ORNL/TM-2023/2903

Coupled Decay Heat and Thermal Hydraulic Capability for Loss-of-Coolant Accident Simulations

Aaron M. Graham
Andrew T. Godfrey

May 10, 2023

OAK RIDGE National Laboratory

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

ORNL/TM-2023/2903

Nuclear Energy and Fuel Cycle Division

COUPLED DECAY HEAT AND THERMAL HYDRAULIC CAPABILITY FOR LOSS-OF-COOLANT ACCIDENT SIMULATIONS

Aaron M. Graham
Andrew T. Godfrey

Date Published: May 10, 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

ABBREVIATIONS

API	application programming interface
BOC	beginning of cycle
C3	cycle 3
CASL	Consortium for Advanced Simulation of LWRs
DOE	Department of Energy
EFPD	effective full-power days
EOC	end of cycle
FFRD	fuel fragmentation, relocation, and dispersion
HBu	high-burnup
HZP	hot zero power
I/O	input and output
LHR	linear heat rate
LOCA	loss-of-coolant accident
LWR	light-water reactor
MOC	middle of cycle
NEAMS	Nuclear Energy Advanced Modeling and Simulation
ORNL	Oak Ridge National Laboratory
ppm	parts per million
RIA	reactivity insertion accident
RMSE	root mean square error
TH	thermal hydraulic
VERA	Virtual Environment for Reactor Applications
WB1	Watts Bar Unit 1

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 for longer durations. 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 has traditionally been allowed in these reactors. Thus, 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 being leveraged to determine the susceptibility of HBu fuel to fuel fragmentation, relocation, and dispersion (FFRD). Improvements have previously been made to fuel performance capabilities to more accurately model these phenomena [2]; multiphysics simulations have also been conducted to determine the power and burnup histories of the HBu fuel [3], which are needed as inputs for 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].

Prior LOCA simulations have relied on the TRACE systems code [7], which can model the core and primary loop during accident conditions. TRACE includes many models for various aspects of the primary loop, but two sets of models are important for this report. First, TRACE uses a lumped-fuel approach for modeling the core. This approximates the ~50,000 fuel rods in the core with a much smaller number of rods. The rods can be lumped in various ways as determined by the user. For example, one lumped rod may be used to represent all rods in an assembly, sometimes with an additional rod representing the hottest fuel rod. However, due to runtime constraints and complexity of modeling, a more common approach is to group several assemblies or larger regions of the core into single lumped rods. These lumping schemes apply not only to fuel rods but to flow channels as well. Second, TRACE has several different models for treating decay heat, ranging from pregenerated decay heat curves based on an ANSI/ANS-5.1 standard [8] (hereinafter abbreviated simply as ANSI) to explicit time-dependent heat inputs from the user. None of these models account for differences in isotopics between different rods, which is an approximation the work in this report seeks to eliminate.

This report focuses on the implementation of coupled decay heat capabilities in the Virtual Environment for Reactor Applications (VERA) code suite [9] to address a gap identified in previous LOCA simulations. This constitutes an improvement for both the lumped-fuel and decay heat models in TRACE. VERA has been developed to perform high-fidelity, whole-core multiphysics simulations for LWRs. Previously, during the Consortium for Advanced Simulation of LWRs (CASL) program, the emphasis was on providing accurate steady-state analysis—with a secondary focus on reactivity insertion accident (RIA) analysis—to address operational challenges in the nuclear energy industry. Under the Department of Energy (DOE) Nuclear Energy Advanced Modeling and Simulation (NEAMS) program, these capabilities are being extended to a broader range of transient analyses with the goal of quantifying the risk of fuel failures such as FFRD. To properly model such conditions with VERA, decay heat calculations have been integrated with the multiphysics to enable rod-by-rod thermal hydraulic (TH) conditions to be driven by the decay heat in long-running accidents such as LOCAs.

2. CODES AND METHODS

2.1 MPACT

MPACT is the neutron transport module of VERA. It uses the 2D/1D method to solve 3D transport problems for nuclear reactors [10]; it also acts as the multiphysics driver for VERA depletion and transient calculations [11]. MPACT calculates sub-pin–resolved neutron flux, typically using around 50 axial levels and 3 rings per level (10 rings for fuel with gadolinia). Isotopics and burnup are calculated on a similar mesh. These quantities can be sent to other codes for multiphysics coupling. The neutron transport library used by MPACT contains cross sections in 51 energy groups for 263 isotopes [12].

2.2 ORIGEN

ORIGEN is the nuclide transmutation and decay module in SCALE [13]. ORIGEN has been coupled with MPACT and can be called as a subroutine for every mesh region in the MPACT model. This results in sub-pin–resolved isotopic distributions at every axial level. The full isotope set for ORIGEN contains 2,226 isotopes. Because of VERA's long run times, a reduced isotope set with 263 isotopes was generated for use when coupled with VERA [14]. This reduced set contains all the isotopes that are important to reactivity in a reactor. Unimportant isotopes were eliminated from the decay chains, and fission yields and decay constants were modified to account for the eliminated isotopes without explicitly modeling them. This results in decreased run time for coupled VERA calculations without sacrificing accuracy.

2.3 CTF

CTF is a subchannel TH code that can be used to solve reactor problems [15]. It contains a wide range of correlations that allow it to solve single- or two-phase flow for both normal operating conditions and accident scenarios. CTF is called by MPACT as a subroutine, with MPACT sending pin-by-pin axial power distributions to CTF and receiving fuel temperatures, coolant temperatures, coolant densities, and other similar TH quantities. Prior to the work presented in this report, the pin powers sent to CTF by MPACT were always calculated as follows:

$$P'_{r,z} = \frac{\sum_{r' \in (r,z)} \sum_{g=1}^{G} \kappa\Sigma_{f,r',g} \phi_{r',g}}{\overline{P}'} \,, \tag{1a}$$

$$\overline{P}' = \frac{P_{total}}{\sum_{r \in R} \sum_{z in Z} l_{z,r}} \,, \tag{1b}$$

where subscript r, z is rod index r at elevation z, subscript r' is an MPACT mesh region index, subscript g is the energy group index, G is the total number of energy groups, R is the total number of rods, Z is the total number of axial levels, $\kappa\Sigma_f$ is the heat production from fission, ϕ is the scalar flux, $\overline{P}'$ is the nominal linear heat rate (LHR), l is the heated length at axial level z, and P_{total} is the total power for the reactor. Numerous options are available to refine the spatial distribution of both power and burnup in the MPACT-CTF coupling [16–18], but these options are largely irrelevant to the work discussed in this report.

3. DECAY HEAT COUPLING IMPLEMENTATION

3.1 DECAY HEAT CALCULATION

The ORIGEN application programming interface (API) provides a straightforward function that can be called to calculate decay heat. The function simply takes in a list of isotopic number densities and returns the total heat generation in watts for that set of isotopes. All of the various decay mechanisms and heat generated by each decay are handled internally by ORIGEN. MPACT was modified to call this ORIGEN interface for every mesh region to obtain a 3D decay heat distribution. These values are then converted to decay heat LHRs to match the format of the prompt power production used for the coupling with CTF. The decay heat values are updated at the start of a calculation and after each call to ORIGEN so that the decay heat generation always matches the current isotopics in the MPACT model.

3.2 CTF POWER CALCULATION

The MPACT-CTF coupling was modified on the MPACT side to incorporate the decay heat distribution. The distinction between decay and prompt heat is made only on the MPACT side; ultimately, CTF only receives total heat from MPACT. First, the total power is decomposed into prompt and decay power:

$$P_{prompt} = P_{total} - P_{decay} \,, \tag{2a}$$

$$P_{decay} = \sum_{r \in R} \sum_{z \in Z} P'_{D,r,z} l_{r,z} \,, \tag{2b}$$

where P_{decay} and P_{prompt} are the decay heat and prompt portions of P_{total}, and $P'_{D,r,z}$ is the LHR due to decay heat at axial level z of rod r as calculated by ORIGEN. Using these definitions, we can devise a modified version of Eq. (1):

$$P'_{r,z} = P'_{P,r,z} + P'_{D,r,z} \,, \tag{3a}$$

$$P'_{P,r,z} = \max\left(\frac{\sum_{r' \in (r,z)} \sum_{g=1}^{G} \kappa \Sigma_{f,r',g} \phi_{r',g}}{\overline{P_P}'}, \; 0 \right) , \tag{3b}$$

$$\overline{P_P}' = \frac{P_{prompt}}{\sum_{r \in R} \sum_{zinZ} l_{z,r}} \,. \tag{3c}$$

These modifications allow the MPACT-CTF coupling to use the same subroutines in the CTF coupling API while separating out the prompt and delayed portions of the pin power distribution on the MPACT side of the coupling.

Note that the above approach does contain one limitation. The κ values in the MPACT represent all recoverable energy produced by a fission event for all time. This includes the fission Q value, the fission neutron energy, energy produced by capture of non-fissionable isotopes, prompt and delayed gammas, and all other delayed energy due to the decay of fission products. Because of this, κ essentially assumes that the current reactor power P_{total} has been the power level for a very long time previously. Thus, when rapid power maneuvers or transient calculations take place, it is possible for the decay power in a region to be larger than the calculated total power based on κ. This is why the $max\,(..., 0)$ function is necessary in Eq. (3b) to prevent negative power values.

3.3 INTEGRATION OF DECAY HEAT WITH MPACT TRANSIENT

3.3.1 ANSI/ANS-5.1 STANDARD

Because delayed power can be quite important for rapidly evolving transients, MPACT already included a treatment of delayed energy based on the ANSI/ANS-5.1 standard [8], which is briefly described in this section for the sake of comparison with the ORIGEN coupling later in this report. This standard decomposes decay heat production into 23 groups. Decay heat precursor isotopes are placed into one of these groups based on their decay constants, and a collective fission yield for the group is determined. The group yield and decay constant are formulated separately for four fissionable isotopes: ^{235}U, ^{238}U, ^{239}Pu, and ^{241}Pu. Fission of any other isotopes are considered to have come from ^{235}.

To use this standard in MPACT, the initial total power P_{total} is calculated at the onset of the transient using the contributions' four fissionable isotopes:

$$P^0_{total} = \sum_{i=1}^{4} \sum_{r\in R} \sum_{z\in Z} P_{i,r,z} \ , \tag{4a}$$

$$P_{i,r,z} = \sum_{g\in G} \kappa\Sigma_{f,i,r,z,g}\phi_{r,z,g} \ , \tag{4b}$$

where $P_{i,r,z}$ is fission power production from fissionable isotope i, the superscript indicates the time step index during the transient, and 0 refers to the steady-state condition. Now, at each time step, an explicit time discretization is applied to the decay heat precursor groups to determine the delayed power:

$$P^k_{delay} = \sum_{r\in R} \sum_{z\in Z} \sum_{i=1}^{4} \sum_{j=1}^{23} E^k_{r,z,i,j}\lambda_{i,j} \ , \tag{5a}$$

$$E^k_{r,z,i,j} = e^{-\lambda_{i,j}\Delta t^k} E^{k-1}_{r,z,i,j} + \left(1 - e^{-\lambda_{i,j}\Delta t^k}\right) \frac{\beta_{i,j} P_{i,r,z}}{\lambda_{i,j}} \ , \tag{5b}$$

where Δt is the time step size for time step k and $\lambda_{i,j}$ and $\beta_{i,j}$ are the decay constants and fission yields from fissionable isotope i for decay heat precursors in precursor group j. The first term on the right-hand side of Eq. (5b) is precursor concentration at time t that remains at time $t + \Delta t^k$, whereas the second term is the is the additional precursor concentration generated between t and $t + \Delta t^k$. The delayed power is then used to modify the total power:

$$\tilde{P}^k_{total} = P^k_{total} + P^k_{delay} - P^{k-1}_{total} \sum_{i=1}^{4} \sum_{j=1}^{23} \beta_{i,j} \ , \tag{6}$$

where the last term is the previous time step's delayed power, which is subtracted to prevent double-counting the delayed power. This updated power is then used for all remaining calculations in the time step.

One limitation in the current implementation in MPACT is that the starting power at the onset of transient conditions is assumed to have been the power for eternity past. This means that running an end of cycle (EOC) hot zero power (HZP) transient is not possible because there should be significant decay heat at EOC, but the code ignores that decay heat because the starting power is 0. This is due only to simplifications made for the MPACT implementation; the ANS-5.1 standard allows for an approximation of decay heat at zero power based on the total time of operation at full power and time since end of operation, but this feature is not implemented in MPACT.

3.3.2 ORIGEN DECAY HEAT COUPLING

With the ORIGEN decay heat coupling enabled, the previous section's approach is modified. Equation (5) is modified as follows:

$$P^k_{delay} = \sum_{r \in R} \sum_{z \in Z} P'_{D,r,z} l_{r,z} \; . \tag{7}$$

This formulation is simply a sum of the ORIGEN decay heat calculations over all MPACT mesh regions, obtained directly from the isotopic concentrations in each region. Equation (6) can then simplified:

$$\tilde{P}^k_{prompt} = P^k_{prompt} + P^k_{delay} \; . \tag{8}$$

It is important that the flux be normalized only to the prompt power in this case; because the decay heat precursor yields and decay constants are not easily attainable in the ORIGEN approach, there is no straightforward way to prevent double counting decay heat remaining from the previous time step if the total power is used. This resolves the limitation regarding EOC transient calculations discussed in the previous section, as the isotopics are used in the calculation of P^0_{delay} with no direct dependence on the current power level.

4. STEADY-STATE DECAY HEAT ANALYSIS

4.1 REDUCED VS. FULL DECAY CHAIN

By default, VERA utilizes a short depletion chain for ORIGEN to minimize computational resources (run time and memory). Although this library has been shown to produce accurate results for typical reactor physics calculations (e.g., eigenvalue, pin powers), it is not accurate for determining reactor decay heat immediately following SCRAM or shutdown. For these calculations, the full ORIGEN depletion chains using all isotopes must be used. Figure 1 shows the fraction of decay heat calculated using the short depletion chain for the nominal 2D fuel lattice depletion, relative to the decay heat calculated from the full depletion chain. Because the decay heat calculated from the short depletion chain is always less than 15% of the value from the full chain, the use of the short chain depletion library is clearly inappropriate for this application. The short chain is sufficient for neutron transport calculations, steady-state multiphysics calculations, and cycle depletion calculations because the decay heat closely follows the prompt heat during steady-state so negligible errors are introduced. This will be further shown in the remainder of the section.

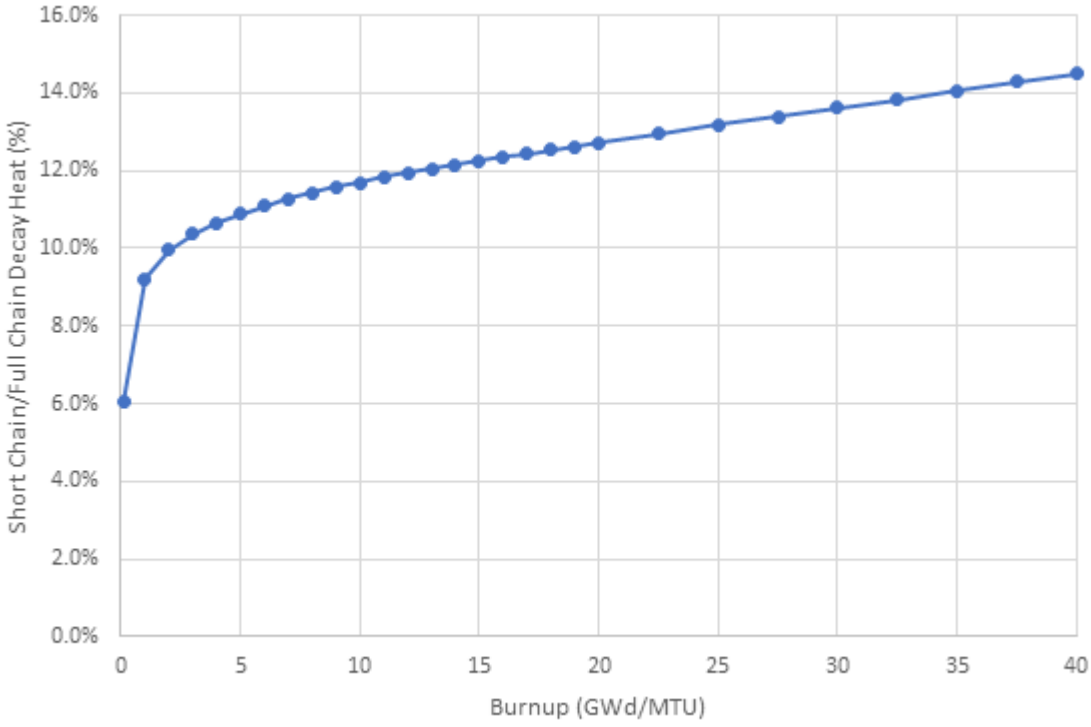

Figure 1. Decay heat using the short depletion chain as a fraction of that from the full chain.

Because typical reactor analysis involves depletion of multiple fuel batches in the reactor, it is desirable to switch between the short and full depletion chains depending on the analysis required—and potentially due to limited computing resources. For decay heat applications, a reactor model that is normally depleted with the default short depletion chain will need to be switched to the full chain sometime before the intended analysis, with sufficient intermediate burnup to establish the full set of needed isotopics. Using the simple 2D fuel lattice case previously discussed, the depletion to 40 GWd/MTU was performed with the short chain ORIGEN library, then an hourly depletion was performed at nominal operating conditions following

that using the full depletion chain library. The error in the total decay heat is shown in Figure 2. The error is initially greater than 90% (consistent with Figure 1) because the isotopes that are not in the short chain are initialized to 0 while reading the restart, but it quickly reduces to ~5% at 4 h and less than 1% at 27 h. In fact, 99.9% of the correct decay heat amount is obtained after 96 h at full power (4 effective full-power days (EFPD)). This supports the use of the short chain for eigenvalue and depletion calculations and using the full chain only for transient situations in which the time scales are short enough for the decay heat and prompt power shapes to be significantly different.

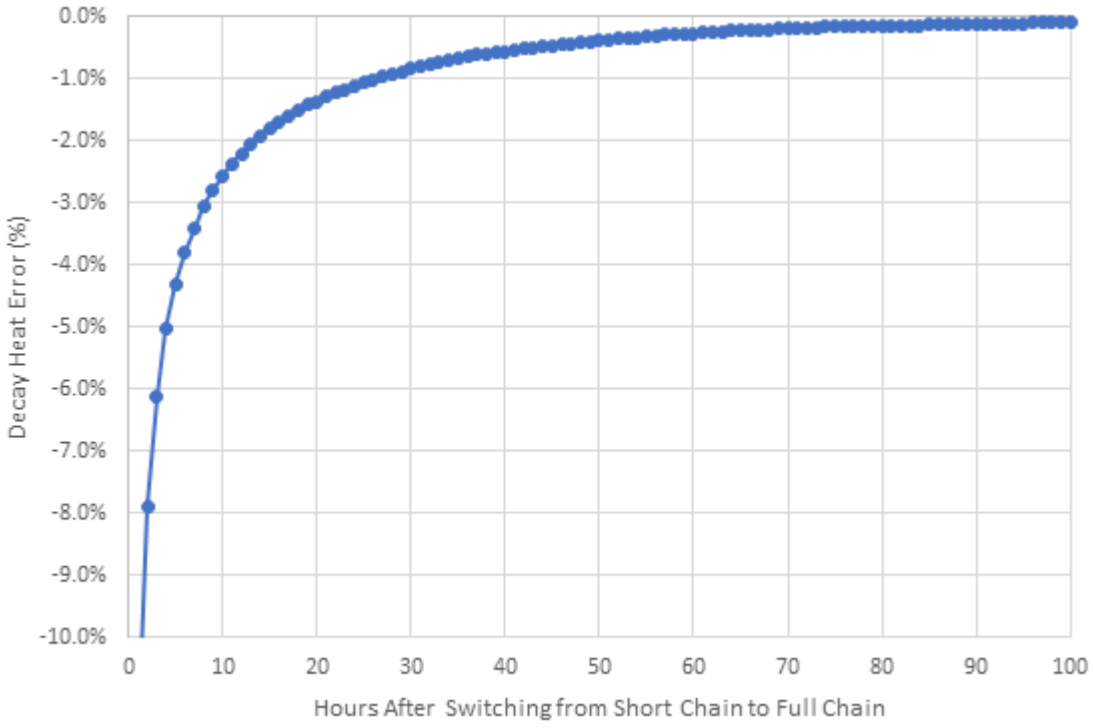

Figure 2. Decay Heat Error when Switching from Short Depletion Chain

4.2 RESTARTING VERA

Another potential source of error for VERA decay heat calculations is the use of a reduced number of depletion isotopes stored on the MPACT restart file. It is very common in LWR analysis to use restart files to store the isotopics of depleted fuel rods between fuel cycles. To save disk space and reduce input and output (I/O) time, the default set of isotopes stored on restart files is significantly reduced. Thus, even after depletion with the full ORIGEN chain, many of the isotopes needed for decay heat calculations are lost through the restart process. An MPACT option has been added to allow storage of all depletion isotopes (restart_isotope_set=depletion). Usage of this option, like usage of the short depletion chain, requires depletion with the full chain after restart to reestablish the equilibrium levels of radioactive isotopes producing the decay heat. The time required for this is shown in Figure 3, which resembles Figure 2 closely.

Based on these lattice calculations, whole-core analysis is subsequently performed using

- the full depletion chain as a reference solution,
- full depletion isotopes stored on the restart file, and
- decay times on the scale of hours.

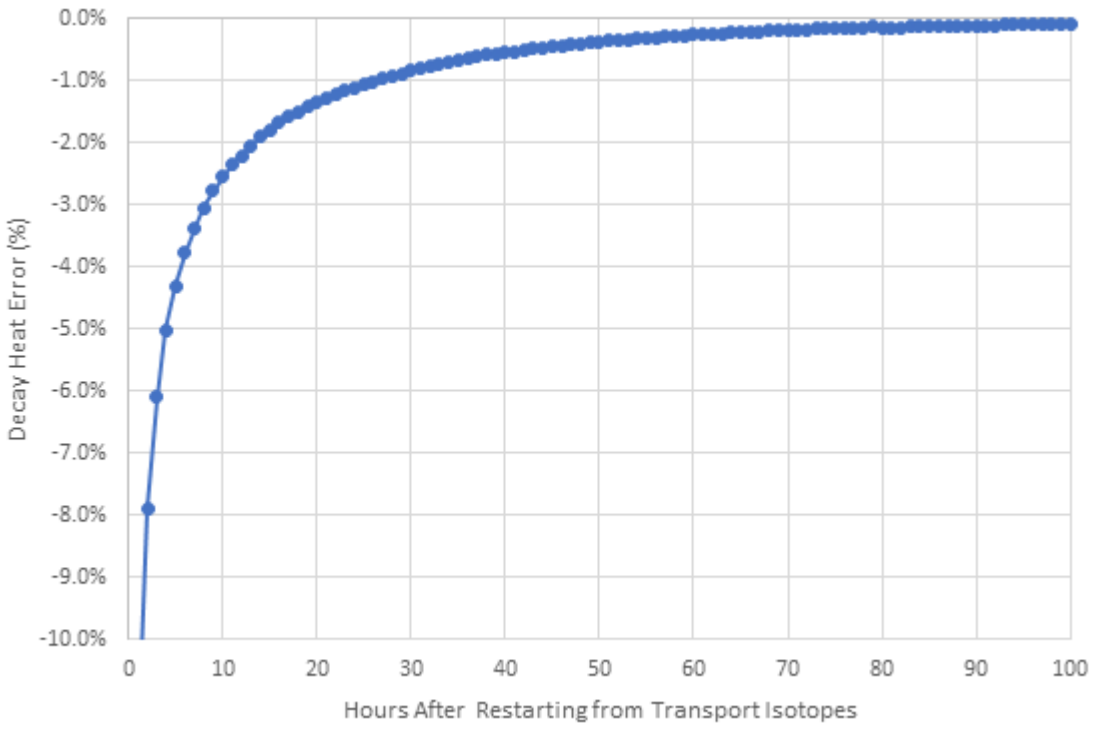

Figure 3. Decay Heat Error when Restarted from Transport Isotopes Only

Note that the decay_heat option was tested on the 2D lattice and was found to produce identical results since TH feedback was not used for the lattice calculations.

4.3 DECAY HEAT COUPLING WITH CTF

A publicly available quarter-core model of Watts Bar 1 Cycle 3 was used for whole-core decay heat testing. Similar to the previous lattice calculations, the following results were obtained with the quarter-core model:

1. For the steady-state nominal cycle depletion, decay heat magnitudes using the short depletion library were 11–14% of the values using the full ORIGEN depletion chain. Calculations with the full chain required approximately 15% more run time and were able to fit on the available 4 GB/core on Sawtooth.
2. The errors due to depletion with the short chain library and restarting from transport isotopes are completely healed by the first subsequent depletion step with the full depletion chain. This was tested at all fuel cycle burnup points, 5.1 to 505.3 EFPD, and the smallest step size in all cases was 5.4 EFPD.
3. The maximum difference in reactivity due to depletion with the short chain verses the full change was −1.8 parts per million (ppm).

Hourly tests that involved switching from the short depletion chain to the long depletion chain were performed at beginning of cycle (BOC), middle of cycle (MOC), and EOC for Watts Bar Unit 1 (WB1) cycle 3 (C3) (31, 241, and 474 EFPD). These tests were performed at realistic full-power operating conditions at each burnup, restarting from a short chain nominal depletion and performing an hourly nominal depletion with the full change. The error (compared to that observed always using the full chain) is shown in Figure 4 for each burnup. This result demonstrates that the errors observed using the short chain depletion library are not dependent on the cycle burnup.

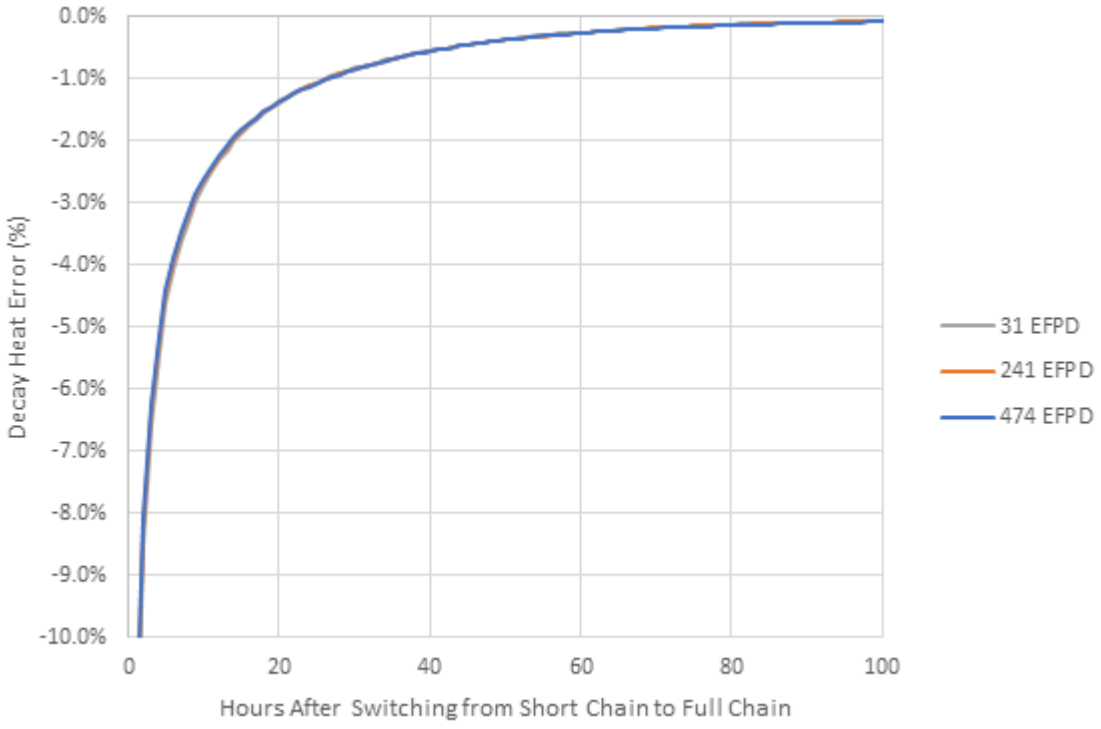

Figure 4. Whole-core decay heat error when switching from the short depletion chain.

The root mean square error (RMSE) and maximum absolute difference in pin-wise normalized decay heat distributions at 474 EFPD are shown in Figure 5. The differences are caused by restarting from the case

using the short chain depletion library and changing to the full chain depletion library. The instantaneous 6% RMSE difference decreases to 0.3% in one hour, and the maximum difference drops below 0.5% after four hours.

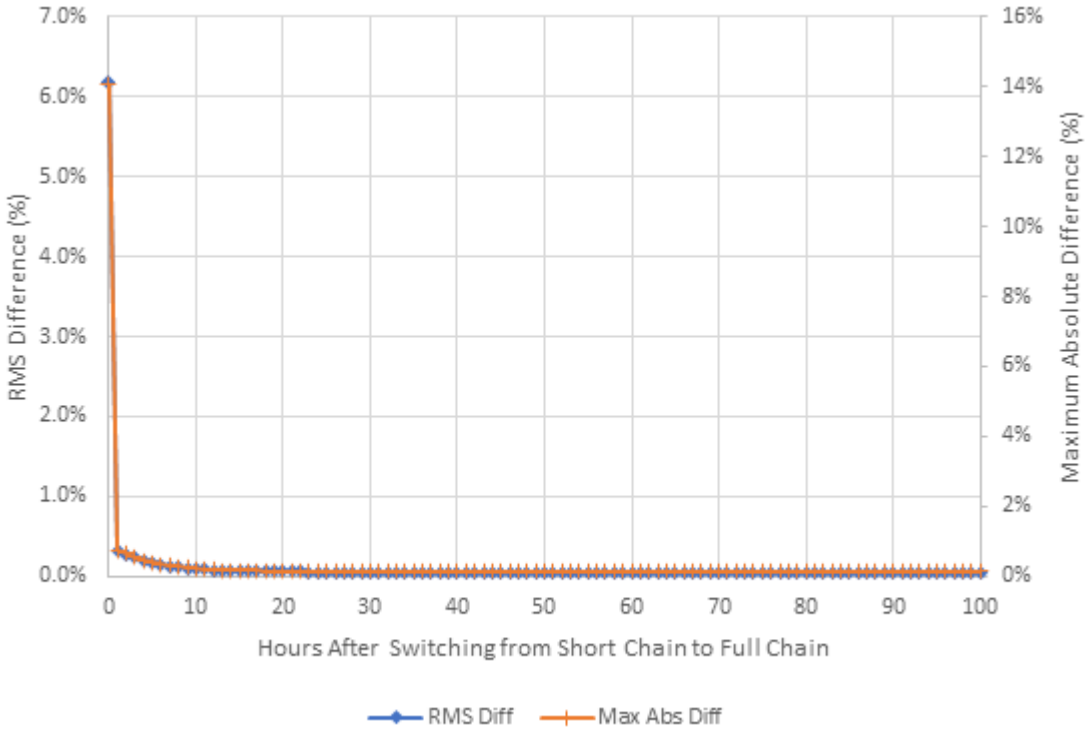

Figure 5. Whole-core decay heat error when switching from the short depletion chain

The whole-core depletion for WB1 C3 was executed with the full ORIGEN depletion chain and the new decay_heat option for pin power and TH feedback. Compared with the same case without this option, the differences were very small:

- maximum reactivity difference of 0.4 ppm
- maximum pin power difference of −0.4% and maximum pin power RMSE difference of 0.2%
- identical decay heat magnitudes
- maximum pin decay heat distribution difference of −0.6% and maximum pin decay heat distribution RMSE difference of 0.2%
- the largest differences occur at state points where power changes occur (i.e., power ramp or coast-downs).

The above results indicate that the implementation of decay_heat has not affected the steady-state pin power distributions significantly, as expected. Only minor differences in heat distribution affecting the TH are observable. Some of these differences may even be within default convergence criteria.

4.4 DECAY HEAT POST-SHUTDOWN

A WB1 C3 EOC test (474 EFPD) was performed by completely shutting the reactor down (0% power input) and observing the effects of the new code modifications. The decay heat calculated in pin_decay_heat was reasonable and unchanged by the decay_heat option. Figure 6 shows the total shutdown decay heat vs. time using hourly time steps, and Figure 7 shows the decay heat distribution at the point of shutdown. There

is a large drop in the power between the first 2 time steps of Figure 6 because the short-lived decay heat precursors decay rapidly; this time step size would need to be greatly refined for transient calculations on the order of seconds or minutes, but it is sufficient for this shutdown demonstration.

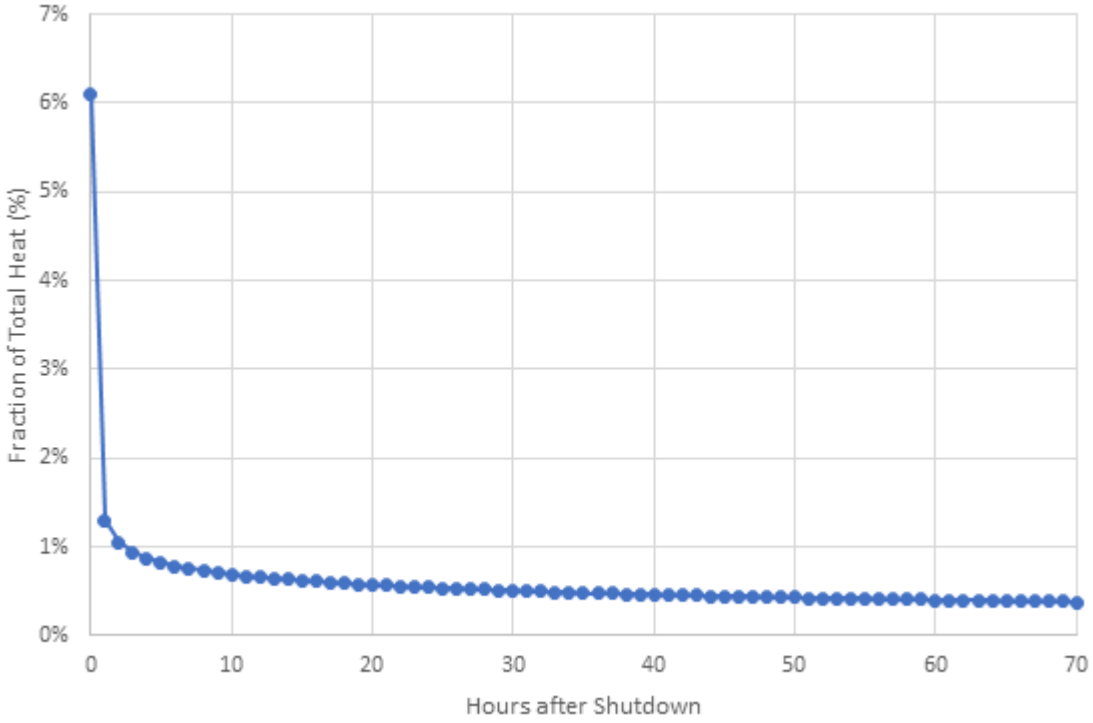

Figure 6. Whole-core decay heat following shutdown.

When TH feedback and decay_heat are enabled, the coolant and fuel temperature distributions at the beginning of the shutdown were also verified. The coolant temperature distribution from CTF (Figure 8) appears as expected, with a variation of only about 3°C. Likewise, the fuel temperature distribution follows the decay heat distribution and has a range of 36°C (Figure 9). Both of the coolant and fuel temperature distributions decay quickly as the decay heat decreases after shutdown.

The new decay_heat option that produced the TH results shown above also resulted in a negative effect on reactivity. The increased coolant and fuel temperature reduced reactivity by up to 19 ppm, which also decreased quickly with the decay heat, shown in Figure 10.

One issue found with the shutdown testing is that the output pin_powers do not reflect the distribution of decay heat. Figure 11 shows the output pin powers, which are very top peaked rather than in the shape of the decay heat shown in Figure 7. The pin_powers data are normalized according to prompt power and neglect the effects of decay heat. Unfortunately, this can be resolved only through the generation of a new cross section library. Since such an effort is outside the scope of this milestone, the frequency and magnitude of the pin_powers error will be quantified in terms of its impact on analyses that are expected to make use of the decay heat coupling.

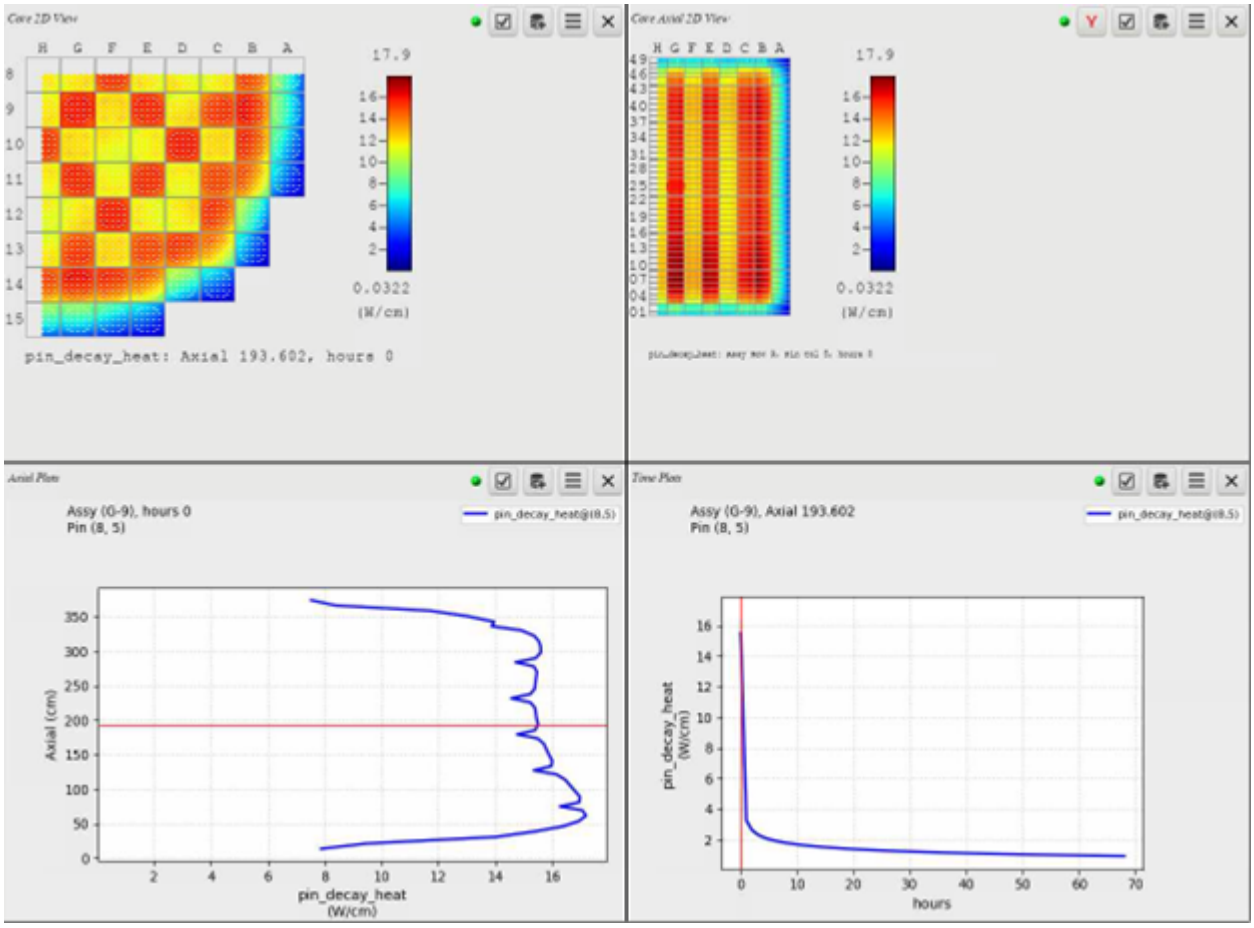

Figure 7. Whole-core decay heat distribution after shutdown (474 EFPD).

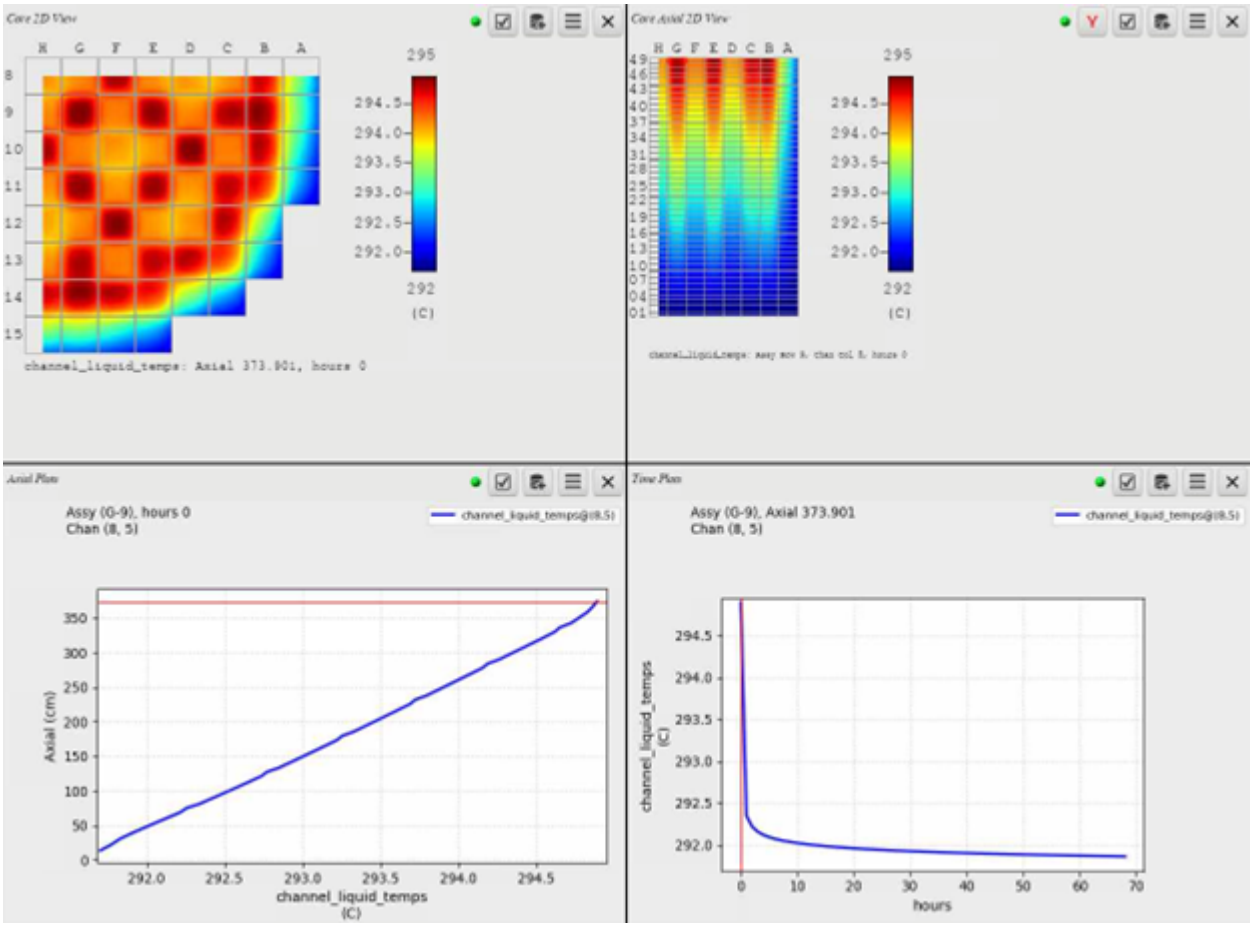

Figure 8. Coolant temperature distribution after shutdown (474 EFPD).

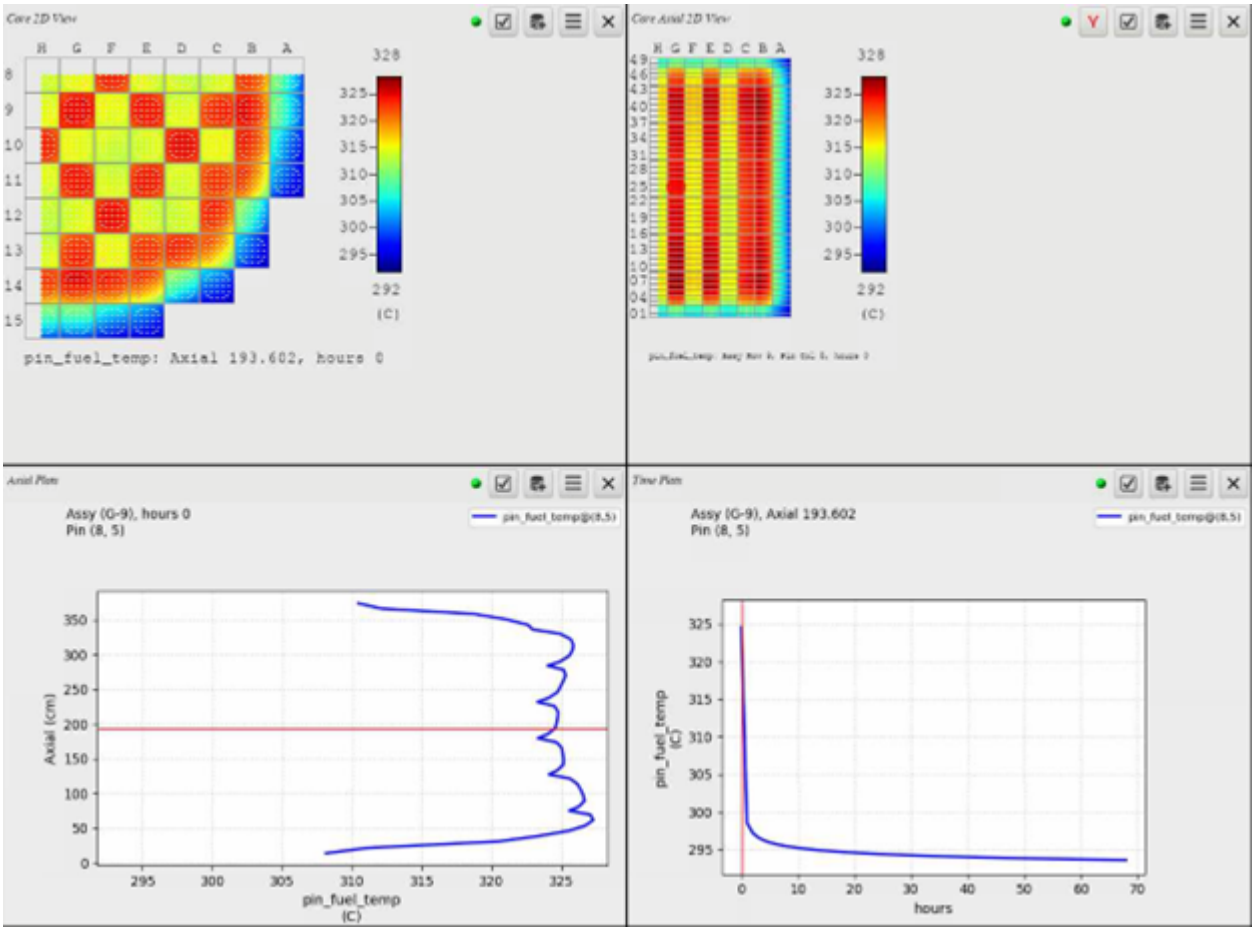

Figure 9. Fuel temperature distribution after shutdown. (474 EFPD)

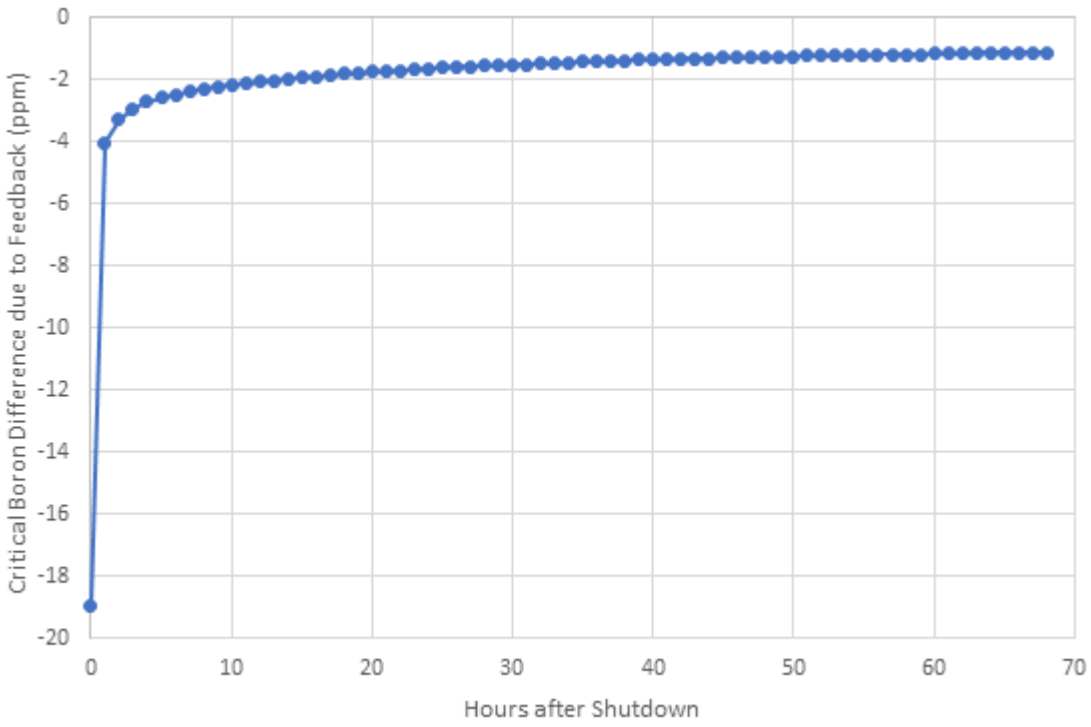

Figure 10. Reactivity impact from decay heat after shutdown (474 EFPD).

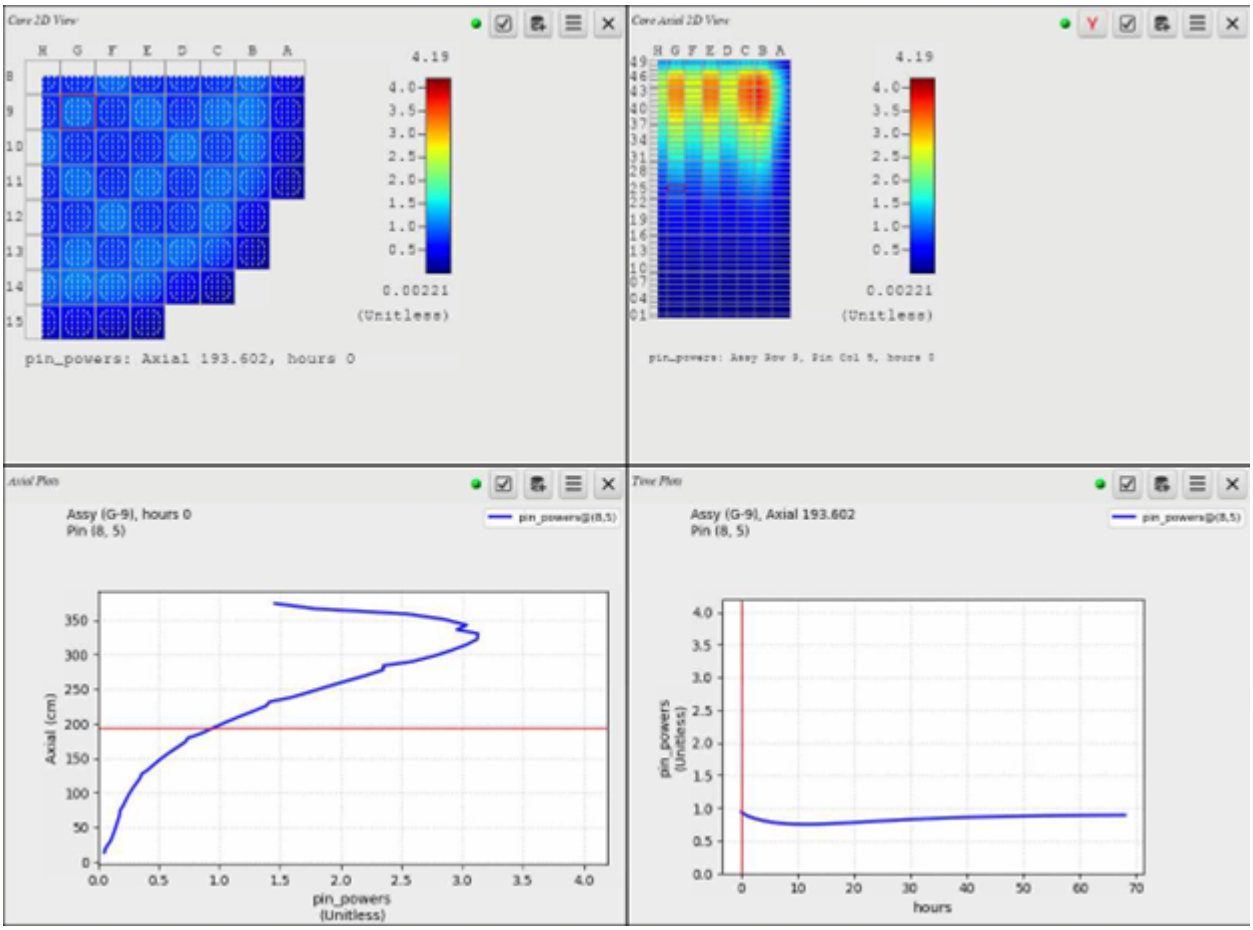

Figure 11. Incorrect pin power distribution after shutdown (474 EFPD).

4.5 POWER MANEUVER

An EOC power maneuver down to 10%, power was simulated for WB1 C3. The power and decay heat profiles with time are shown in Figure 12. The maneuver was performed by an instantaneous drop to 10% power (for which ~6% should be from decay heat initially), and then an instantaneous return to full power after 12 h. The magnitude of "prompt" heat is calculated by subtracting the calculated decay heat magnitude from the total input power. The figure shows the fraction of the total power being produced that comes from decay and prompt sources. At the initial power drop, the decay heat is 61% of the total power. When the power is ramped back up, the prompt power fraction must increase dramatically above the normal fraction because it takes a few hours for the decay heat to catch up to the equilibrium level.

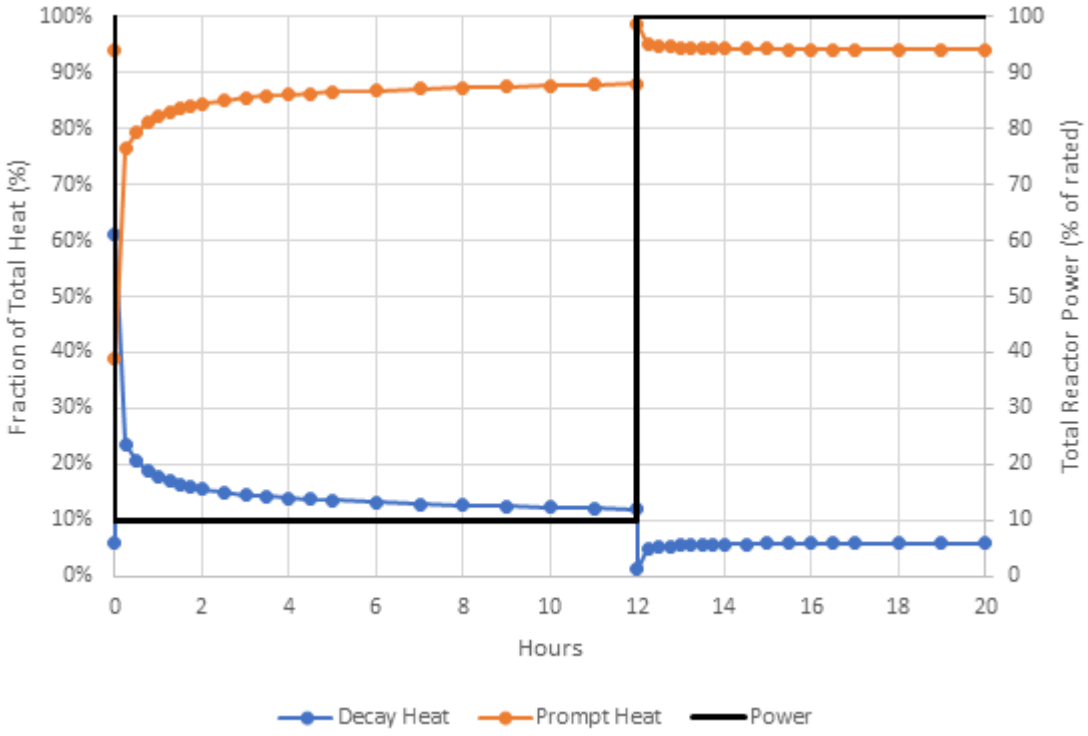

Figure 12. Components of total heat during a power maneuver to 10%.

The new decay_heat feature was tested for this power maneuver, which should demonstrate an impact on the core TH and power distribution resulting from the decay heat, which is a significant portion of the core total power at the initial power reduction to 10%. Figure 13 shows the power differences at the core level over time, in terms of total RMSE difference and maximum absolute pin power difference. At the initial power reduction, the pin power distribution showed a percentage difference of 4% RMSE and nearly 12% for the maximum. The distribution of these differences is also shown in Figure 20.

Furthermore, because of the effects of the decay heat on the TH, there appears to be a downstream impact on future decay heat distributions. This is shown in Figure 14. Initially, the decay heat is the same with and without the new model, but later varies by up to 0.5% RMSE and 1.5% maximum difference. This shows that the secondary effects of decay heat on the core TH is not a significant contributor to the decay heat distribution itself.

Figures 15 and 16 show the prompt and decay heat distributions before and after the reduction to 10% power, respectively. In Figure 15, the prompt and decay heat distributions are very similar. Note that the prompt

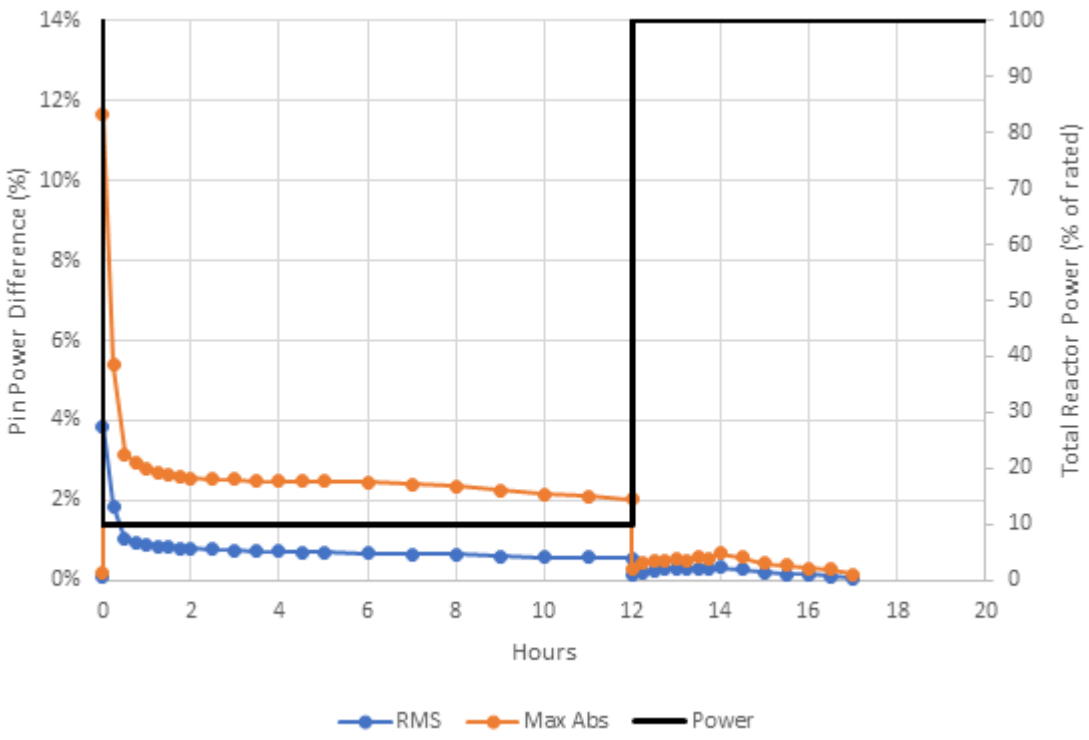

Figure 13. Pin power distribution differences from new decay heat model.

distribution is calculated by subtracting the decay heat distribution from the total power (pin_powers times nominal LHR). In Figure 16, after the power reduction and control rod insertion, the prompt heat distribution shifts up and away from the control rods, whereas the decay heat distribution is unchanged. However, a problem is apparent: the prompt data have negative values, which means there are locations where the decay heat is greater than the total power; this is simply a different manifestation of the shortcoming in pin_powers where the distribution is normalized only to the prompt power shape instead of accounting for decay heat.

Likewise, Figures 17 and 18 show the prompt and decay heat distributions just prior and subsequent, respectively, to returning to full power at 12 h into the maneuver. By 12 h, the decay heat distribution shifts somewhat higher in the core and begins to show impacts from the control rods. This is, then, much different from the power distribution produced once power ascension is completed.

Figure 19 plots the local decay and prompt heat (W/cm) in a fuel rod near the control rods for the power maneuver. The relative differences in the different heat sources and how they contribute with time are as expected.

Figures 20–23 show the differences in relative simulation results obtained using the new decay_heat option. Because it allows for the decay heat distribution to impact the core TH, many of the typical core operating parameters are affected in different degrees. Pin power differences are seen up to 12%, decay heat differences up to 1.5%, and coolant temperature differences up to 0.8°C.

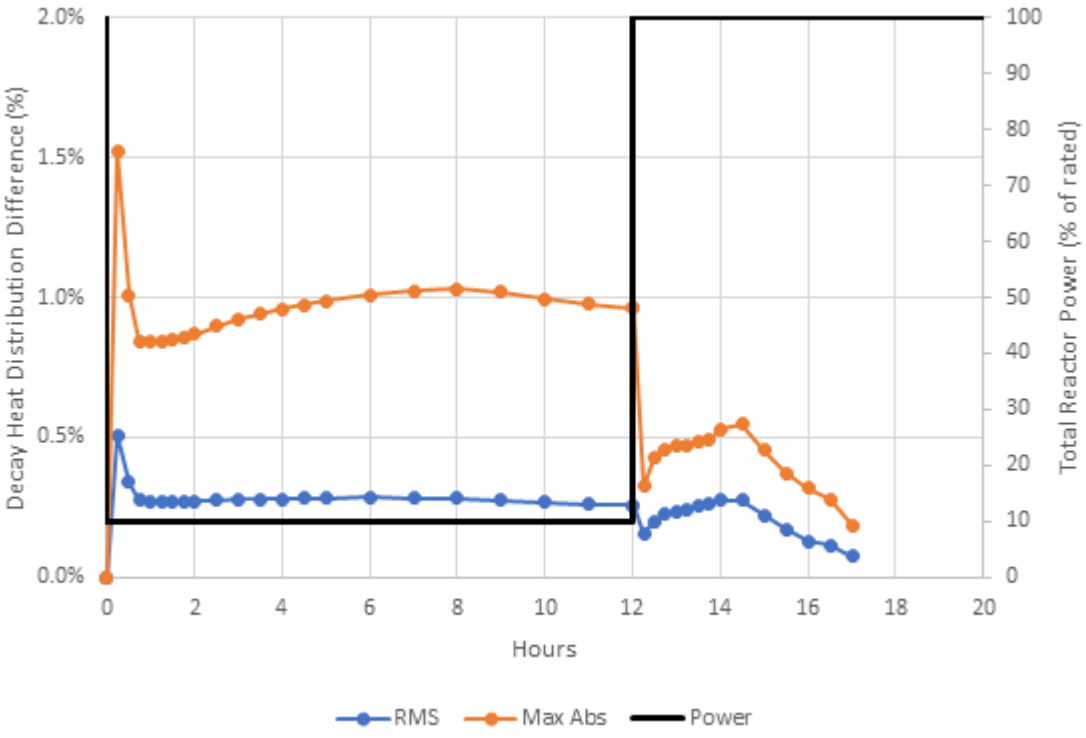

Figure 14. Decay heat distribution differences from new decay heat model.

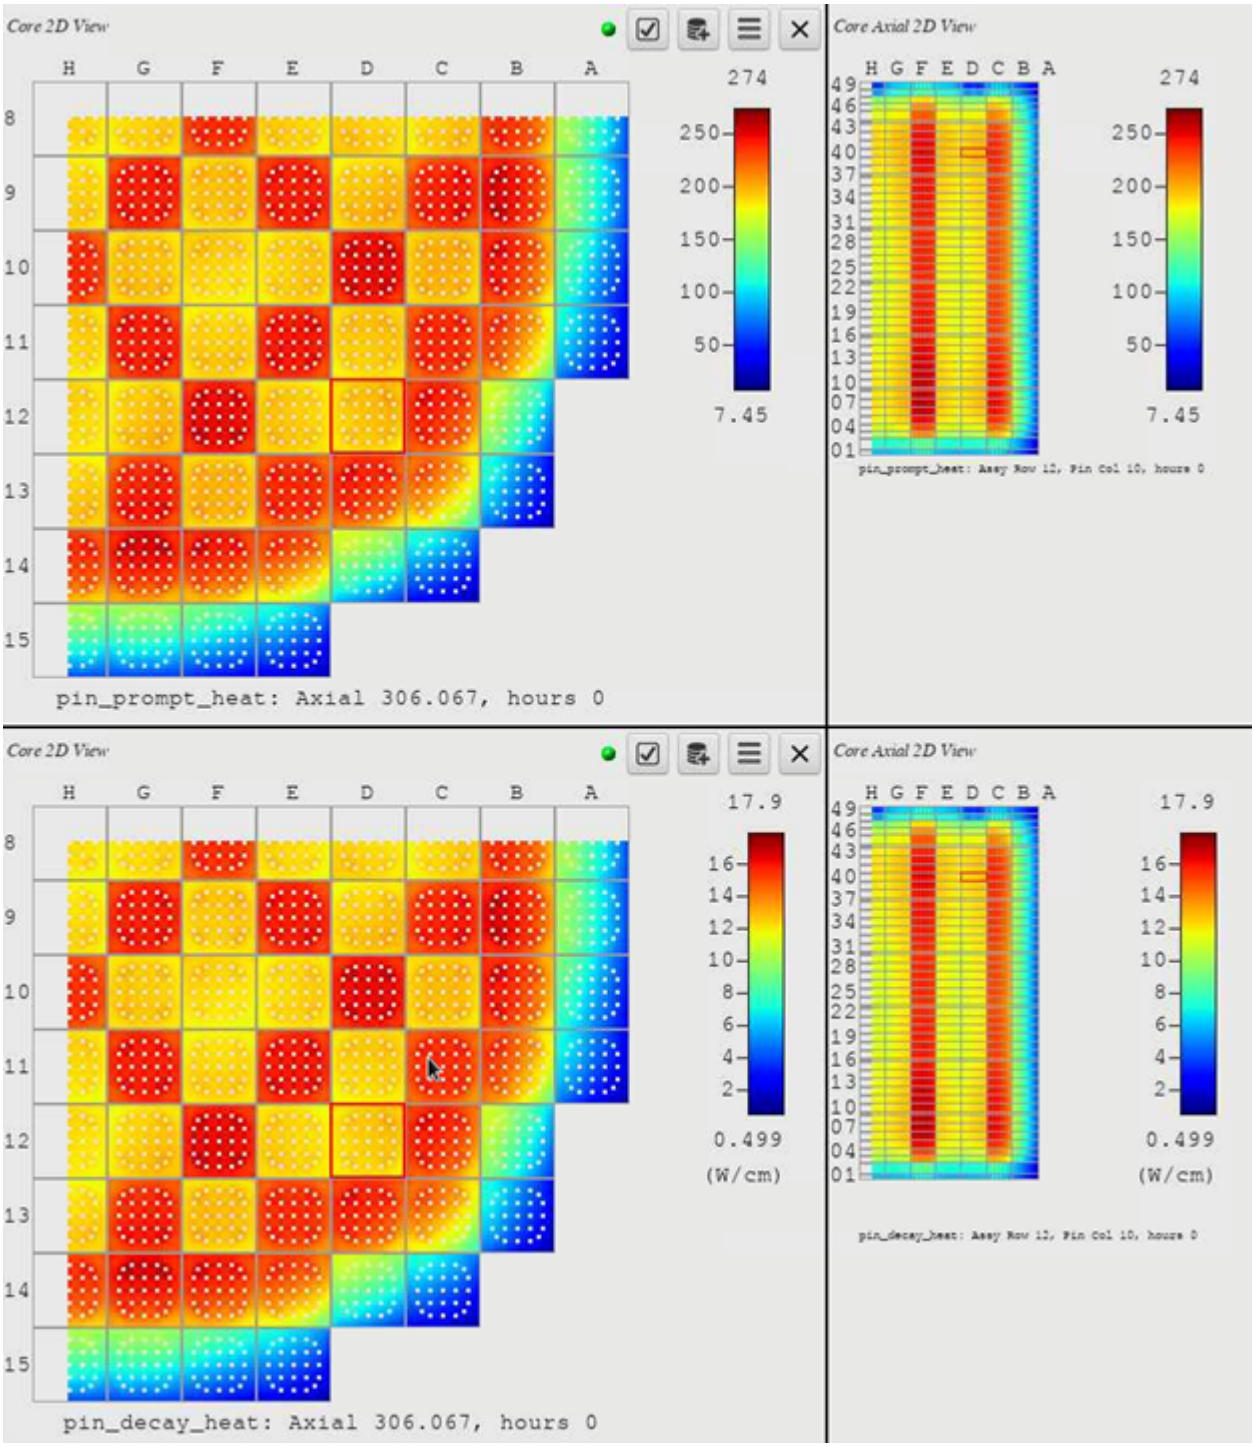

Figure 15. Prompt (top) and decay (bottom) heat distributions (W/cm) prior to the power reduction.

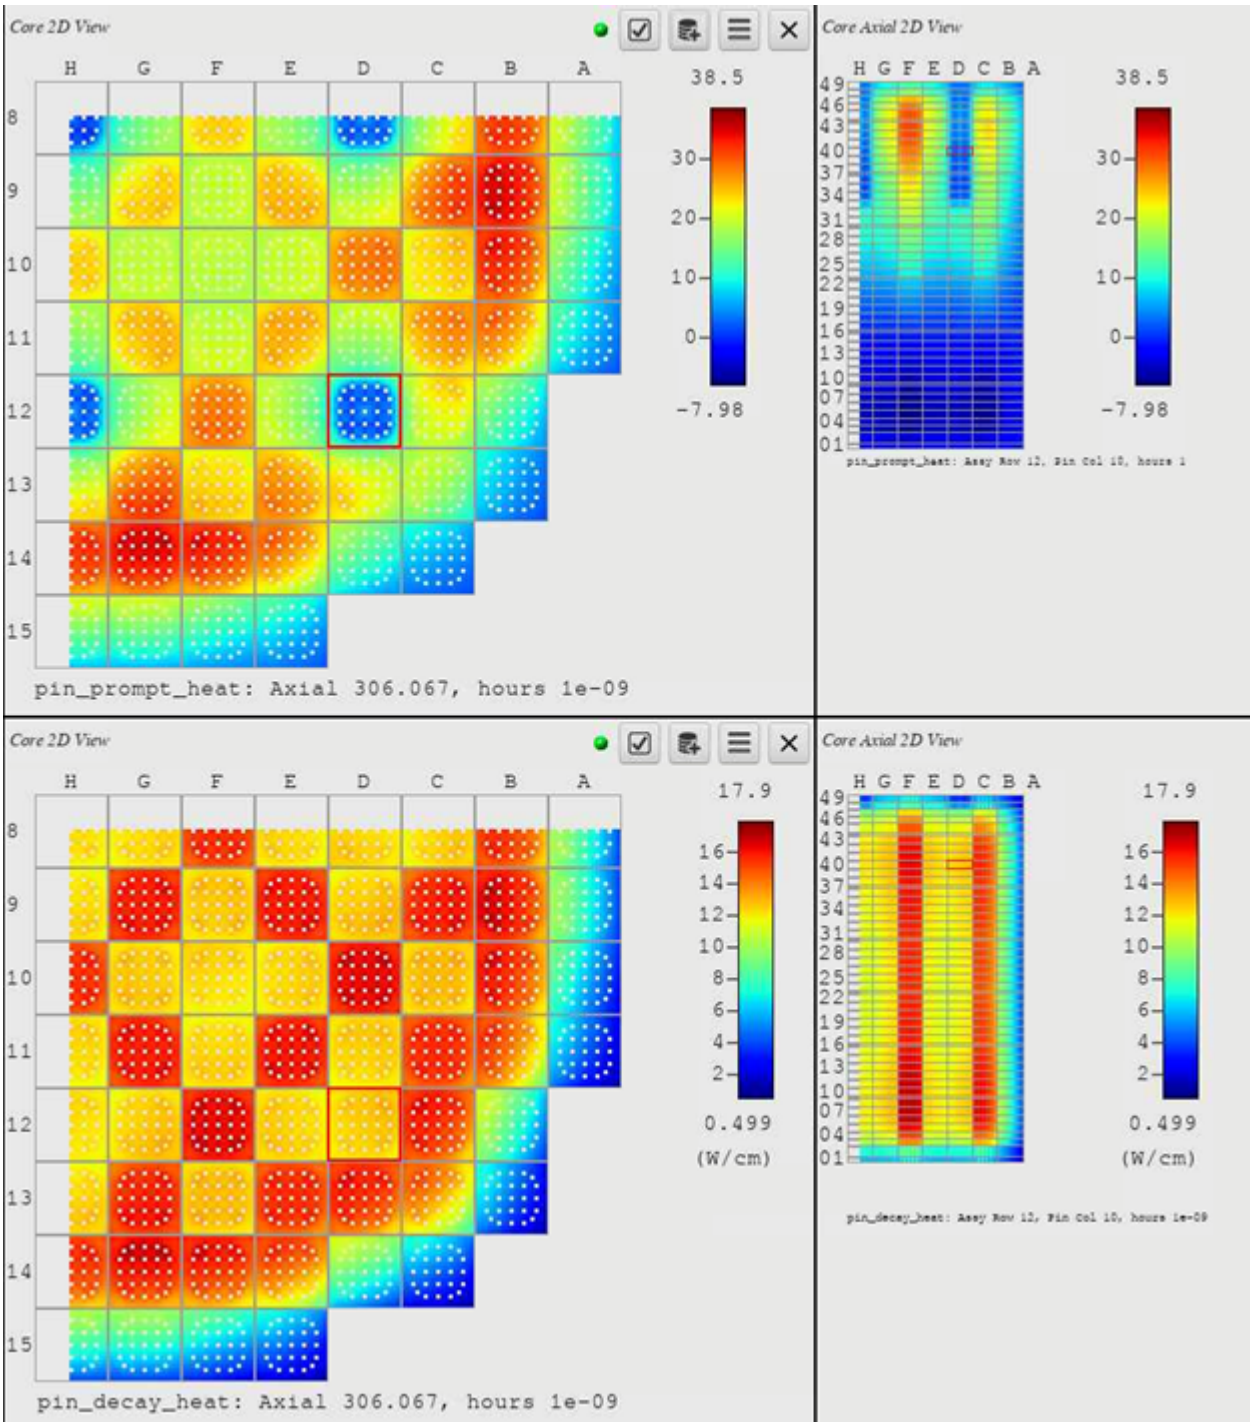

Figure 16. Prompt (top) and decay (bottom) heat distributions (W/cm) subsequent to the power reduction.

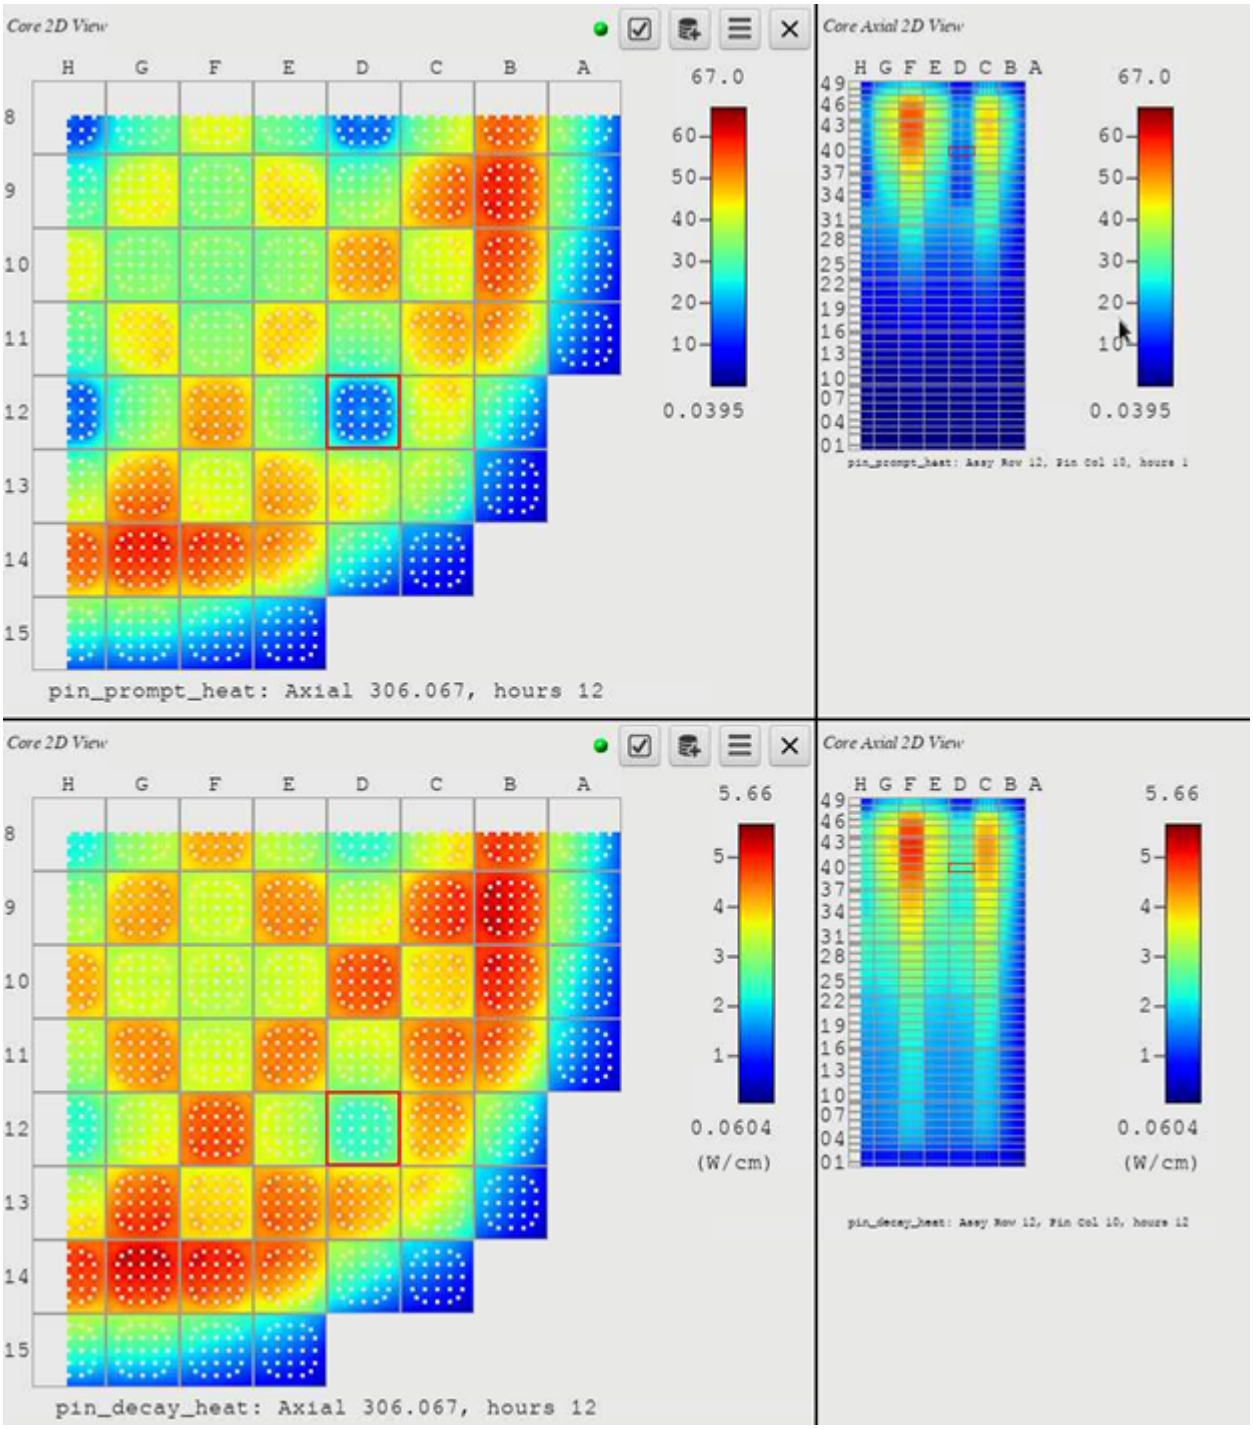

Figure 17. Prompt (top) and decay (bottom) heat distributions (W/cm) prior to returning to full power.

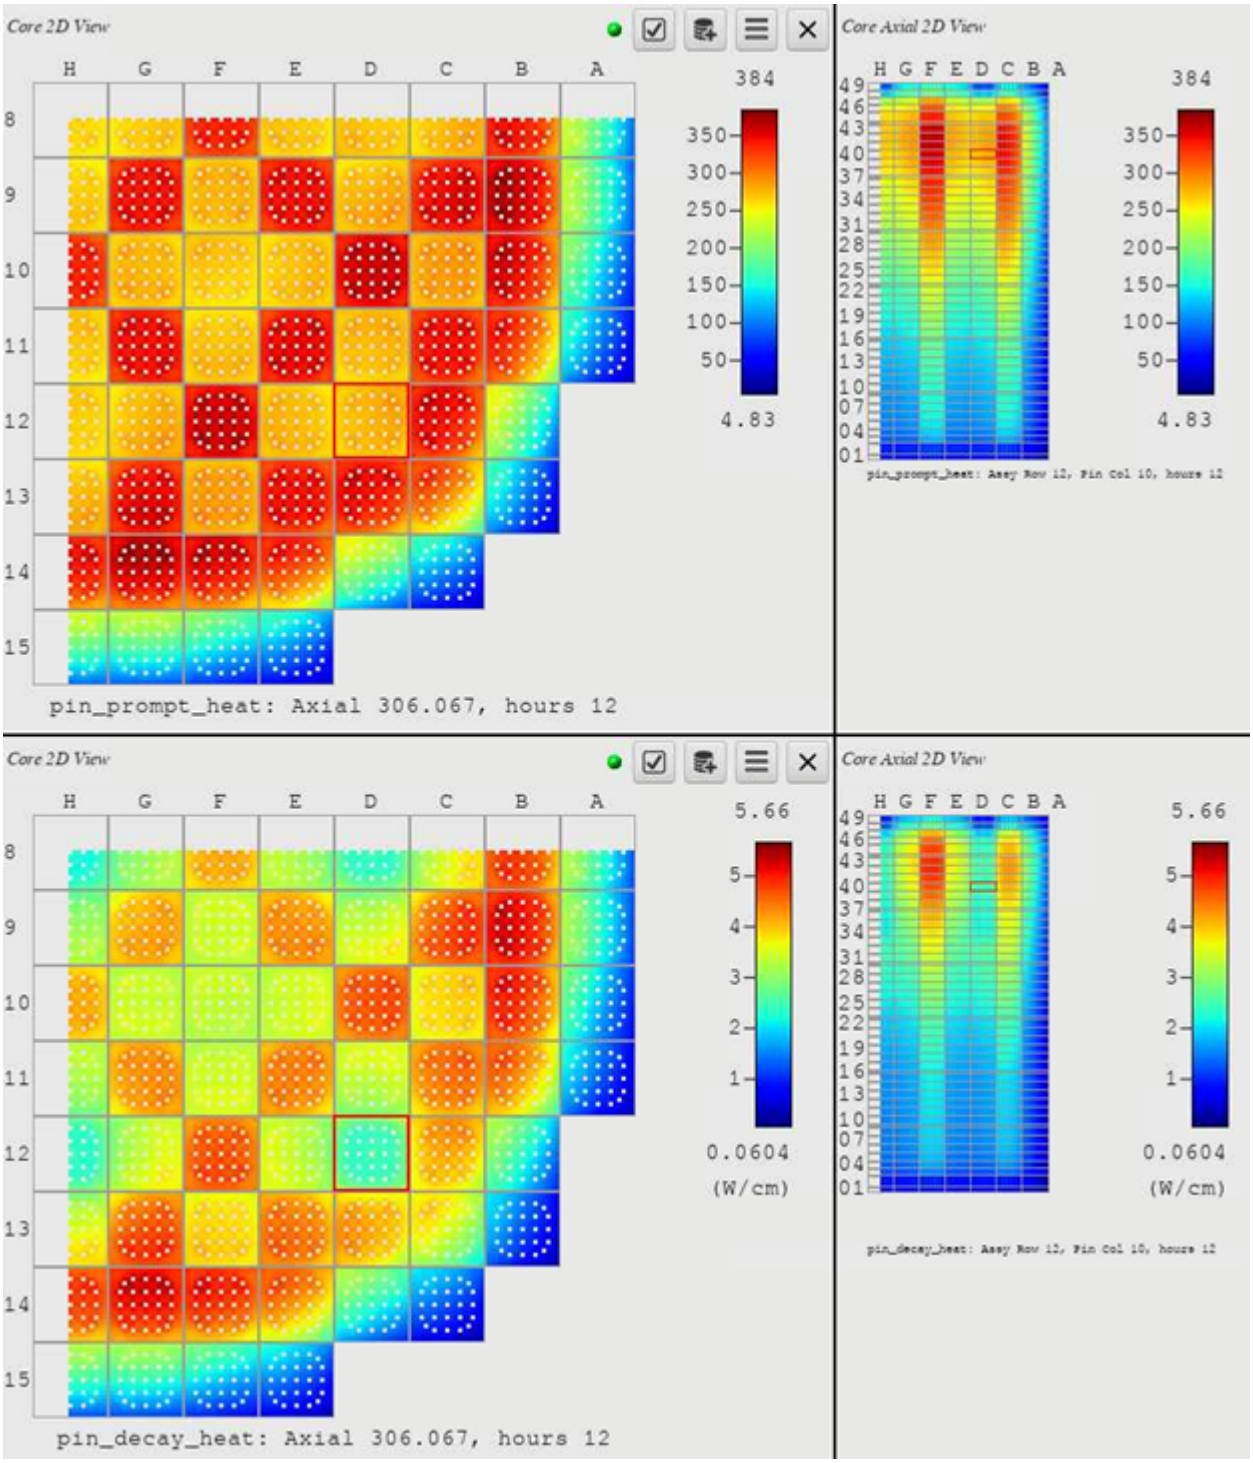

Figure 18. Prompt (top) and decay (bottom) heat distributions (W/cm) subsequent to returning to full power.

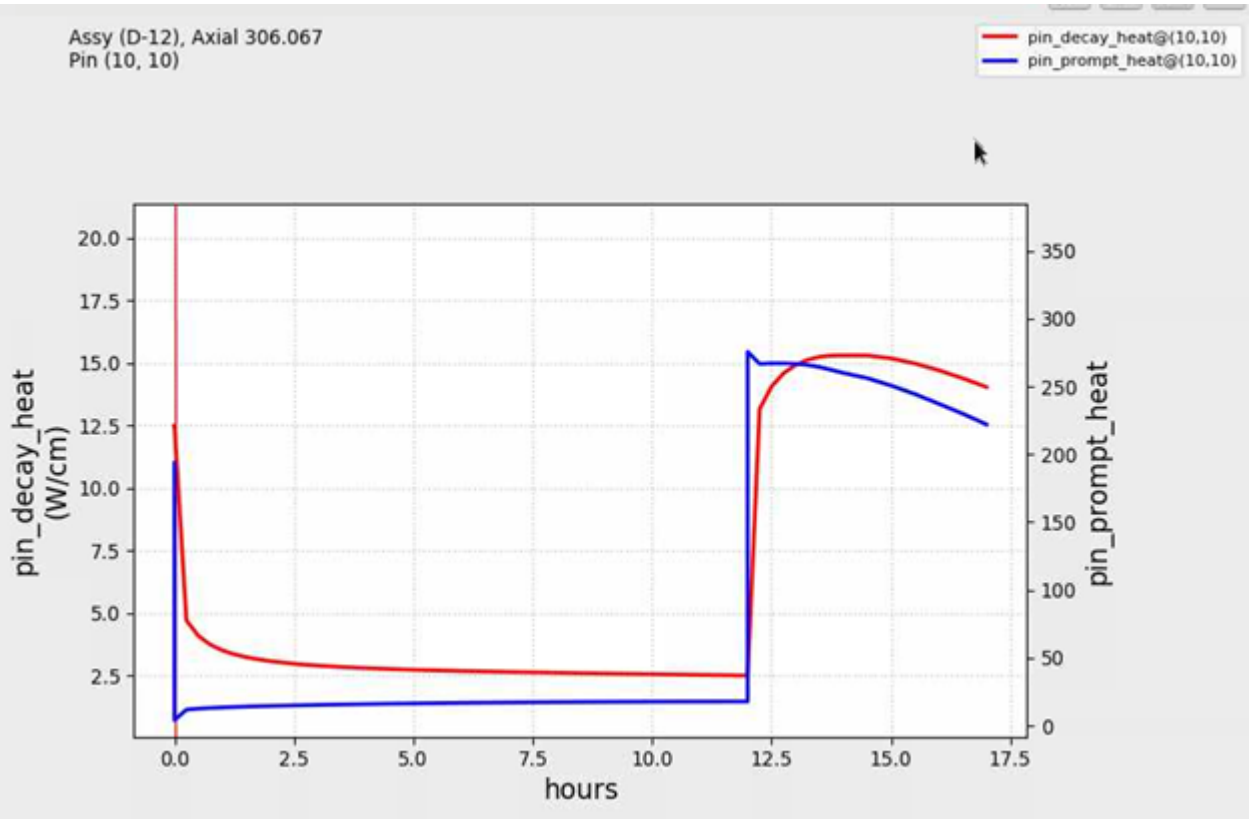

Figure 19. Local prompt and decay heat in location near control rod.

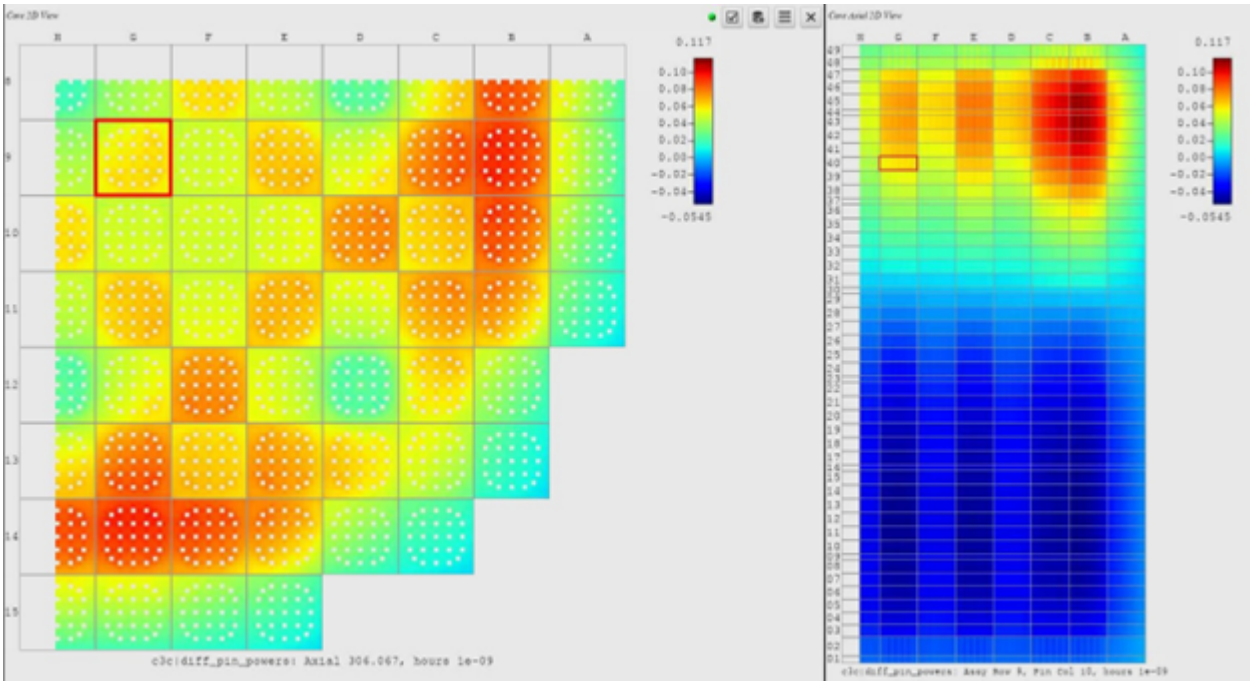

Figure 20. Pin power distribution differences after power reduction using the new decay heat method

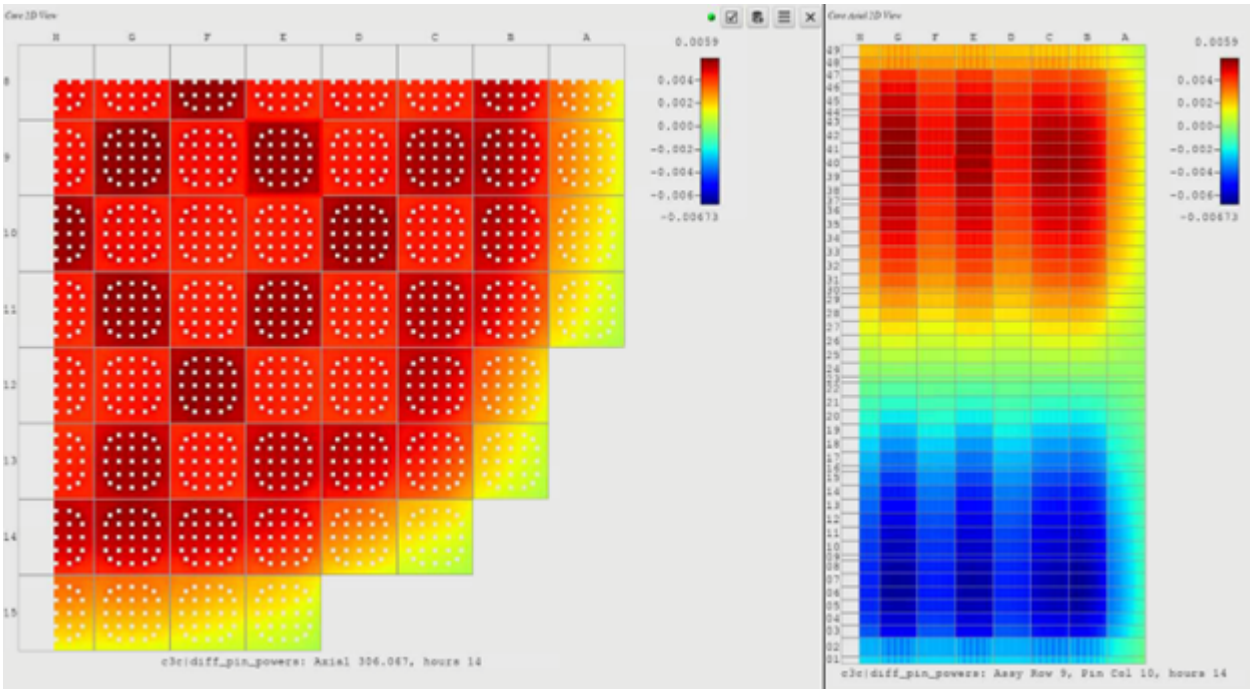

Figure 21. Pin power distribution differences after power ascension using new the decay heat method.

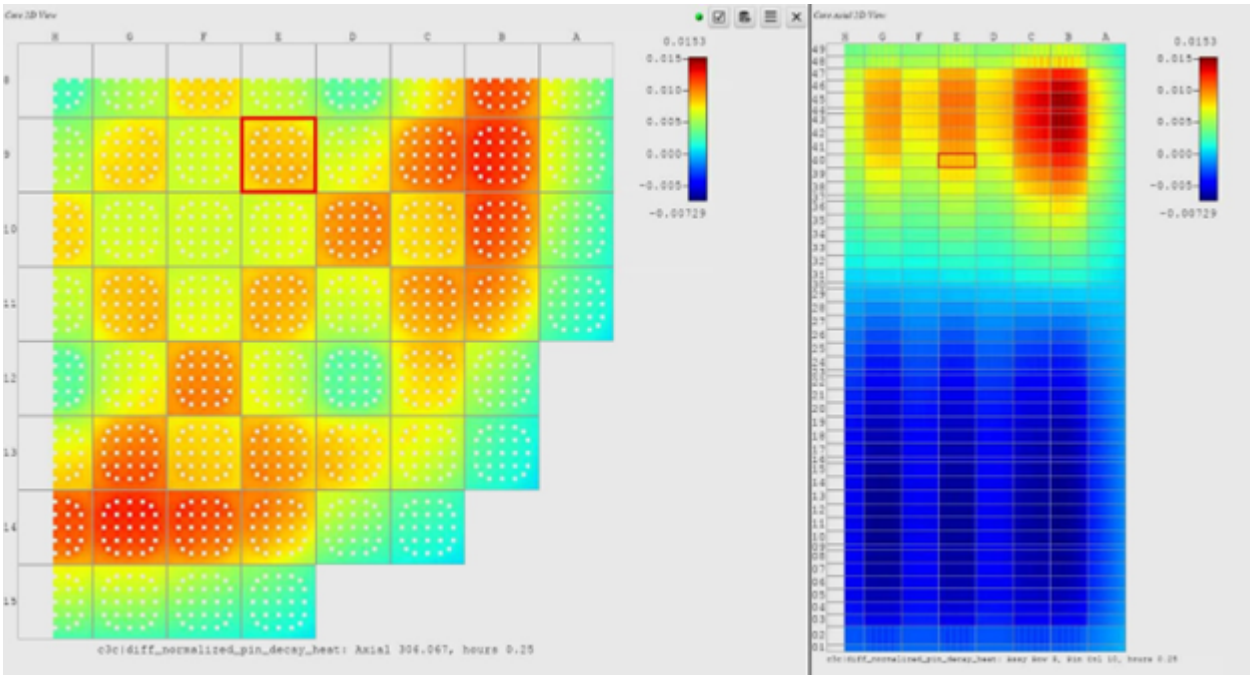

Figure 22. Normalized decay heat distribution differences after power reduction using new decay heat method.

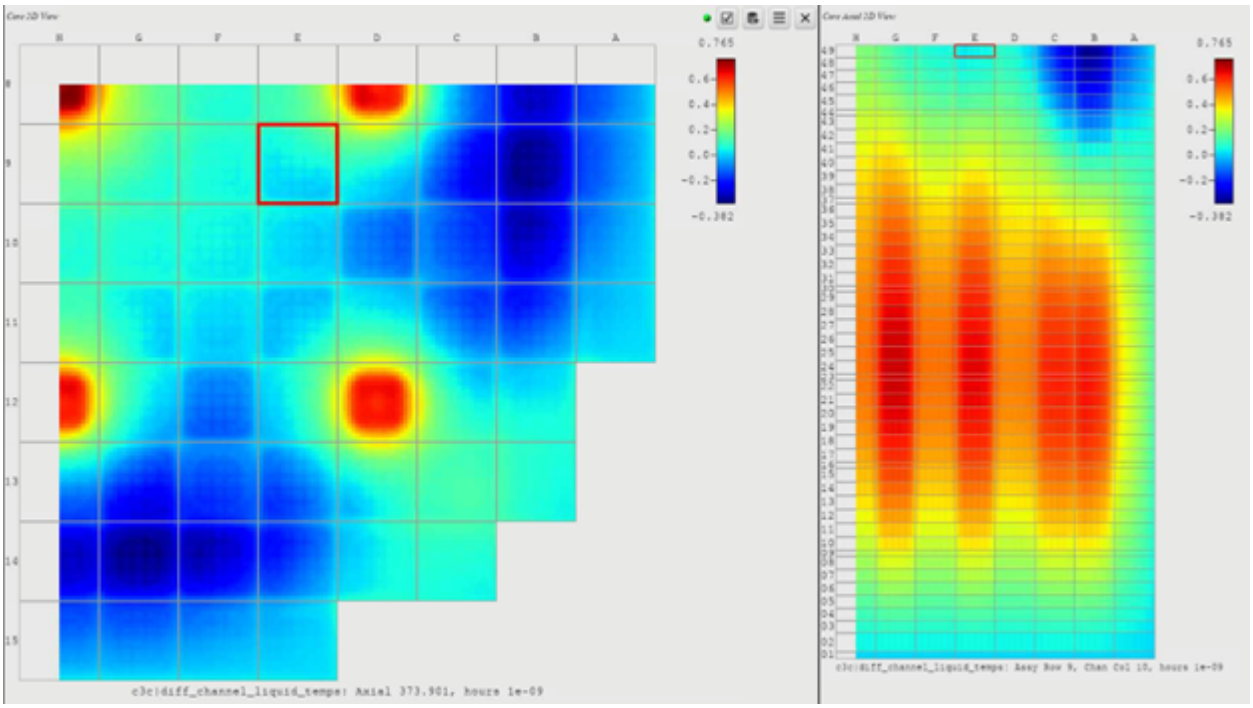

Figure 23. Normalized decay heat distribution differences after power ascension using new decay heat method.

5. TRANSIENT DECAY HEAT ANALYSIS

5.1 TEST PROBLEM

To test the new capability, a small problem called 4-mini was used. This is based on the problem 4 of the VERA progression problems [19]. To reduce computational expense, the height of problem 4 is reduced, 8 energy groups are used instead of 24, and the lattice is reduced from 17×17 to 7×7. Simplified control rods are in the center assembly as in the original problem; however, the corner assemblies are modified to also have assemblies. The model is depleted to a burnup of 20 MWd/kgHM to initiate the transient with EOC-like isotopic distributions. For the test, the center rod is partially ejected from 0 steps to 4 steps (out of 230 total) in 0.025 s to induce a positive reactivity and power rise; a SCRAM then occurs at 0.05 s. It is expected that the ANSI and ORIGEN approaches to calculating the decay heat ought to be fairly similar; using this small problem allows an initial check of that assumption before moving to larger, more computationally expensive problems.

Three variations of this problem were run: (1) one with no delayed power, (2) an ANSI-based delayed power calculation, and (3) an ORIGEN-based delayed power calculation. The results for the three calculations are shown in Figure 24. It is evident from this plot that the ANSI and ORIGEN approaches are quite similar in their prediction of total core power. The ORIGEN approach shows slightly higher power during the SCRAM by about 0.001% of the nominal core power; for the rest of the calculation, it predicts lower power than the ANSI approach. The largest difference occurs during the initial power rise, ranging from −0.005% to −0.003%. In the portion of the calculation that is driven by the decay heat, the differences start out smaller than −0.003% and trend toward 0% fairly quickly. Overall, these differences are very minor and would have very little impact on the TH response for a long-running transient. Additionally, it appears possible that the ANSI calculations might be more conservative than those of ORIGEN. This must be confirmed through more extensive full-scale simulations, but it would be good news for industry if the high-fidelity calculations show that they do not need to be any more conservative than they would be based on the ANSI standard.

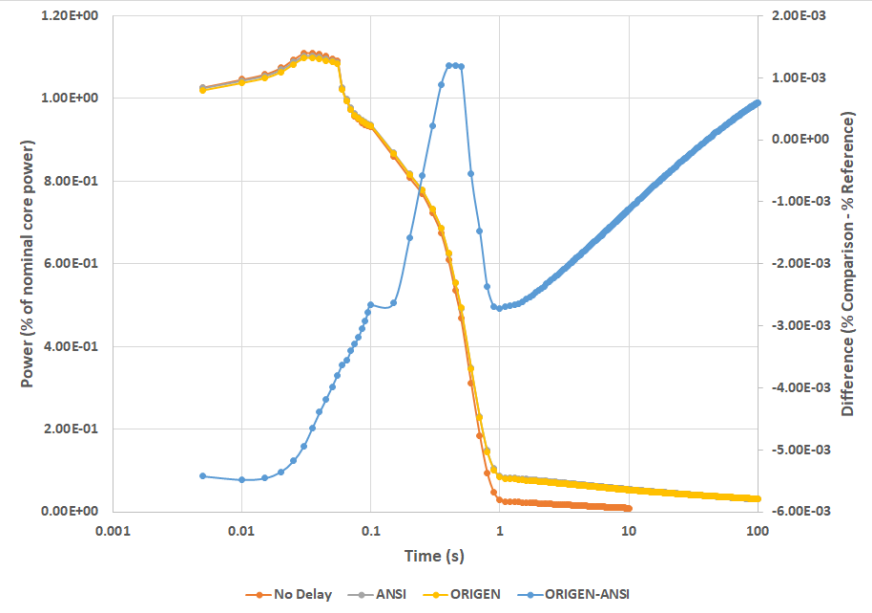

Figure 24. Comparison of various delayed power calculation approaches for transient 4-mini case.

Figure 25 shows the ORIGEN case compared to the case with no delay power, including the separate prompt and delayed components of the power. The delayed portion of the power is nearly constant, which is anticipated over the short duration of this transient. The prompt portion of the power matches the total power in shape before the SCRAM. After the SCRAM, it matches the no delay power. Thus, the decay heat coupling has very little impact on the total prompt power but adds a delayed component. By about 10 s into the transient, the total ORIGEN power is made up almost entirely of decay heat.

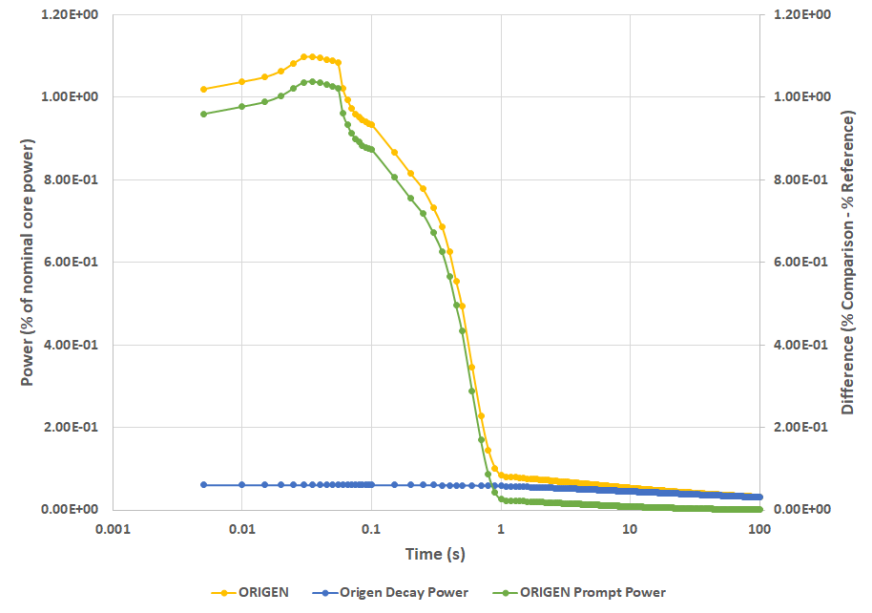

Figure 25. Power components using ORIGEN-based delayed power calculation for transient 4-mini case.

The power distribution can also be examined. Figures 26–31 show the heat distributions and differences at the following five times:

1. t = 0.0 s, transient initiation
2. t = 0.03 s, peak core power
3. t = 0.055 s, SCRAM initiation
4. t = 1.0 s, end of rapid power descent due to SCRAM
5. t = 10.0 s, power is primarily decay heat and SCRAM is complete
6. t = 100.0 s, end of simulation and no remaining prompt power

Each figure shows the slices of the total heat calculated by the ORIGEN-coupled simulation first, then the ORIGEN result minus the ANSI result. Both data sets are shown in units of W/cm. At each time step, the slice for both the total heat and the difference was taken at the location where the difference had the largest magnitude.

At transient initiation, the differences in the heat are negligibly small, peaking less than 0.15 W/cm compared to powers over 400 W/cm in that region. The errors are larger near the peak core power, but still small; the largest difference is −2.86 W/cm compared to powers over 500 W/cm. No significant change is seen at the time step immediately preceding the SCRAM. At t = 1.0 s, the differences are a little more significant. The ORIGEN coupling predicts higher total power by up to 2.2 W/cm; however, the powers in this location are less than 60 W/cm, so this difference is close to 4%. This is still not an especially large difference, but it is not immediately negligible like the previous errors. After 10 s, ORIGEN predicts around 4 W/cm more

heat compared to less than 20 W/cm in those locations. This constitutes a 30% increase in heat production at those locations, which could certainly be significant under certain TH conditions. After 100 s, ORIGEN calculates about a 50% increase in heat compared with ANSI at the locations of the largest difference. However, these are not in the hottest rods like they were earlier in the transient. The hottest rods now show very little difference compared with ORIGEN.

These results indicate that the ANSI-based calculations might be overestimating how rapidly the decay heat precursors decay. This is seen in the the slightly faster rise in power during the early portion of the transient, indicating that the precursors are decaying more quickly than for the ORIGEN case. If this result holds at whole-core scales, then this will be important information to factor into safety analyses. The next section presents initial results for a whole-core test of this capability.

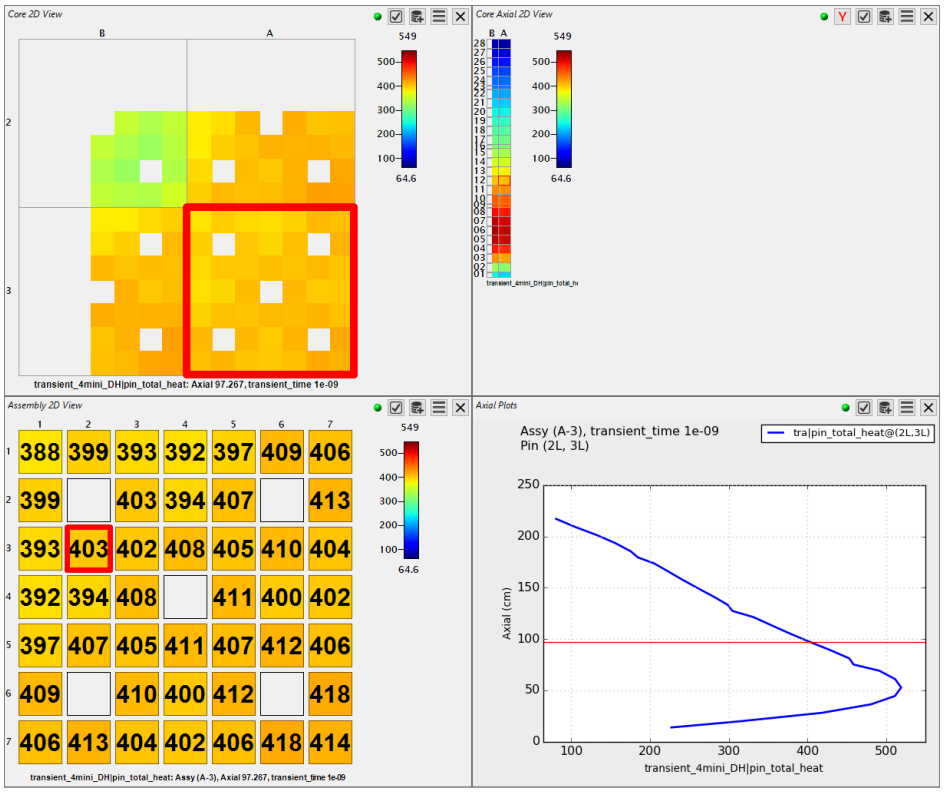

(a) Total heat predicted by ORIGEN decay heat coupled calculation.

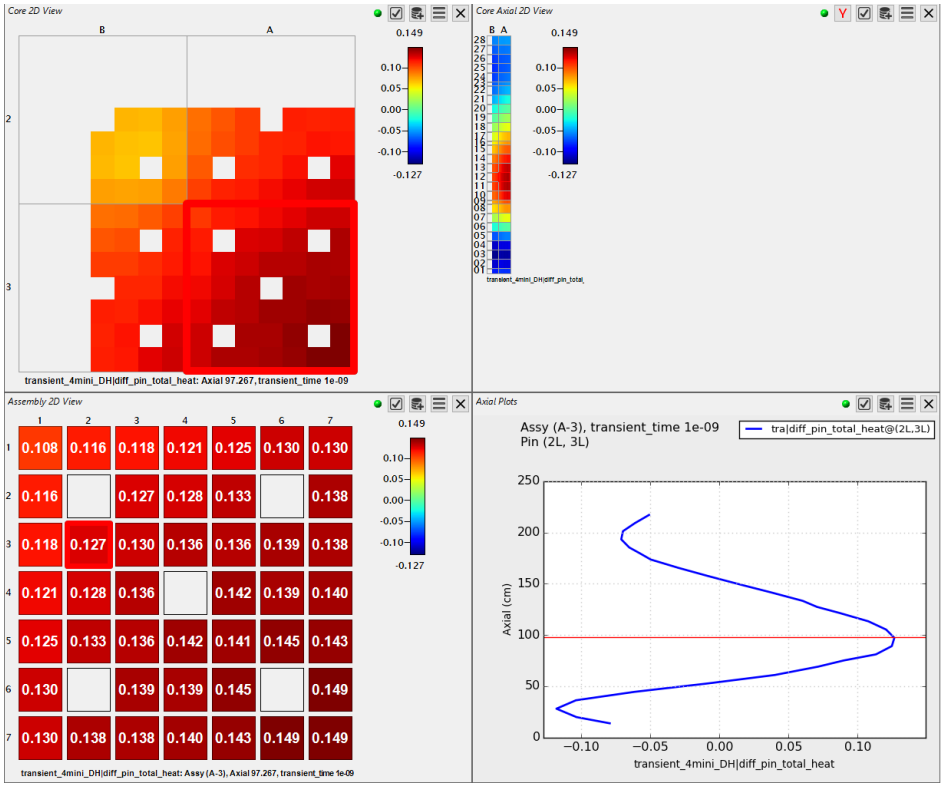

(b) Difference between ANSI and ORIGEN total heat calculations.

Figure 26. Heat production results during 4-mini transient at t=0.0

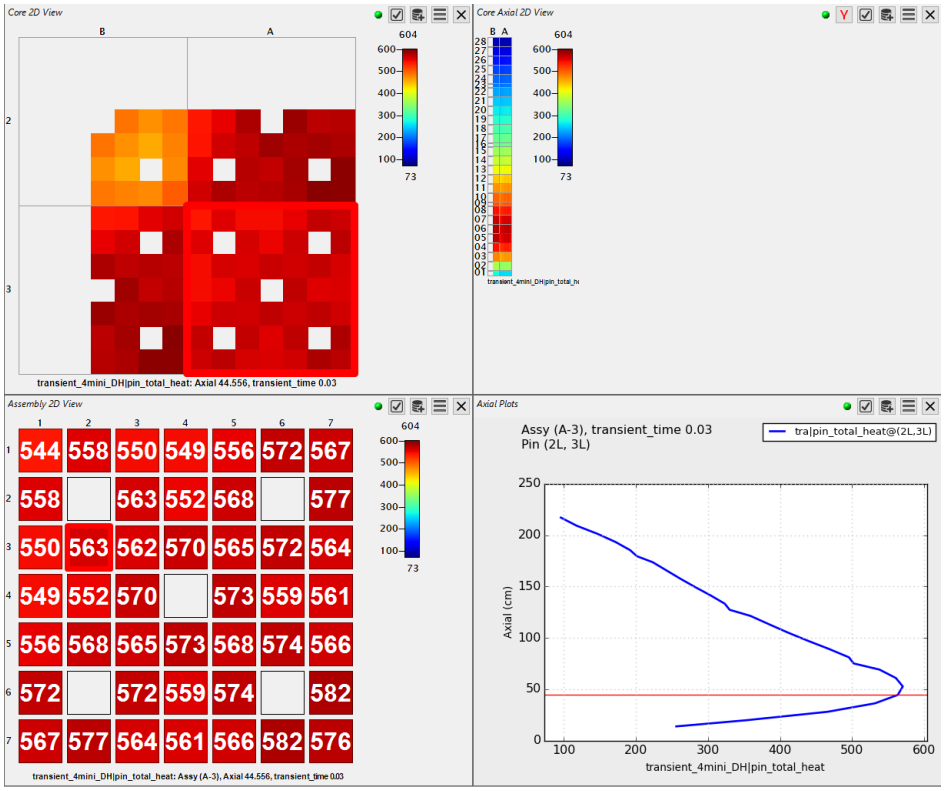

(a) Total heat predicted by ORIGEN decay heat coupled calculation.

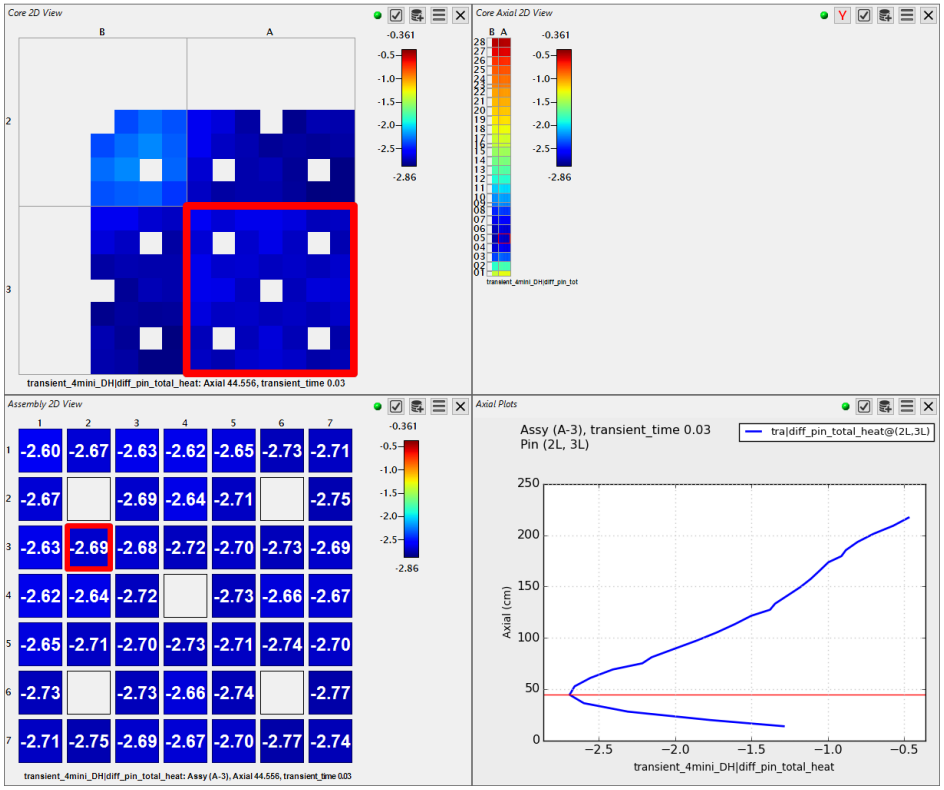

(b) Difference between ANSI and ORIGEN total heat calculations.

Figure 27. Heat production results during 4-mini transient at t=0.03.

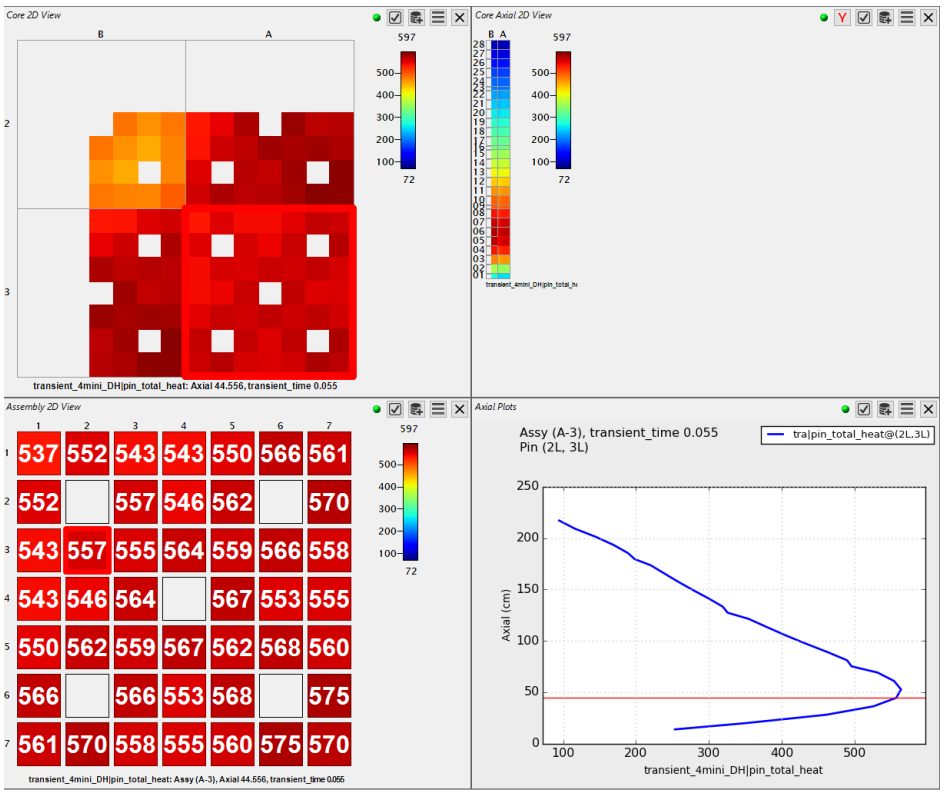

(a) Total heat predicted by ORIGEN decay heat coupled calculation.

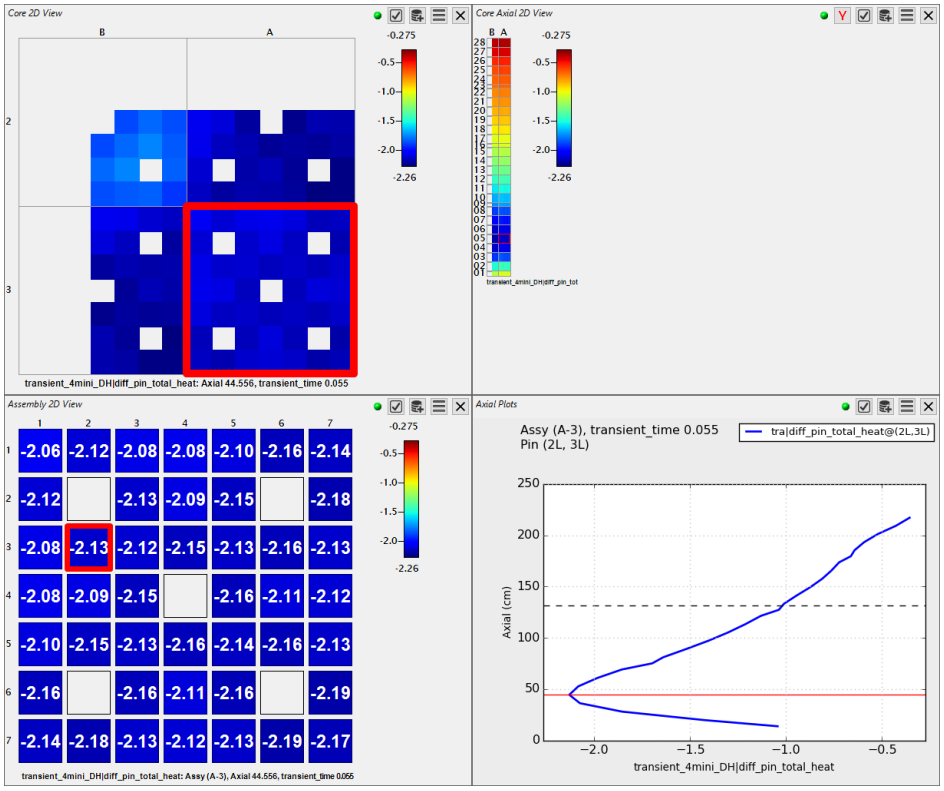

(b) Difference between ANSI and ORIGEN total heat calculations.

Figure 28. Heat production results during 4-mini transient at t=0.055.

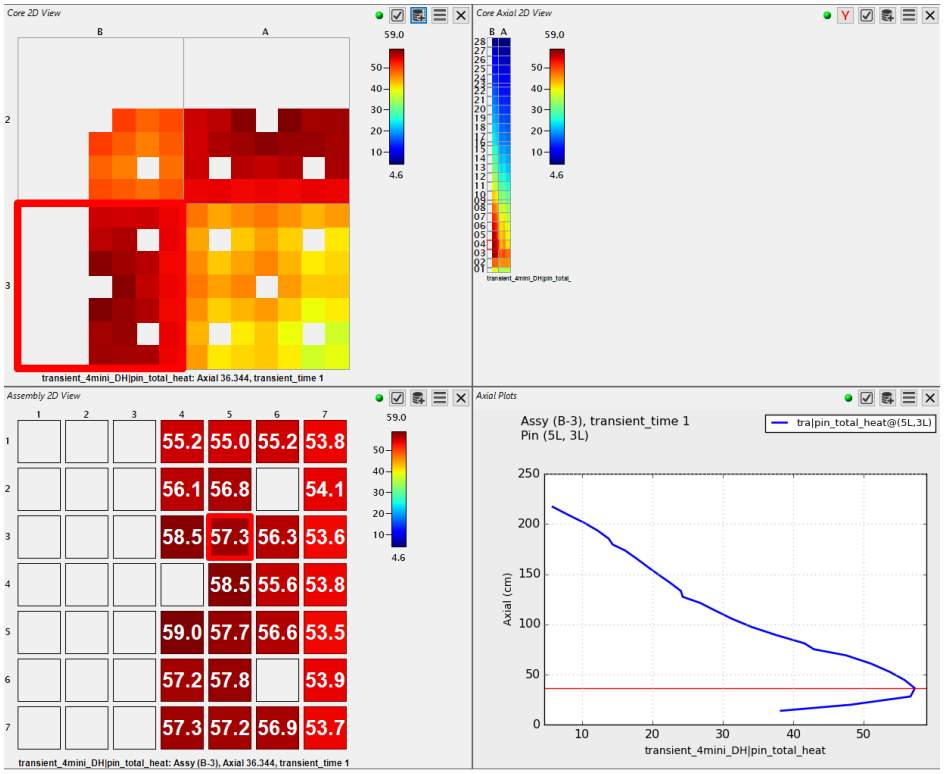

(a) Total heat predicted by ORIGEN decay heat coupled calculation.

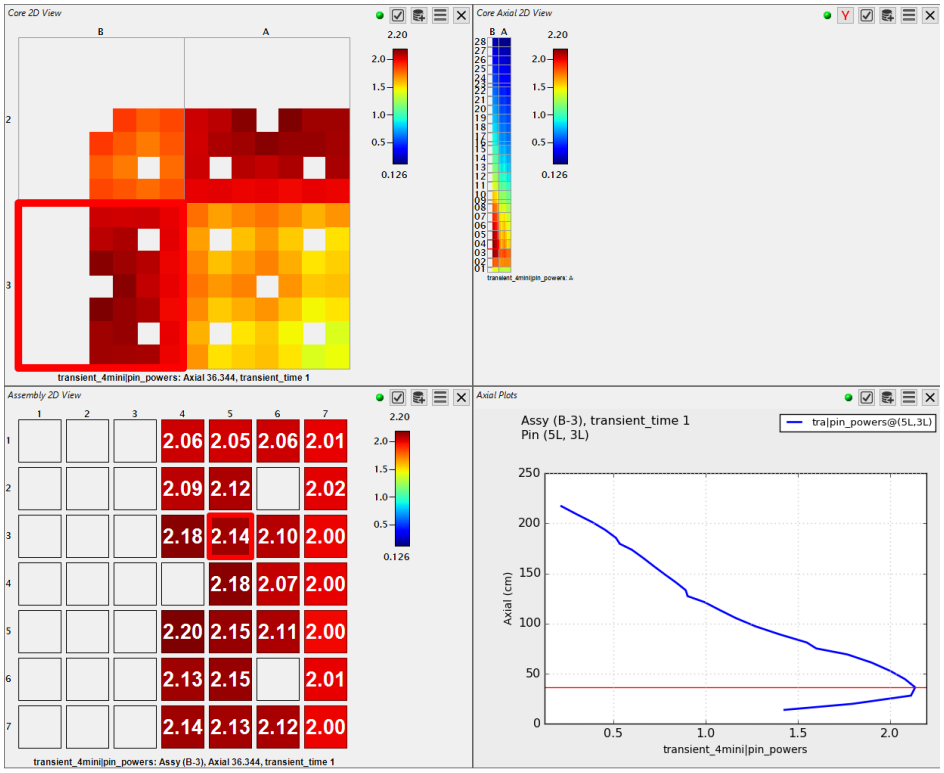

(b) Difference between ANSI and ORIGEN total heat calculations.

Figure 29. Heat production results during 4-mini transient at t=1.0.

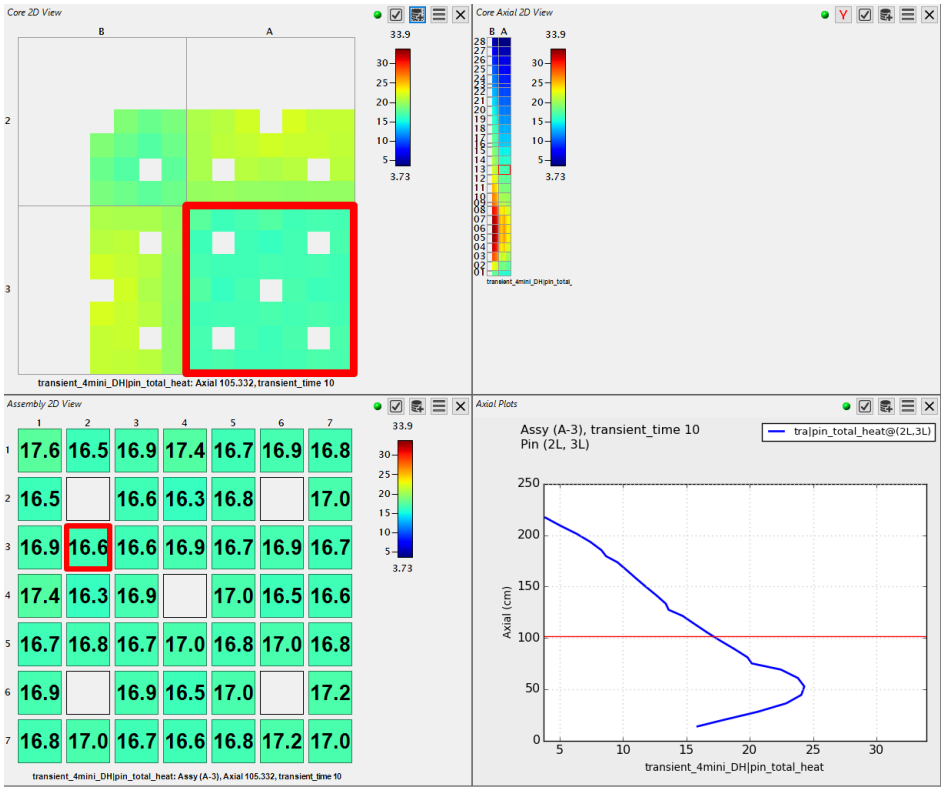

(a) Total heat predicted by ORIGEN decay heat coupled calculation

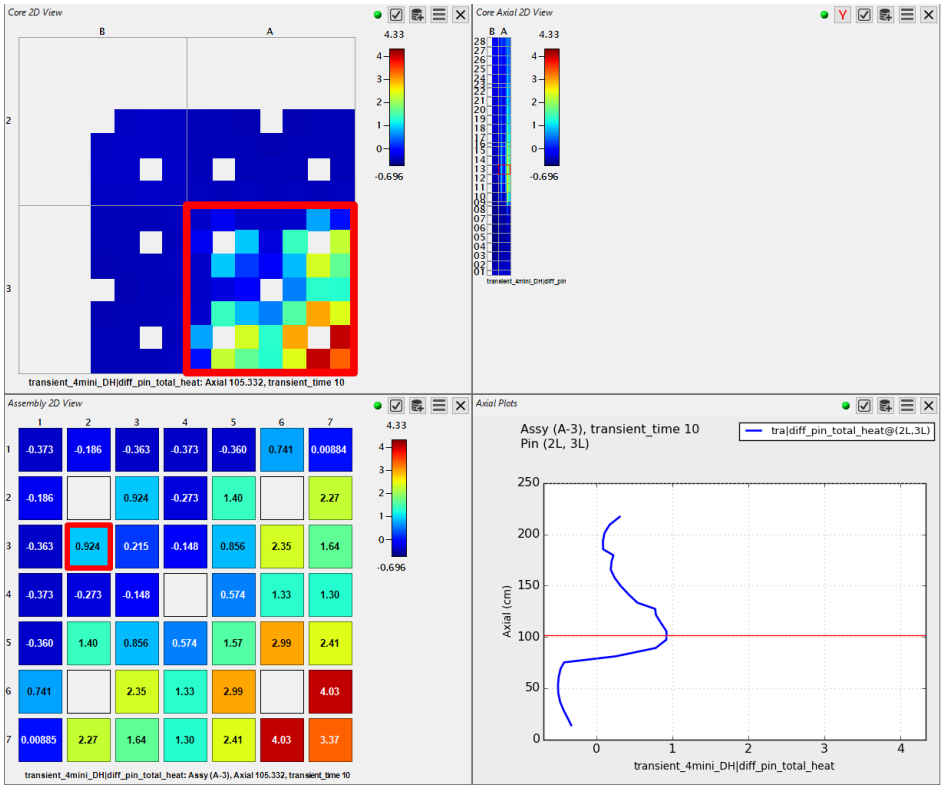

(b) Difference between ANSI and ORIGEN total heat calculations.

Figure 30. Heat production results during 4-mini transient at t=10.0.

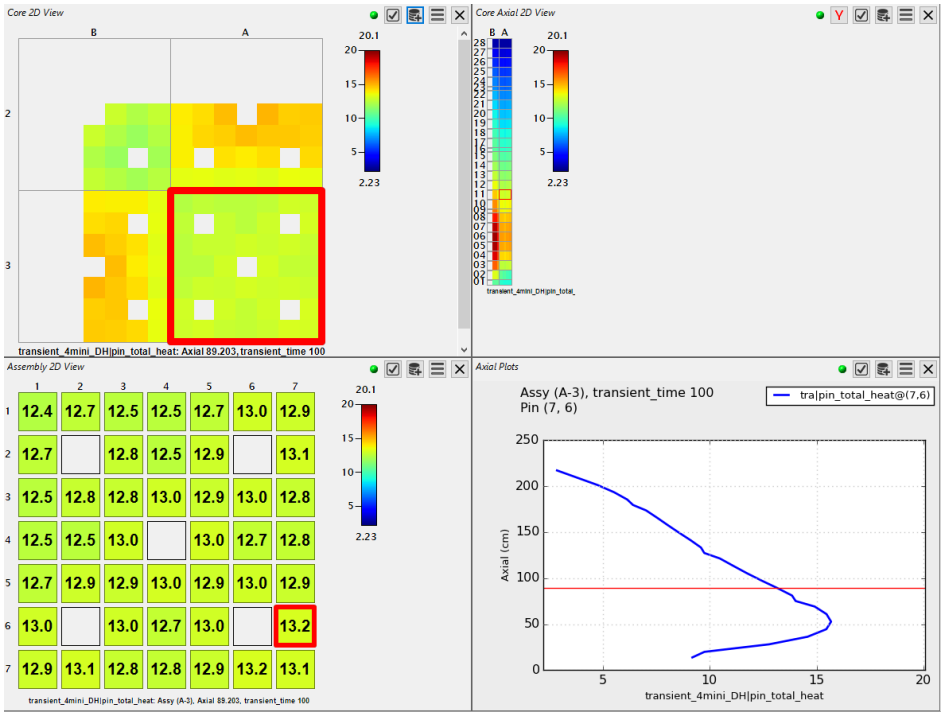

(a) Total heat predicted by ORIGEN decay heat coupled calculation.

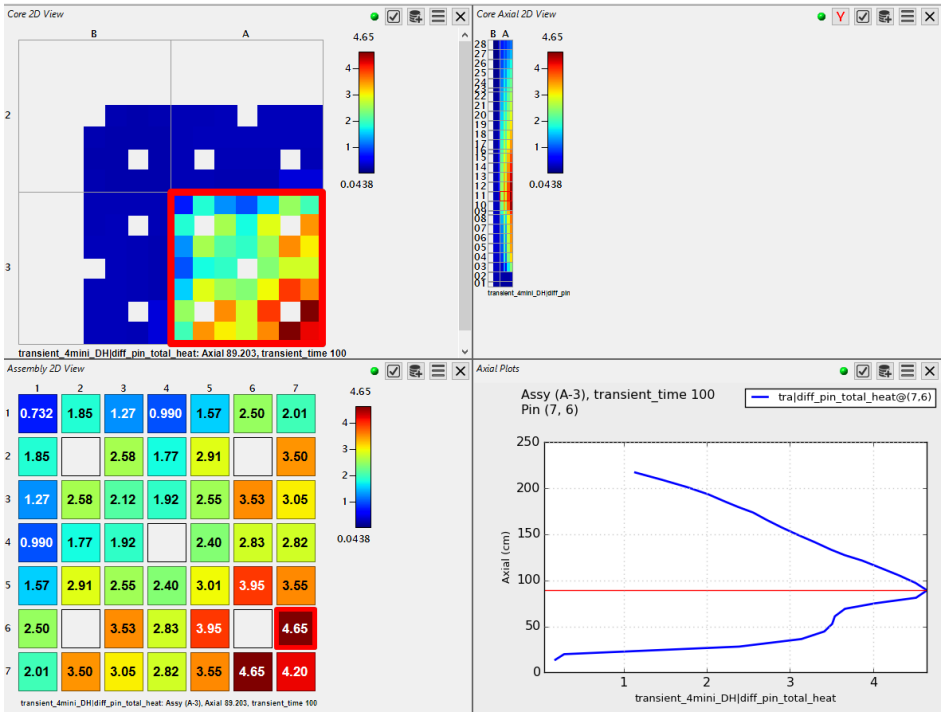

(b) Difference between ANSI and ORIGEN total heat calculations.

Figure 31. Heat production results during 4-mini transient at t=100.0

5.2 WATTS BAR UNIT 1 CYCLE 3

A full-scale demonstration was attempted using Watts Bar Unit 1. The first three cycles were depleted, then a LOCA was simulated at the end of cycle 3. This LOCA was simulated using both the ANSI and and ORIGEN approaches to decay heat calculation. The boundary conditions are shown in Figure 32. A SCRAM is initiated immediately at the start of the transient so that once the SCRAM is complete, the transient is driven exclusively by the decay heat instead of the prompt power. Unfortunately, at this time VERA is not able to model these conditions. An ongoing NEAMS activity (Milestone M2MS-23OR0701056: Complete Pin-by-Pin Transient Thermal-Hydraulic Assessment for High Burn-up Loss of Coolant Accident) focuses on enhancement of TH capabilities for LOCA conditions. It is expected that after completion of this milestone, this demonstration LOCA would be possible.

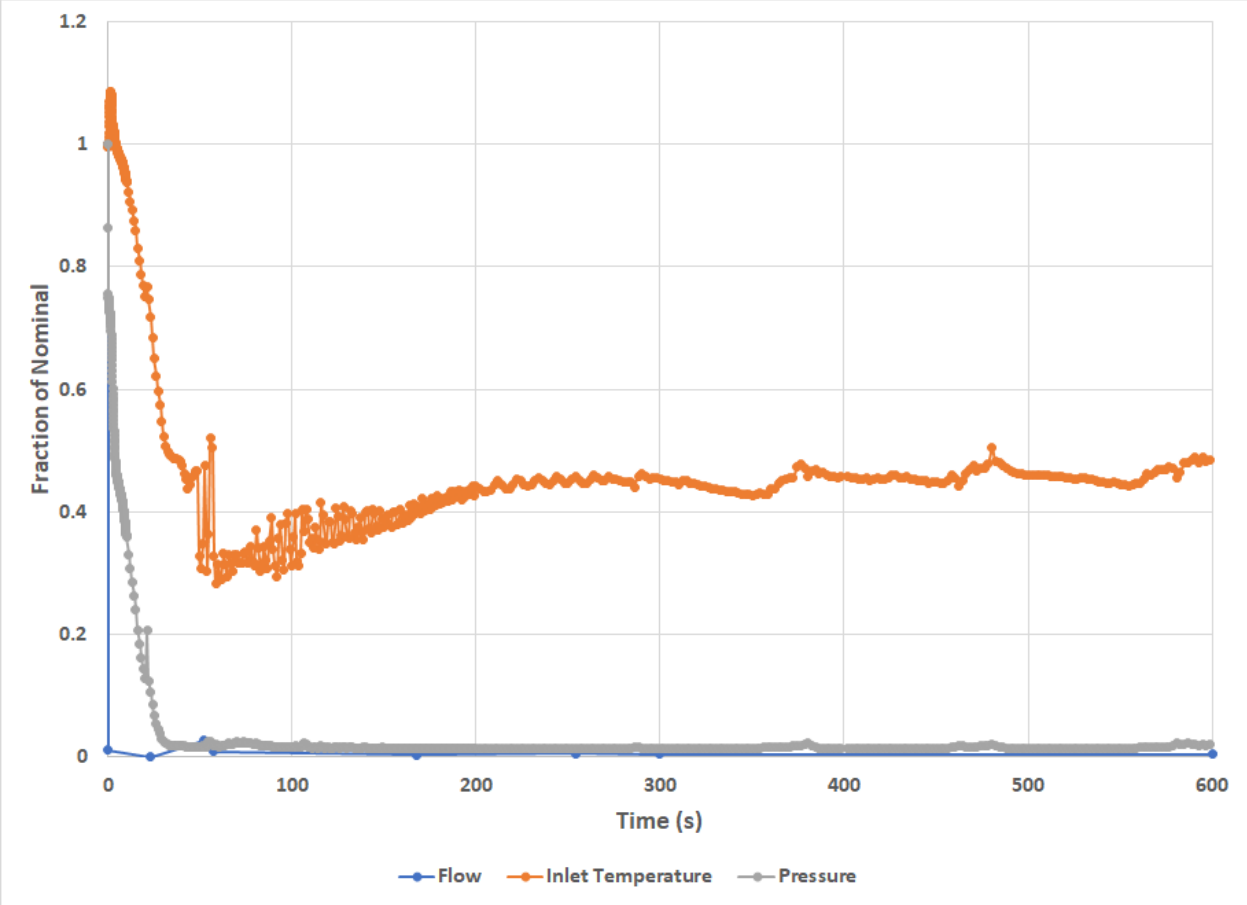

Figure 32. Simplified boundary conditions used for the LOCA demonstration.

To work around these limitations, the LOCA conditions were removed. The TH coupling remained enabled but used nominal boundary conditions instead of those shown in Figure 32. Instead, the reactor was simply shut down with a SCRAM at time $t = 0.0$ to calculate decay heat as a function of time post-shutdown. While the conditions are not the same as during a LOCA, the neutronics and decay heat should be similar enough to make some valid comparisons between the ANSI and ORIGEN decay heat calculations. Figure 33 shows the core power, prompt power, and decay power over time as percent of nominal operation power for the SCRAM test. Overall, the results are similar between ANSI and ORIGEN. ORIGEN's decay heat is slightly higher during the SCRAM but quickly drops below the ANSI power and remains there. Based on this, ORIGEN predicts more build-up of decay heat precursors than ANSI, but also more rapid decay for the

most short-lived precursors. This is consistent with the 4-mini transient results.

To more carefully investigate the differences between the two decay heat approaches, spatial heat distributions are shown at the following times for the SCRAM calculation:

1. t = 0.0 s, SCRAM initiation
2. t = 2.3 s, SCRAM completion
3. t = 10.0 s
4. t = 60.0 s
5. t = 120.0 s
6. t = 350 s, end of ORIGEN simulation

At each time, the total heat, decay heat, and total heat difference distributions are shown for the location of the maximum total heat difference. The total heat is taken from the ANSI simulation, the decay heat is taken from the ORIGEN simulation, and the differences are ORIGEN minus ANSI. Thus, a positive difference indicates ORIGEN predicts more heat than ANSI, which is the situation of greatest interest for fuel performance. These distributions are shown in Figures 34–39 and help to show that there are no surprises in the distributions that were not apparent in the core power levels in Figure 33.

At the initiation of the transient, the greatest difference in total heat is just 0.907 W/cm, less than 1% higher than ANSI. At that location the total heat is essentially the same between the two simulations, which is consistent with prior steady-state simulations. Immediately after the SCRAM completes, the maximum differences are around -0.1 W/cm. It is noteworthy that at this point in the simulation, ORIGEN predicts less heat than ANSI at every location. While the differences shrink with time, the trend of ORIGEN predicting less heat never changes for the remainder of the simulation. This is consistent with the core powers in Figure 33, but it is in contrast with the 4-mini tests. Again, this shows that ORIGEN may predict a greater build-up of decay heat precursors than ANSI, along with a faster decaying away of the short-lived precursors. Thus, the 4-mini transient has additional build-up during the rod ejection that is not observed during the SCRAM test. This build-up would also not be expected to be seen in a LOCA simulation

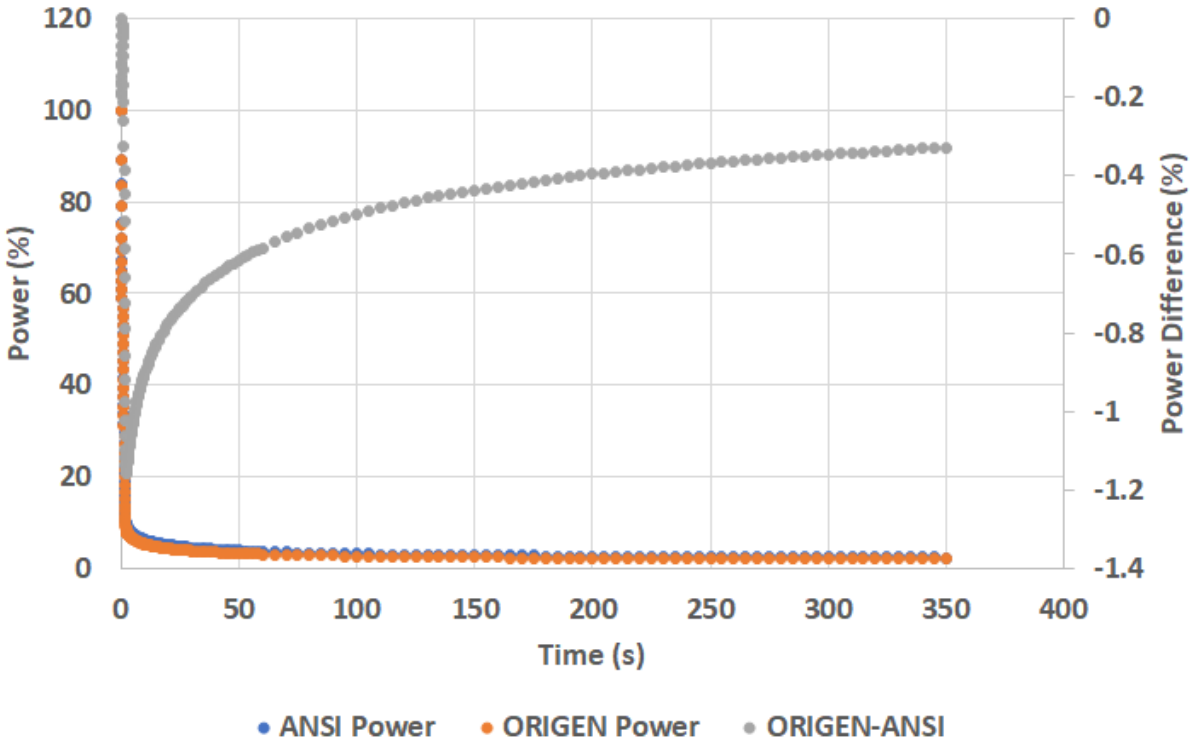

(a) Total Power

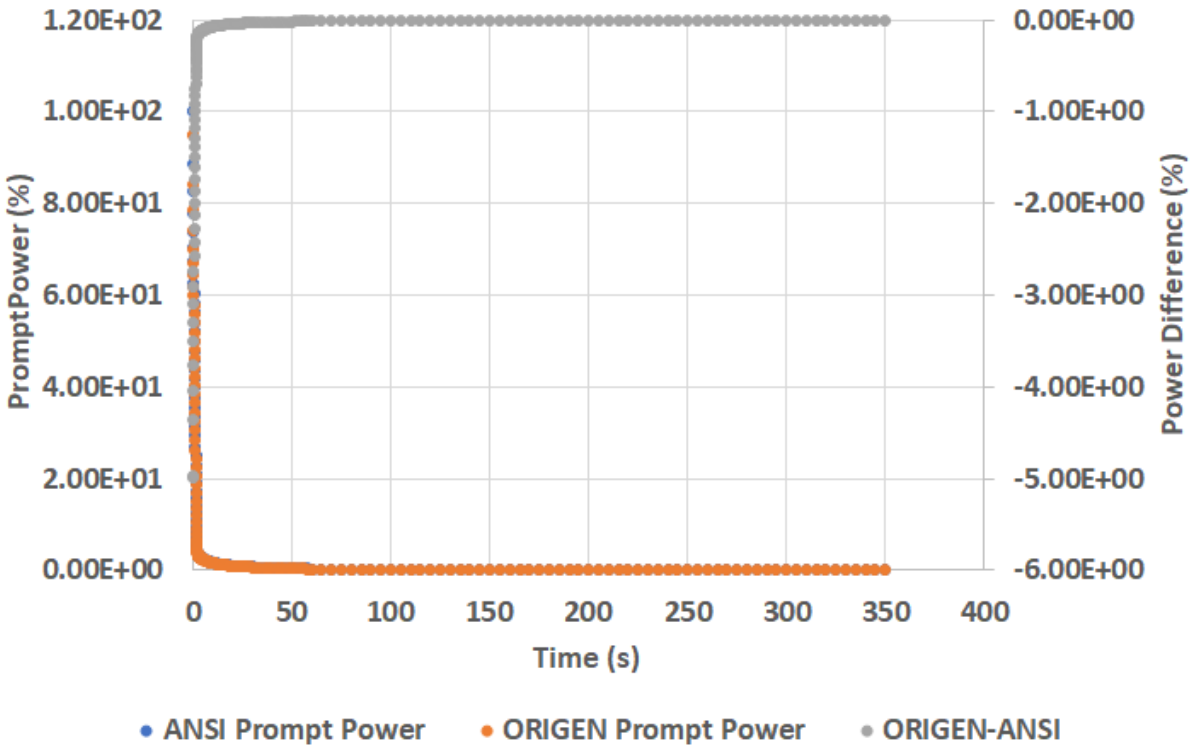

(b) Prompt Power

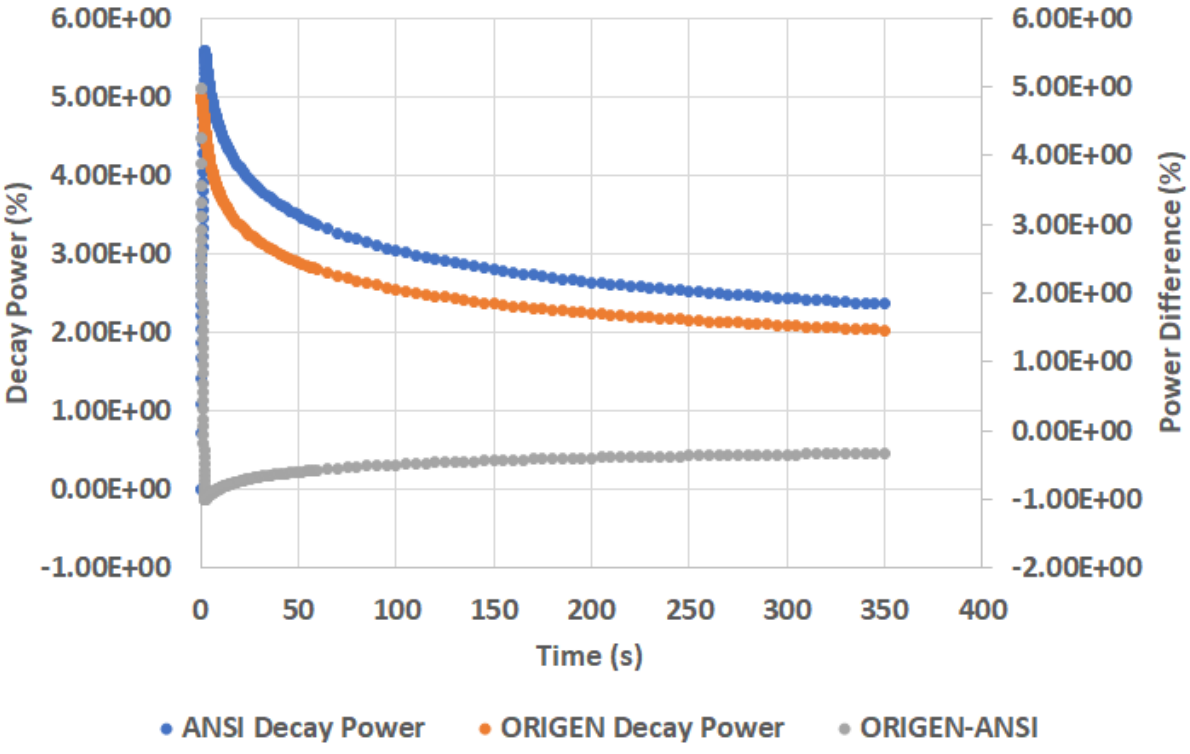

(c) Decay Power

Figure 33. Core power levels over time for SCRAM test

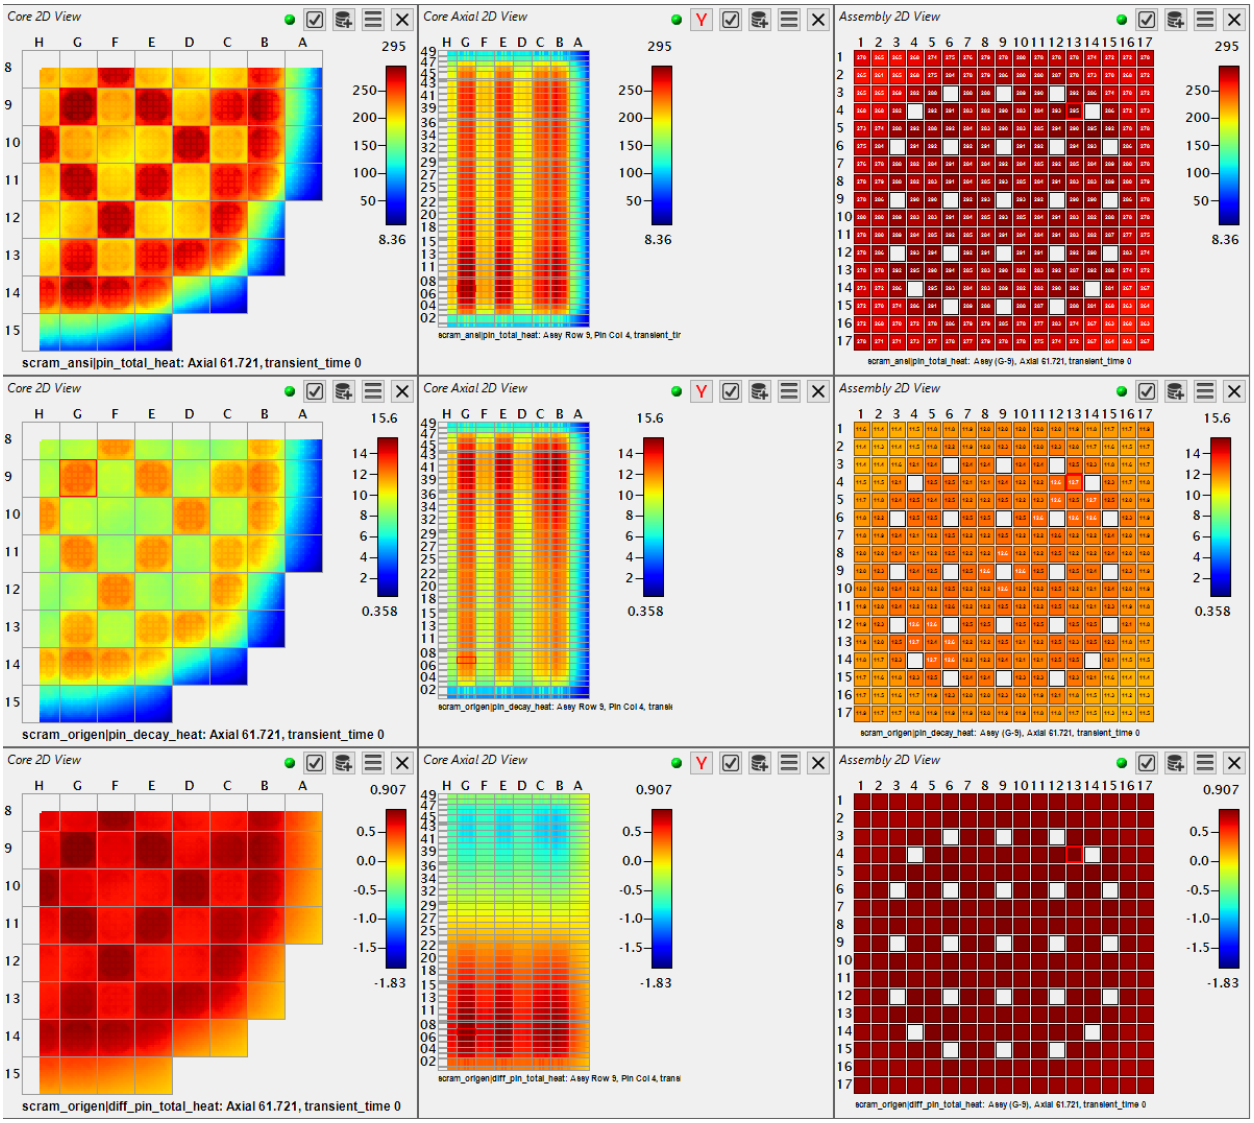

Figure 34. Total heat (top row), decay heat difference (middle row), and total heat difference (bottom row) for ORIGEN compared to ANSI decay heat calculations at t=0.0

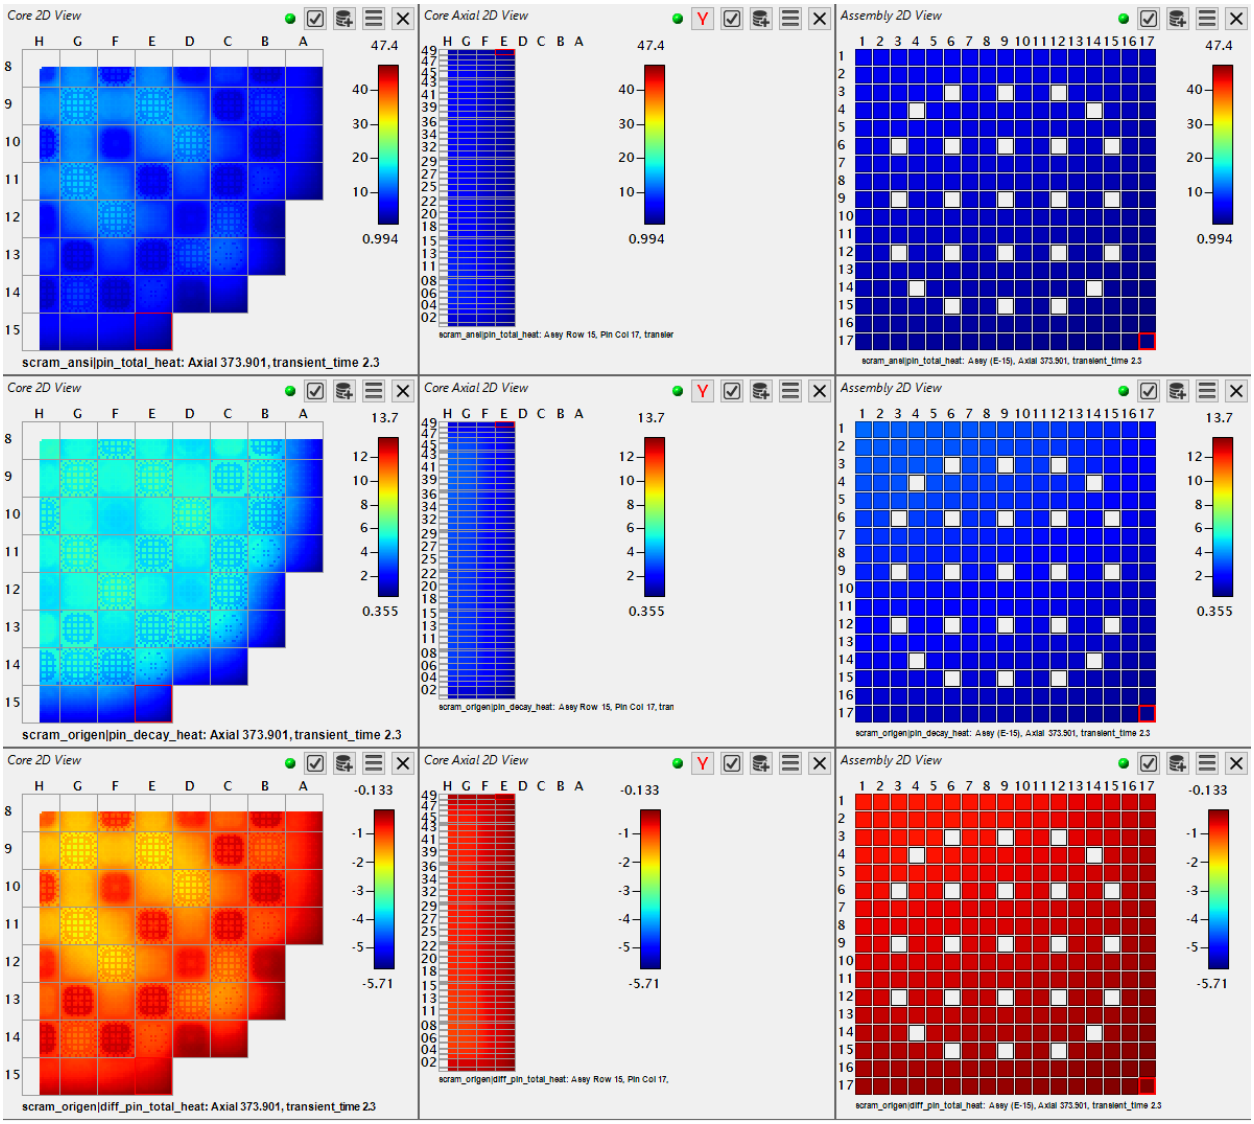

Figure 35. Total heat (top row), ORIGEN decay heat (middle row), and total heat difference (bottom row) for ORIGEN compared to ANSI decay heat calculations at t=2.3

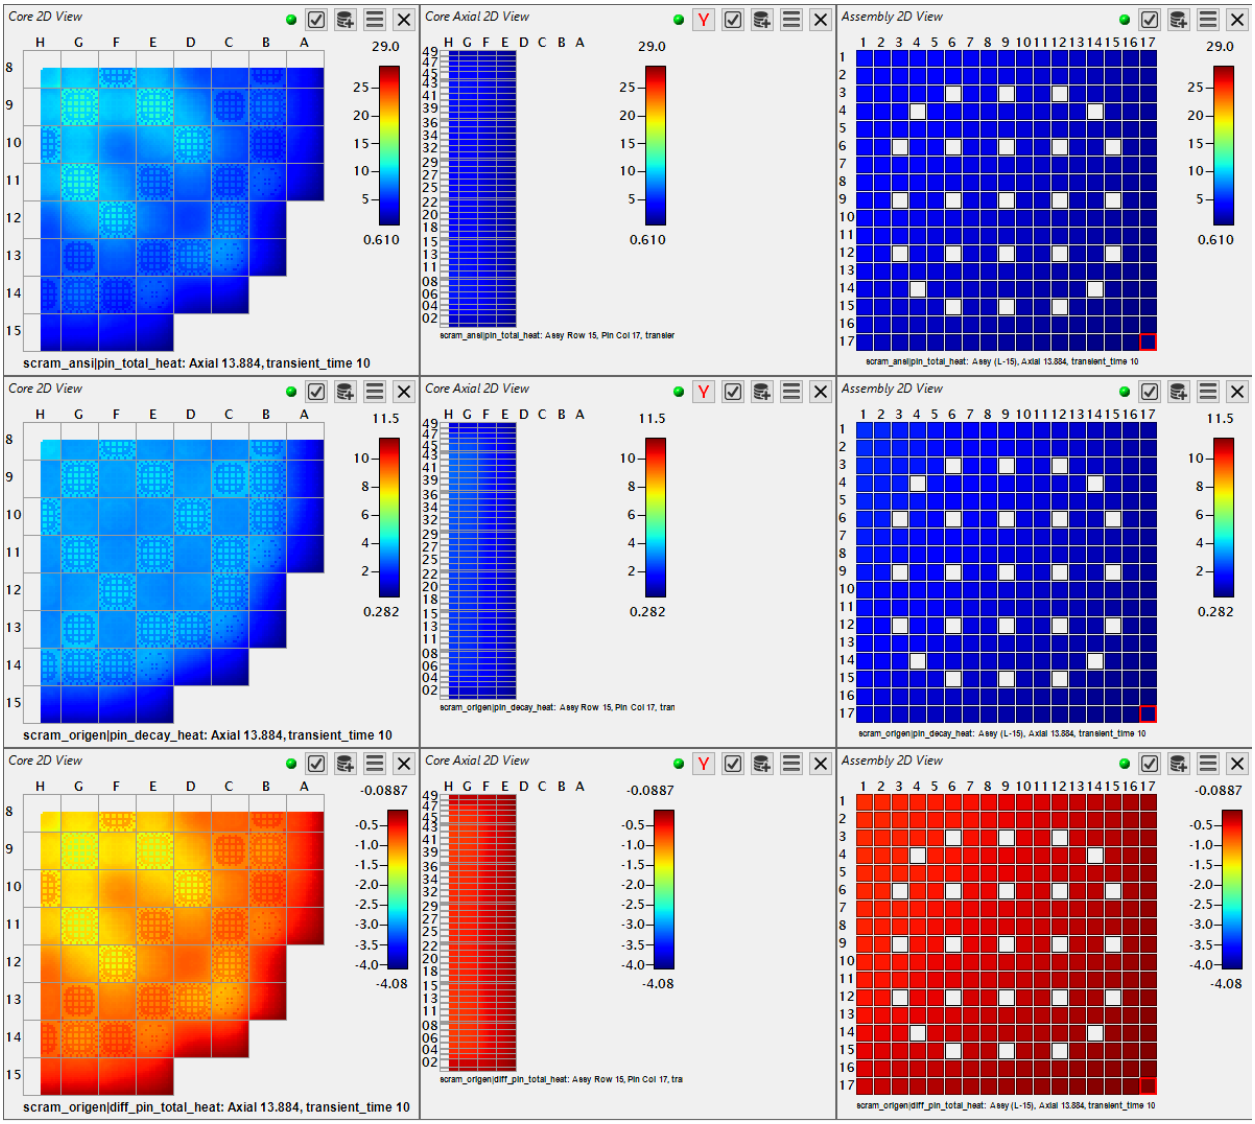

Figure 36. Total heat (top row), ORIGEN decay heat (middle row), and total heat difference (bottom row) for ORIGEN compared to ANSI decay heat calculations at t=10.0

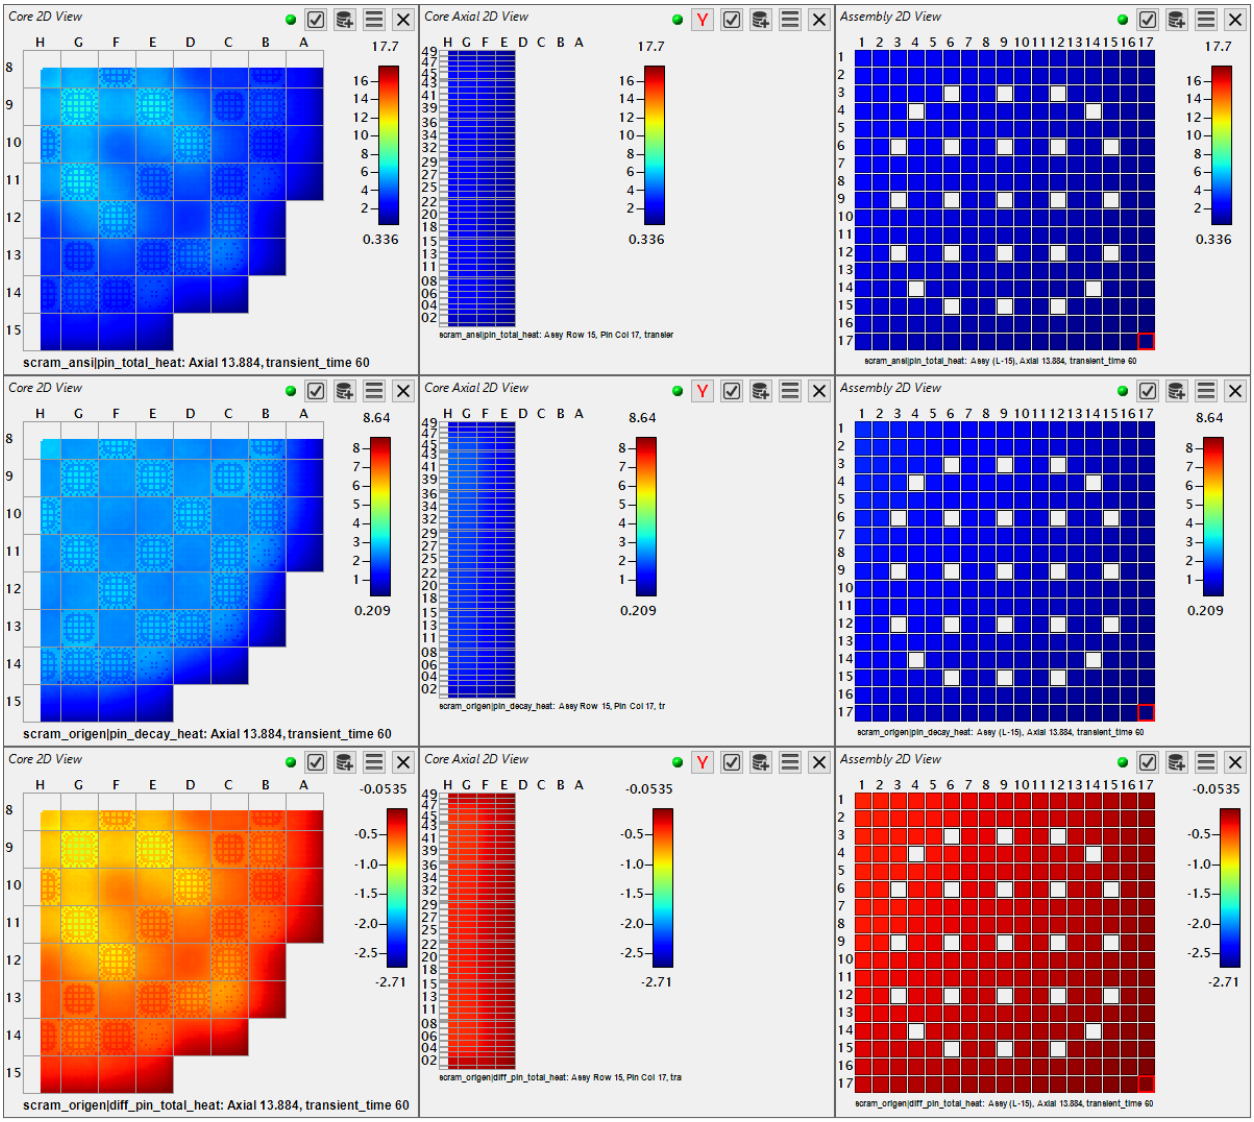

Figure 37. Total heat (top row), ORIGEN decay heat (middle row), and total heat difference (bottom row) for ORIGEN compared to ANSI decay heat calculations at t=60.0

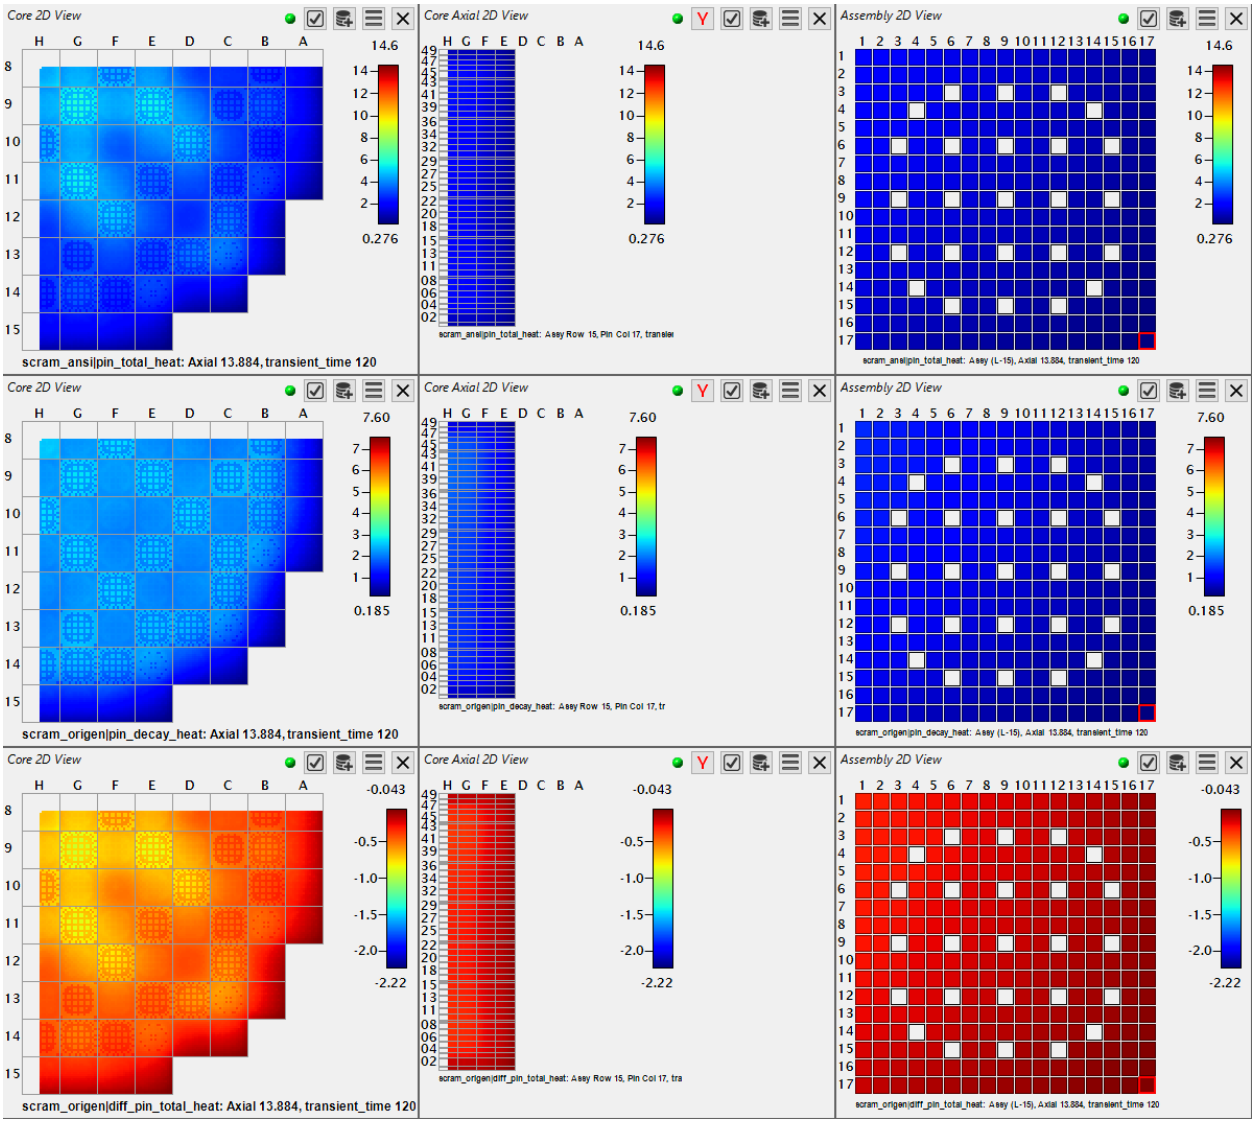

Figure 38. Total heat (top row), ORIGEN decay heat (middle row), and total heat difference (bottom row) for ORIGEN compared to ANSI decay heat calculations at t=120.0

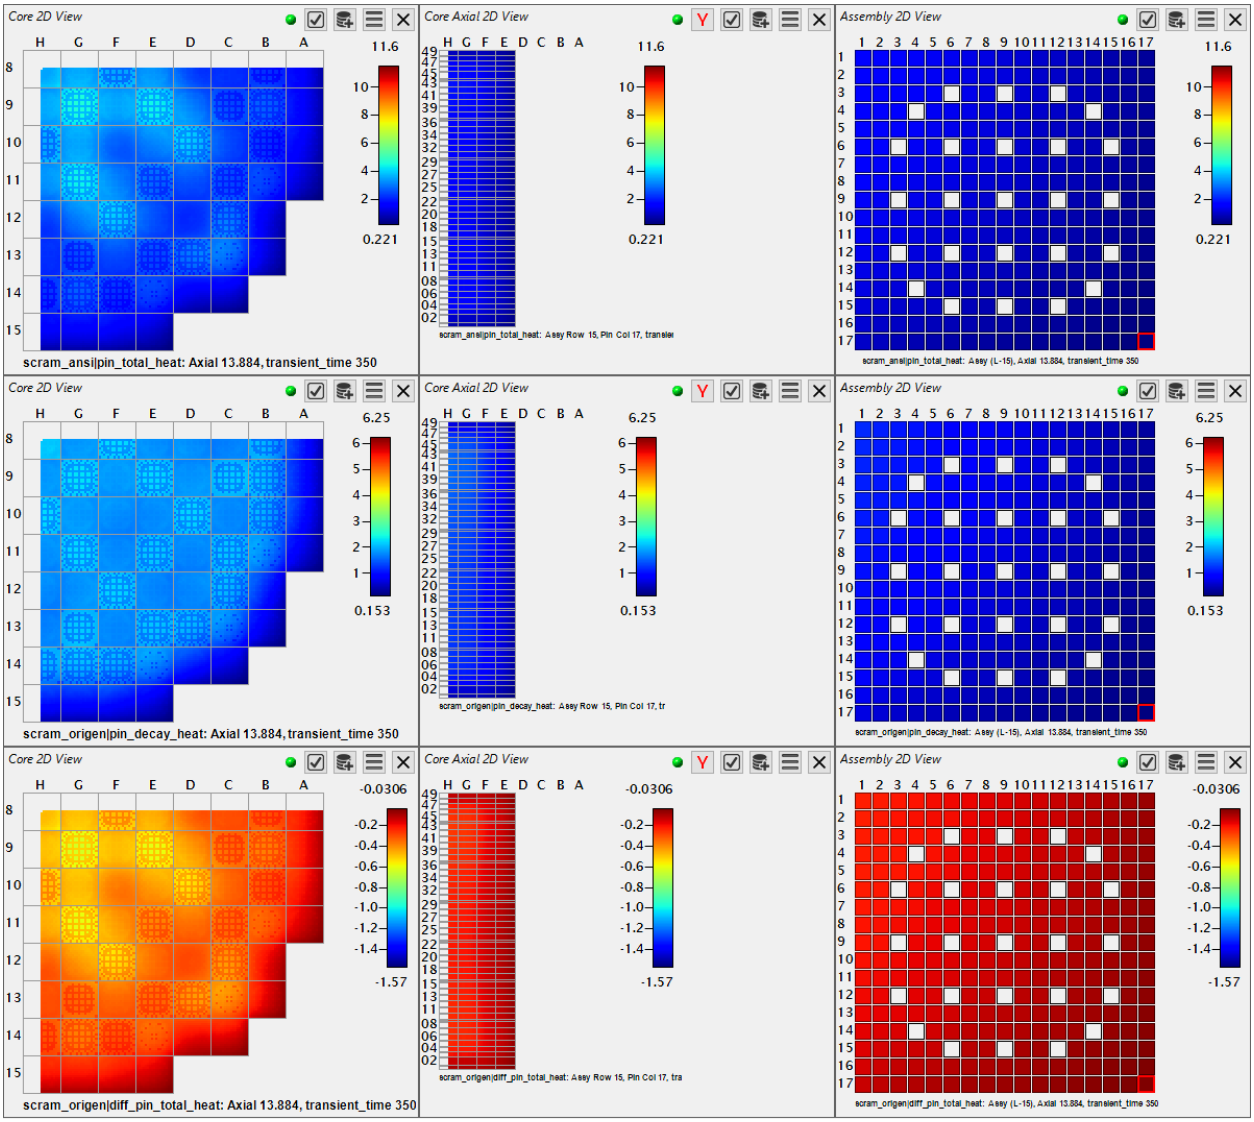

Figure 39. Total heat (top row), ORIGEN decay heat (middle row), and total heat difference (bottom row) for ORIGEN compared to ANSI decay heat calculations at end of ORIGEN simulation

6. CONCLUSIONS

Initial tests of this capability on a small control rod ejection indicate that the ANSI standard may be less conservative for coupled decay heat calculations during the first 10 min after a SCRAM. Using the ORIGEN coupling consistently results in higher power after shutdown compared to ANSI. Although there is some variation depending on which part of operation and which part of the core is analyzed, the largest errors at times when the core is driven primarily by decay heat indicate that ORIGEN predicts more heat than ANSI. However, when testing a whole-core SCRAM in preparation for LOCA simulation, the reverse was observed: throughout the entirety of the simulation, ORIGEN predicted slight lower heat than the ANSI-based approach. This is a positive result since it implies that the current approach was *more* conservative, and therefore switching to the higher fidelity approach will not suddenly require stricter decisions about cycle design and operation. Based on the 4-mini test, this positive result may not hold for a transient in which there is a prolonged power rise that drives additional build-up of precursors. However, for something like a LOCA where the reactor SCRAMs almost immediately, the ORIGEN decay heat calculations should predict similar or slightly lower heat sources than the ANSI calculations.

7. FUTURE WORK

7.1 ANALYSIS

One future work item planned for the near future is to perform additional analysis to compare the new ORIGEN decay heat coupling with the ANSI approach. Prior work by Wysocki et al. [6] performed LOCA analysis with the TRACE systems code. TRACE uses a lumped fuel rod approach to model the core and model the decay heat using the ANSI standards. Work is planned to repeat that analysis using the decay heat results from a VERA simulation to determine whether one approach is more conservative than the other.

There is also a known issue with the power normalization in VERA when using the decay heat coupling, as discussed in Section 3.3.2. The $\kappa\Sigma_f$ cross section in the VERA library includes all heat that will eventually be produced from fission, including decay heat. To correct this, the library should be modified to include a version of κ that excludes the decay heat contributions. This is a non-trivial task. Steady-state analysis revealed that this effect was most pronounced immediately after a major power maneuver and that the effects dissipate quickly; it has not yet been determined if the effects can be neglected for transient calculations as well; the limited transient testing that has been possible so far has shown a similar effect.

7.2 PERFORMANCE IMPROVEMENTS

For performance improvement, the addition of an option to disable neutron transport during certain transient conditions is planned. Once the TH is primarily driven by decay heat, no benefit is gained from running the high-fidelity neutron transport. Instead, VERA will be modified to skip these expensive calculations and perform only decay heat updates for the TH calculations.

A related performance improvement that can be pursued is to interpolate the decay heat source. The implementation presented here calls ORIGEN to update the decay heat every time step. These calculations require non-trivial compute time, but the decay heat changes very slowly compared to the size of the VERA time step. Instead, an approach will be pursued that stores the decay heat at the current time and a future time, and all intermediate time steps will interpolate between the two. This approach will save significant time in the ORIGEN calculations without sacrificing any significant accuracy.

8. ACKNOWLEDGMENTS

The authors wish to acknowledge Ben Collins for his initial efforts in the decay heat calculations in VERA prior to departing Oak Ridge National Laboratory (ORNL).

This work was supported through the DOE Office of 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 U.S. Department of Energy and the Nuclear Science User Facilities under Contract No. DE-AC07-05ID14517

REFERENCES

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

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

[8] ANS/ANS-5.1-2014. Decay Heat Power in Light Water Reactors, 2014.

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

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

[11] B. Kochunas et al. VERA Core Simulator Methodology for PWR Cycle Depletion. *Nuclear Science and Engineering*, 185(1):217–231, 2017.

[12] 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.

[13] I. Gauld et al. Isotopic Depletion and Decay Methods and Analysis Capabilities in SCALE. *Nucl. Technol.*, 174(2):169–195, 2017.

[14] William A Wieselquist. The scale 6.2 origen api for high performance depletion. Technical report, Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States), 2015.

[15] Robert Salko, Aaron Wysocki, Taylor Blyth, Aysenur Toptan, Jianwei Hu, Vineet Kumar, Chris Dances, William Dawn, Yixing Sung, Vefa Kucukboyaci, William Gurecky, Travis Lange, Xingang Zhao, Jordan Rader, Caleb Jernigan, Benjamin Collins, Maria Avramova, Jeffrey Magedanz, Scott Palmtag, Kevin Clarno, and Martin Pilch. CTF: A modernized, production-level, thermal hydraulic solver for the solution of industry-relevant challenge problems in pressurized water reactors. *Nuclear Engineering and Design*, 397, 2022.

[16] Yuxuan Liu, Brendan Kochunas, William Martin, and Thomas Downar. Delayed fission energy effect on LWR normal operation and transients. *Annals of Nuclear Energy*, 128:84–93, 2019.

[17] Erik Walker, Andrew Godfrey, Shane Stimpson, and Benjamin Collins. Effects of fuel temperature-shaping functions on xenon oscillations. In *Proc. of PHYSOR 2020*, 2020.

[18] Yuxuan Liu, Robert Salko, Kang Seog Kim, Xinyan Wang, Matthew Kabelitz, Sooyoung Choi, Brendan Kochunas, Benjamin Collins, and William Martin. An improved energy deposition model in MPACT with simplified gamma smearing and time-dependent delayed energy. In *Proc. CASL Symposium 2021*, 2021.

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