

The next transient test is for a miniature version of VERA problem 4. This problem reduces the height of the fuel and changes from 17×17 assemblies with 24 guide tubes to 7×7 assemblies with 4 guide tubes. A control rod is ejected from the center assembly of this simplified model, inducing a rapid power rise. These types of transients are often extremely computationally expensive because of the small time steps required for numerical stability during the power rise; thus, this test constitutes an important test for the OTF capability.

Figure 31 shows the core power for the rod ejection test. Overall, the OTF capability is similar, but it underpredicts the power increase. Whether this underprediction is too great may be application-dependent. Some additional investigation may be warranted to determine whether there is a way to improve the transient accuracy. The pin power distribution is different by less than 0.5% for every position and time step, so it is primarily the magnitude of the power that is an issue.

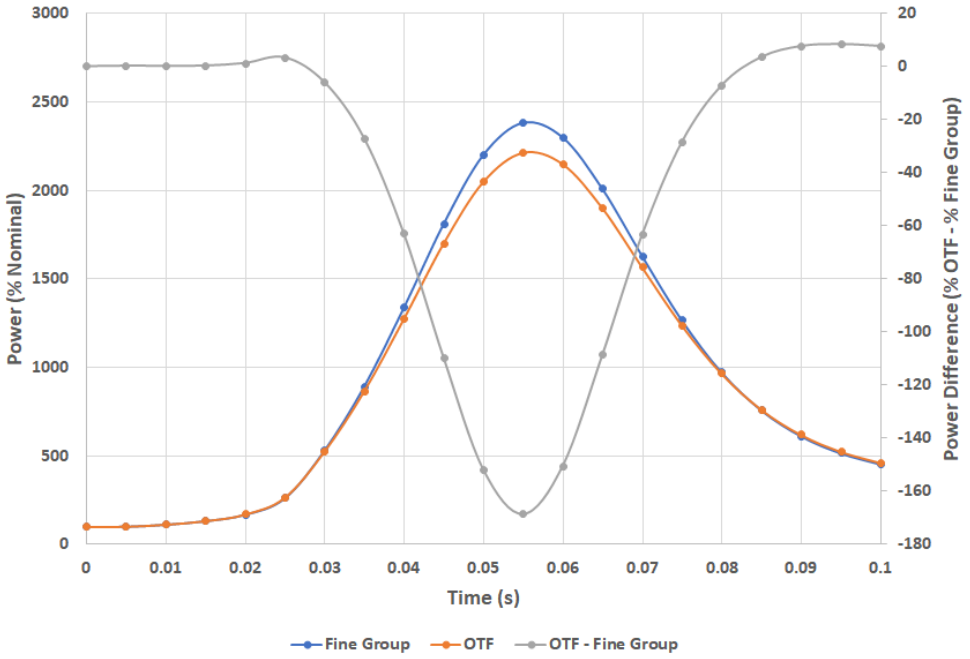

Figure 31. Core power for fine group and OTF transport for problem 4-mini rod ejection.

5.3 PROBLEM 4-MINI SCRAM

This test uses the same problem 4-mini model but instead immediately SCRAMs the model at the start of the transient. This tests the OTF capability for a transient with a rapid drop in power. This model also performs decay heat calculations such as would be needed to support certain long-running transients such as LOCAs. This is another important application for the OTF capability because these transients can run for minutes instead of just several seconds like a rod ejection.

The 51-group fine group calculation required 397 seconds; the 51→19 OTF calculation required 295 seconds. This is a speedup of about 26%. This is less than that for similar steady-state problems, largely because the transient calculation involves many cross section updates compared to the steady-state calculation. For this particular problem, cross sections—and thus pin cell spectra—were updated three times each time step, in contrast to once per steady-state coupled iteration. This is required for the multilevel timestepping procedures used for the transient calculation. Although this multilevel timestepping is necessary, it does have the side effect of eating into the speedup that OTF would otherwise demonstrate for transient calculations.

For this type of calculation, the quantity of greatest interest is the decay heat after shutdown. For all the time steps in the calculation, the maximum difference in decay heat between the fine group and OTF calculations was just 0.065 W/cm. This is a trivial difference and indicates that using OTF to accelerate the post-shutdown transient calculation should have no detrimental impacts on the accuracy of the decay heat prediction.

5.4 WATTS BAR UNIT 1 CYCLE 3 SCRAM

Recent work to support FFRD analysis using VERA has focused on integration of decay heat calculations with TH coupling and transient simulations [15]. In this work, one of the decay heat demonstrations was to SCRAM WB1 cycle 3 at EOC to obtain decay heat profiles as a function of time. The details of this

calculation are left to the reference, which is publicly available. This calculation is one of great interest currently in the Nuclear Energy Advanced Modeling and Simulation (NEAMS) program. It is a long-running calculation in VERA and is, therefore, a good candidate for the OTF capability. Thus, this problem was rerun using the 51→19 OTF capability to determine whether the accuracy was similar and how much faster the calculation is.

The original calculation made it through 350 s of transient time before hitting the 48 h walltime. The OTF calculation made it through 495 s of transient time in the same walltime; it about 40.5 h for the OTF calculation to make it to 350 s, which is about 1.18× speedup. As mentioned previously, the speedup is reduced by the number of times cross section updates are required; there is likely room for improvement in the calculation flow and cross section update logic to remove some of these excess updates.

The primary quantity of interest for this calculation is the decay heat. For the steady-state solve, the decay heat started out at about 0.07% difference between OTF and 51-group transport. This peaked around 0.9% 1.85 s into the transient. This is early in the transient while the SCRAM is still executing, so the differences in the transport solution because of the rod movement influence the decay heat buildup. After this point, the differences drop to as low as 0.3%, then stay between 0.3% and 0.5% for the remainder of the transient. Overall, this is a small difference in the decay heat, indicating that the OTF capability can be safely used to generate decay heat curves with high accuracy.

The total power behaves similarly, with differences peaking at 1.3% just before the peak decay heat differences. These differences then drop rapidly to a fraction of a percent. Some larger fluctuations in the total power differences are seen throughout the transient, but they occur late enough in the transient such that the total power is very low, so these differences no longer matter.

6. ON-THE-FLY ENERGY CONDENSATION PERFORMANCE AND DEMONSTRATION

6.1 TIMING AND MEMORY TRENDS

The speedup trends for VERA problems 1–5 are shown in Figure 32, with speedup defined as the fine group transport runtime divided by the OTF runtime. MOC is essentially linearly in the number of energy groups. CMFD involves inversion of a matrix, which scales cubicly in the number of energy groups in the most extreme case; however, VERA uses the generalized minimum residual (GMRES) linear solver method with 100 restart vectors and a limit of 20 power iterations, In practice, the CMFD scaling is somewhere between linear and quadratic for large calculations. This plays out as expected in the speedup trends because the smaller problems tend to have speedups limited by the ratio of the number of groups—2.68× for 51→19 and 4.94× for 252→51—but can exhibit greater speedup for large problems such as problem 5-3D. The 51→19 OTF showed speedup of nearly 3.5× for problem 5-3D because of the increased runtime of CMFD for the larger 3D calculations. Unfortunately, the 252-group calculations fail to converge for 3D calculations, so speedup cannot be calculated for 252→51 OTF for any of the 3D calculations.

The speedups of the lattice calculations suffer, especially for the 252→51 OTF. This is because the lattice calculations are so fast that the solve time becomes dominated by the spectral calculations more than for other problems. Optimization of the spectral calculations could improve this speedup significantly.

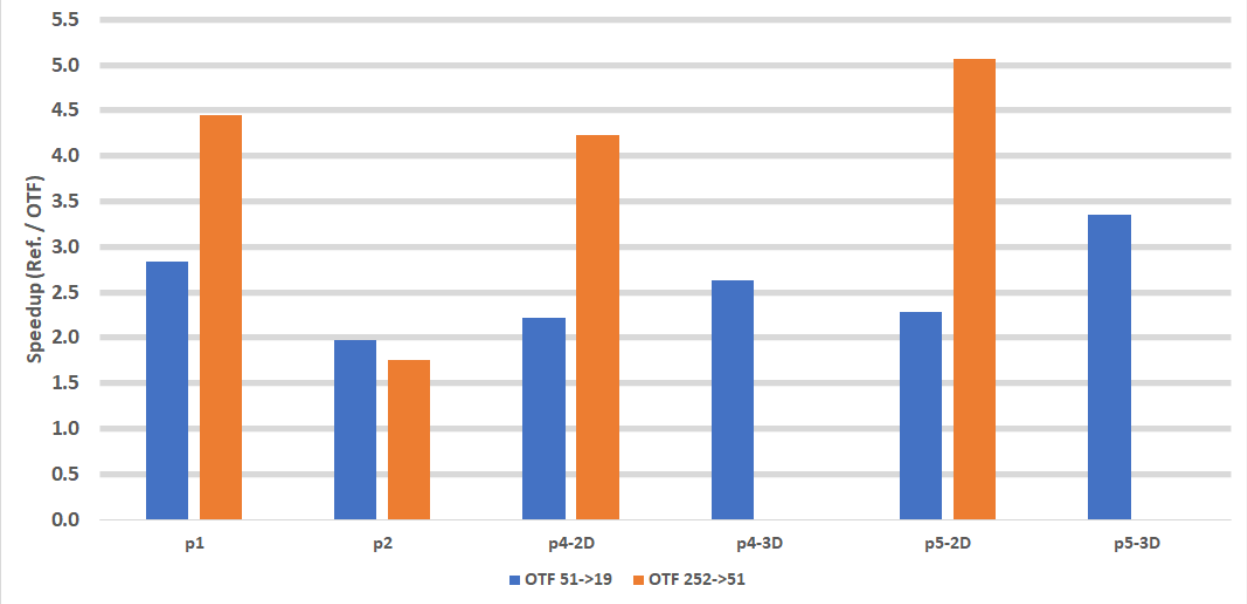

Figure 32. Speedup for 51→19 and 252→51 OTF for neutronics-only calculations.

Although speedup was the primary goal of the OTF work, a secondary benefit is reduced memory usage. The OTF calculations slightly increase the solution storage because they store the fine group spectrum everywhere and the the coarse group flux solution, but they dramatically reduce the cross section and angular flux boundary condition storage. Table 13 summarizes the memory usage for a selection of representative cases: 2a, 5-3D, 6, and 9. For non-coupled cases, 51→19 OTF reduces the memory by around one third, while 252→51 OTF reduces the memory by 40%–60%, depending on the size of the problem. For the coupled cases, the memory improvements are smaller because the TH solver uses non-trivial memory that is unaffected by OTF. For 51→19 OTF there is a reduction of about 20% for both

p6 and p9, whereas for 252→51 OTF there is 45% reduction for p6 and 58.6% reduction for p9. The 252-group calculations require so much memory that the savings scales close to the number of groups. The presence of TH coupling is less noticeable in the memory usage of large 252-group calculations.

Table 13. Summary of memory requirements for fine group and OTF calculations

Case	Cores	51 Groups			252 Groups		
		51 Groups (GB/core)	51→19 OTF (GB/core)	OTF Reduction	252 Groups (GB/core)	51→19 OTF (GB/core)	OTF Reduction
p2a	1	0.238	0.154	35.3%	1.023	0.601	41.3%
p5-3D	464	2.670	1.909	28.5%	16.116	5.958	63.0%
p6	29	0.392	0.312	20.4%	1.907	1.049	45.0%
p9	928	1.934	1.548	20.0%	9.137	3.783	58.6%

6.2 PROBLEM 9 TIMING BREAKDOWN

6.2.1 51→19 OTF

Figure 33 shows the timing breakdowns for problem 9 with 51 groups and 51→19 OTF. The first half of the plot shows the percentage of the solve that was spent in each part of the calculation. The total time using 51 groups was 19.35 h. The second part of the plot shows the percentage of the OTF solve spent in each part of the calculation *as a percent of the total 51-group solve time*. Thus, an extra piece of the pie called *Savings* is added for the reduction in the total solve time. The total time for OTF was 17.38 h, which constitutes a savings of just over 10%.

Examining the second pie chart demonstrates how much faster individual pieces of the code are with the 51→19 OTF. The MOC calculation went from 14.6% to 4.8%, which is about a 3.04× speedup; the CMFD solve went from 32.1% to 9.1% for a speedup of 3.53×; and axial transport went from 3.1% to 1.3% for a speedup of 2.38%. The self-shielding, depletion, and feedback calculations all experienced only minor speedup, which is to be expected since they only loosely depend on the number of energy groups. Finally, the cross section generation portion of the calculation is represented by the expanded piece. In this section, the pin cell building increased from 1.3% to 7.0% because in the 51-group calculation, the pin cells are rebuilt only when shielding calculations are performed; but for the OTF calculations, pin cells must be rebuilt every time cross sections are updated. For the 51-group calculation, pin cells were rebuilt only 160 times out of 546 iterations; for the OTF calculation, pin cells were rebuilt in every one of the 604 iterations, leading to a substantial increase in the time spent rebuilding them. The "XS Generation" label refers to the calculation of fine group cross sections in the 51-group calculation; and for the OTF calculation, it also includes generation of buffer region fine group cross sections and collapsing the fine group cross sections down to the coarse group structure after completion of the spectrum calculations. The increased number of iterations and extra work to collapse cross sections account for the increase from 16.2% to 19.2%.

Finally, the "Spectrum Generation" component exists only in the OTF calculation. This component includes both the fine group spectrum calculation and the coarse group CPM calculations to obtain SPH factors. The time spent in this portion of the code is 15.1% of the total time for the 51-group calculation. Thus, we have a 15% brand new calculation time that did not exist at all in the 51-group calculation but still achieved more than 10% speedup. Considering that the efforts so far have been focused only on accurate implementation of the OTF capability, this is a promising result. The total time spent in generating cross sections for the transport calculations increased from 17.5% of the total run time for the 51-group calculation to 46.5% of the OTF calculation. Efforts should be focused on this area for further optimization.

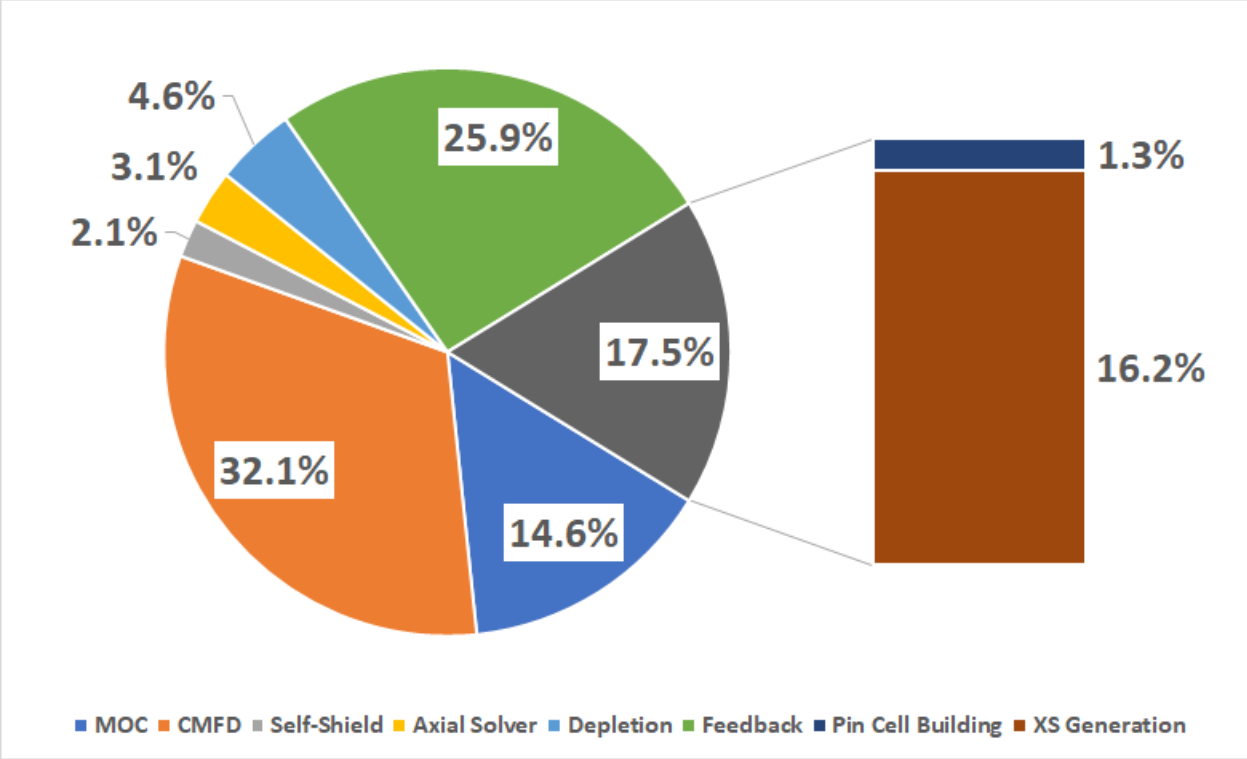

(a) 51 group timing breakdown

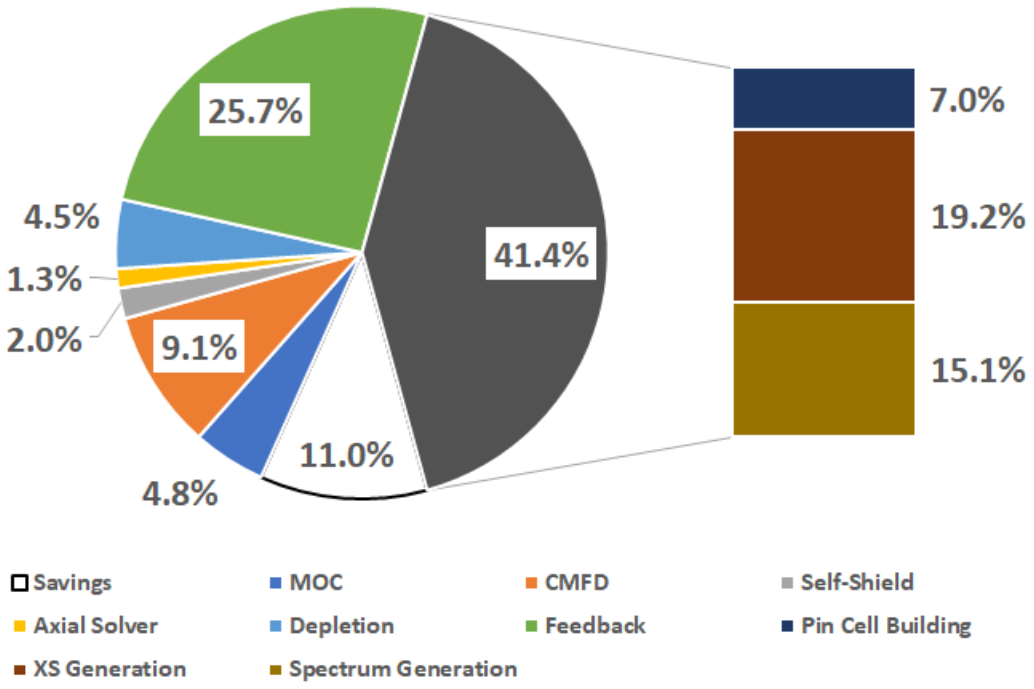

(b) 51→19 OTF timing breakdown as percentage of 51 group calculation

Figure 33. Timing breakdown for VERA problem 9.

6.2.2 252→51 OTF

Without the OTF capability, any 3D calculation with the 252-group library could suffer numerical stability issues. Even if one managed to get the calculation to converge, the computational expense of the calculation made doing so intractable for large problems. Problem 9 is not realistic with the 252-group library and would likely require greater than an order of magnitude increase in computing resources and walltime. However, the OTF capability managed to run this calculation with the 252-group library, though it is still a slow calculation.

Figure 34 shows the timing breakdown for 252→51 OTF; all percentages are shown as a fraction of the total runtime for the calculation since a fine group calculation could not be completed with 252 groups. The total calculation time was about 5.5 days on 2784 cores. While it is encouraging that it's possible to run this calculation now, this is still intractable for a cycle depletion. However, it is important to note that 90% of the runtime was spent generating cross sections, and almost 75% was spent specifically in the spectrum generation. This means that the actual core computations only took about half a day. Once again, there is significant room for improvement in the performance of the cross section generation. If the cross section generation time could be improved by a factor of 5×—which seems reasonable given the disproportianate time currently spent in that part of the code—then the total calculation time would drop to about 1.5 days, a much more reasonable amount of time to deplete a cycle starting from 252 groups.

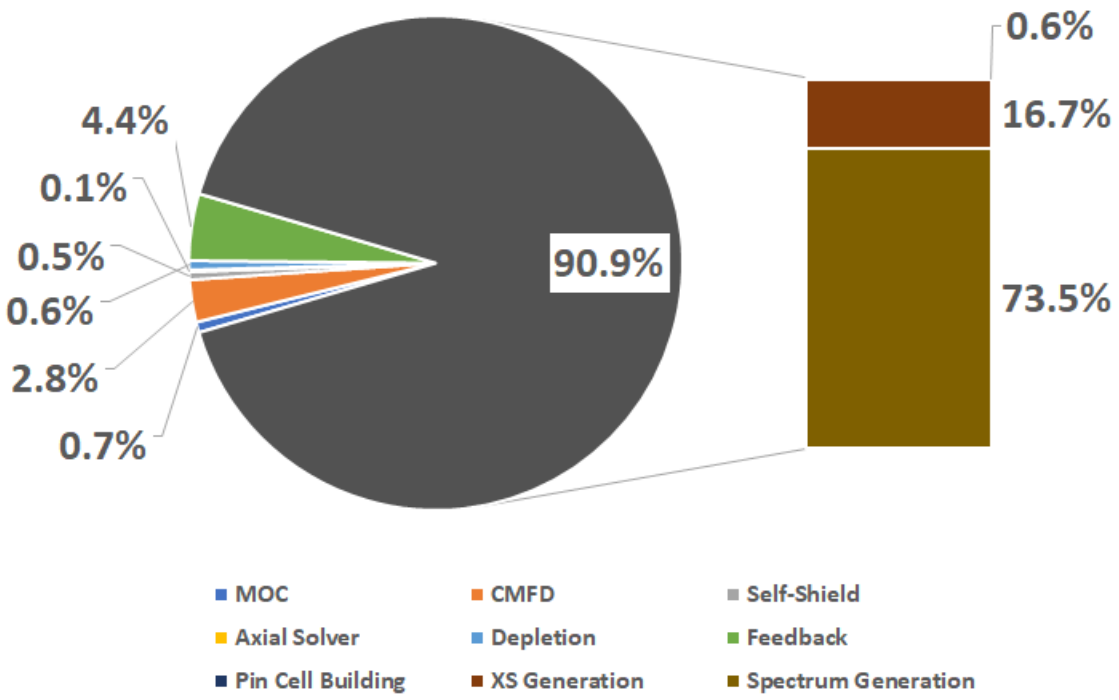

Figure 34. Timing breakdown for VERA problem 9 using 252→51 OTF

7. DISCUSSION

The OTF energy condensation capability has been successfully implemented in MPACT and used for multiphysics coupling calculations within VERA. This capability has been demonstrated on a range of PWR and BWR problems, with and without TH feedback, and for two different sets of group structures. The validation for neutronics-only calculation showed good agreement for PWR cases; the BWR cases had reasonable k_{eff} comparisons, but the power distributions suffered. The multiphysics validation showed better agreement in many cases because the TH feedback flattened the power shapes some, lessening the impact of the the isolated pin cell approximation in the cell spectrum calculations. However, interfaces between burnt and fresh fuel did show some larger differences for BOC in problem 10. This warrants some additional validation to see whether those differences remain at operating temperatures.

Several transient calculations were run that showed good agreement, though the OTF capability did underpredict the power a bit. This should be investigated to determine if it can be fixed to obtain a more accurate prediction of core power. Overall, the power and decay heat shapes agreed well and should generally follow the same patterns seen in the steady-state calculations. Thus, for post-shutdown transients such as LOCAs, the OTF capability will perform very well; the same cannot currently be said for events such as a reactivity insertion accident (RIA).

The performance of the OTF capability was very good overall. The speedup consistently approached the expected speedup, which is the ratio of the number of fine groups to the number of coarse groups. The primary performance-related issue remaining is dealing with the expense of the spectrum calculations. This expense was not terribly important for neutronics-only calculations, but for coupled calculations the repeated spectrum updates were more expensive. Even with these repeated update, several hours were saved on problem 9 calculations, which was a better result than the authors anticipated for a first pass on those calculations. The memory savings due to OTF could also be important for allowing more energy groups, more isotopics, or calculations on smaller computers. For neutronics-only calculations, a speedup approach the ratio of fine groups to coarse groups should be expected for most cases; the speedup will be much less than this for multiphysics calculations because of the repeated spectrum generation.

Regarding the application of the OTF capability to FFRD analysis, it is expected that the capability would perform well for PWRs right now. However, some additional validation should be done on the capability for highly burned fuel to ensure that no new issues arise in such a scenario. Since current FFRD analysis focuses on LOCA situations, the mis-predictions of core power during transient scenarios should not be important. For BWRs, some additional work is warranted to improve the accuracy of OTF. It is expected that improved modeling of the radial reflector and bypass flow regions would significantly improve the results. Additionally, LOCAs are limiting accidents for PWRs, but for BWRs, RIAs—specifically, blade drops—are one of the limiting accidents. Thus, the issues with core power prediction during transients will be much more important for BWRs. Thus, the OTF capability should not be considered ready for BWRs FFRD analysis or for any PWR transient that has a significant power rise.

8. FUTURE WORK

The OTF implementation has been quite successful overall, but there are a number of areas for improvements. These can be classified into two categories: accuracy and performance.

8.1 ACCURACY

The most obvious accuracy shortcoming is the negative reactivity bias. For the 51-group calculations, the biases were around −100 pcm for VERA problem 1 and tended to get smaller with larger problems. This would be acceptable, but it is surprising to see such errors for pin cells given the use of SPH factors. Additionally, the errors are much more significant for the 252-group library, as great as −150 pcm even for p5-2D. Thus, these errors should be investigated to determine if they can be reduced more effectively.

A second significant area of improvement is in the treatment of the reflector regions. The axial power shapes seem to show that the axial reflector is probably not as important, but the radial reflector shows significant errors in the rods near the baffle for the 51 group calculations. For the 252 group calculations, the reflector treatment seems to induce a power tilt for the entire core. The $\frac{1}{v}$ approximation should be replaced with a more robust physics-based calculation. One plausible option would be to use a 1D Cartesian CPM calculation on the baffle, then apply that spectrum to all the baffle pin cells.

Related to handling the baffle, improvements could also be made to better handle BWR geometry, specifically channel boxes and control blades. It is unclear what the best approach for this would be, but perhaps approximating control blades as 1D slabs would be sufficient and more accurate than the current cylindricization method.

Another area for improvement is related to strong absorbers. The results showed quite clearly that both the k_{eff} and power distribution accuracy were negatively affected by the presence of strong absorbers, especially thermal absorbers such as B_4C or gad (as opposed to AIC control rods). It is unknown at this time why the CPM calculations do not resolve these spectra as well as expected. It is also interesting that the error is not always the same direction: IFBA rod powers are underpredicted, gad rod powers are overpredicted, and rods near pyrex or WABA are overpredicted. It is possible that there are multiple effects that need to be identified.

Spacer grids also seemed to have some unusual behavior. The dips in power at the spacer grid locations were much larger for OTF than for the fine group calculation. This should be investigated to determine whether there is a way to improve this or not.

Transient calculations produced good agreement between OTF and fine group calculations for the power and decay heat shapes. The agreement was not as good for the core power level. This must be investigated since an underprediction of the core power level undermines the reliability of the OTF capability.

The final area of accuracy improvement is the group structure. The target 19-group structure used for most of the calculations in this report was developed as a modified form of a more well-known 18-group structure. However, the accuracy of this modified structure was not tested for accuracy in any way. It is likely that some errors could be reduced simply by selecting a more robust group structure.

8.2 PERFORMANCE

Throughout this work package, efforts were made to optimize the calculation of macroscopic cross sections in VERA and the CPM calculations. Significant improvements were seen in those efforts. However, it has been postulated by the authors that many of the CPM calculations are redundant. At any point in time, it is

likely that the spectra of several pin cells scattered about the model are very similar. If these pin cells could be identified and grouped such that only one CPM calculation was required for the group of cells, then this could result in significant speedup for the OTF capability, as well as for the WSA self-shielding calculations. There would likely be significant work involved in doing this to ensure that the pin cells are group appropriately without negatively impacting error, so this optimization was not attempted as part of this work package.

It is also likely that the spectra do not need to be updated every single time the cross sections are updated. There are probably several iterations at the end of each state point where the pin cell spectrum does not change significantly even though the global problem has not fully converged. Thus, it is expected that the spectra could simply be reused after a few iterations and save significant time.

9. 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.

The authors also wish to acknowledge Ben Collins (Veracity Nuclear, LLC) and Will Gurecky (Oak Ridge National Laboratory) for several extended discussions about this capability and possible performance improvements. These discussions helped the authors to prioritize various facets of this work and generate ideas for follow-on work.

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] 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.

[3] Nathan Capps, Aaron Wysocki, Andrew Godfrey, Benjamin Collins, Ryan Sweet, Nicholas Brown, Soon Lee, Nicholas Szewczyk, and Susan Hoxie-Key. Full core LOCA safety analysis for a pwr containing high burnup fuel. *Nuclear Engineering and Design*, 379, 2021.

[4] 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.

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

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

[7] Kang Seog Kim, Mark L. Williams, Dorothea Wiarda, and Kevin T. Clarno. Development of the multigroup cross section library for the CASL neutronics simulator MPACT: Method and procedure. *Annals of Nuclear Energy*, 133:46–58, 2019.

[8] Kang-Seog Kim, Aaron Graham, and Matthew Jessee. Dancoff-based Wigner-Seitz Approximation for the Subgroup Resonance Self-Shielding in the VERA Neutronics Simulator MPACT. In *Proc. of PHYSOR 2022*, 2022.

[9] Hansol Park and Han Gyu Joo. Practical resolution of angle dependency of multigroup resonance cross sections using parametrized spectral superhomogenization factors. *Nuclear Engineering and Technology*, 49(6):1287–1300, 2017.

[10] William Boyd, Nathan Gibson, Benoit Forget, and Kord Smith. An analysis of condensation errors in multi-group cross section generation for fine-mesh neutron transport calculations. *Annals of Nuclear Energy*, 112:267–276, 2018.

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

[12] Chase Lawing, Scott Palmtag, and Mehdi Asgari. BWR progression problems. Technical Report ORNL/TM-2020/1792, Oak Ridge National Laboratory, 2021.

[13] J. Leppanen, M. Pusa, T. Viitanen, V. Valtavirta, and T. Kaltiaisenaho. The serpent monte carlo code: Status, development, and applications in 2013. *Annals of Nuclear Energy*, 84:142–150, 2015.

[14] SCALE Development Team. SCALE: A comprehensive modeling and simulation suite for nuclear safety analysis and design, version 6.1. Technical Report ORNL/TM-2005/39, Oak Ridge National Laboratory, 2011.

[15] 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.