

Bypass Flow Model Implementation for VERA BWR



Aaron Graham
Bob Salko
Mehdi Asgari
Dave Kropaczek

**Approved for public release.
Distribution is unlimited.**

August 1, 2022



DOCUMENT AVAILABILITY

Reports produced after January 1, 1996, are generally available free via US Department of Energy (DOE) SciTech Connect.

Website osti.gov

Reports produced before January 1, 1996, may be purchased by members of the public from the following source:

National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
Telephone 703-605-6000 (1-800-553-6847)
TDD 703-487-4639
Fax 703-605-6900
E-mail info@ntis.gov
Website classic.ntis.gov

Reports are available to DOE employees, DOE contractors, Energy Technology Data Exchange representatives, and International Nuclear Information System representatives from the following source:

Office of Scientific and Technical Information
PO Box 62
Oak Ridge, TN 37831
Telephone 865-576-8401
Fax 865-576-5728
E-mail reports@osti.gov
Website osti.gov

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Nuclear Energy and Fuel Cycle Division

BYPASS FLOW MODEL IMPLEMENTATION FOR VERA BWR

Aaron Graham
Bob Salko
Mehdi Asgari
Dave Kropaczek

Date Published: August 1, 2022

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	iv
List of Tables	v
Abbreviations	vi
1. Introduction	1
2. Bypass Heating Models	2
2.1 Simple Models	2
2.2 HEM Model	2
2.3 STH Bypass Model	3
3. ANTS Model Improvements	6
3.1 Numerical improvements	6
3.2 Performance improvements	12
3.3 User interface improvements	15
4. Results	17
4.1 Fixed TH Conditions	18
4.2 Bypass Heating Model	20
4.3 TIP Comparisons	24
4.4 Performance	38
5. Conclusions and Future Work	40
6. Acknowledgments	41
References	41

LIST OF FIGURES

1	Calculation flow when using bypass heating with inlet orifice loss coefficient iteration	5
2	Peach Bottom Unit 3, Cycle 18, beginning of cycle void difference when ANTS outer iteration stopping criteria is decreased from $1 \cdot 10^{-5}$ to $1 \cdot 10^{-6}$	10
3	Profile of time spent in different sections of CTF/Alternative Nonlinear Two-phase Subchannel Solver (ANTS) for a full-core, steady-state solution of a depletion point of Peach Bottom Unit 3, Cycle 18 with 3D power distribution after performance ANTS improvements.	14
4	VERA rendering of H1C1 geometry	17
5	TIP comparisons for all state points for fixed bypass flow TH conditions with STH	19
6	K_{eff} for each TH treatment for H1C1	20
7	Pressure drop for bypass heating treatments for H1C1	21
8	Converged bypass loss coefficients for bypass heating treatments for H1C1	22
9	Maximum bypass void fraction for bypass heating treatments for for H1C1	22
10	Bypass exit void distribution at MOC for H1C1	23
11	RMS differences in TIP measurement comparisons for each TH treatment for H1C1	24
12	Max differences in TIP measurement comparisons for each TH treatment for H1C1	25
13	TIP comparison map for saturated conditions at 0.702176 GWD/MTU	26
14	TIP comparison map for STH+heating conditions at 0.702176 GWD/MTU	27
15	TIP comparison map for ANTS+heating conditions at 0.702176 GWD/MTU	28
16	TIP comparison map for saturated conditions at 1.28199 GWD/MTU	29
17	TIP comparison map for STH+heating conditions at 1.28199 GWD/MTU	30
18	TIP comparison map for ANTS+heating conditions at 1.28199 GWD/MTU	31
19	TIP comparison map for saturated conditions at 5.79378 GWD/MTU	32
20	TIP comparison map for STH+heating conditions at 5.79378 GWD/MTU	33
21	TIP comparison map for ANTS+heating conditions at 5.79378 GWD/MTU	34
22	TIP comparison map for saturated conditions at 10.3111 GWD/MTU	35
23	TIP comparison map for STH+heating conditions at 10.3111 GWD/MTU	36
24	TIP comparison map for ANTS+heating conditions at 10.3111 GWD/MTU	37
25	Number of coupled neutronics/TH iterations required to solve each state point for each TH model.	39

LIST OF TABLES

1	Description of STH input parameters	3
2	Summary of numerical controls that were developed and tested for ANTS.	7
3	Results sensitivity study to the crossflow turbulent mixing limiter for Peach Bottom Unit 3, Cycle 18	11
4	Results for the ANTS validation cases for different values of the crossflow turbulent mixing limiter. Only bundle cases measuring void are considered and all results shown are in units of percent void.	11
5	Change in results after adding the crossflow ceiling numerical control.	12
6	Summary of statistics comparing ANTS to experimental results before and after numerical controls (Items 1, 5, and 6) are added	12
7	Summary of key state points for H1C1	18
8	RMS of TIP comparisons for all state points for fixed bypass flow TH conditions and STH .	18
9	Summary of performance data for each H1C1 simulation	38
10	Time (s) per coupled iteration for each component of for H1C1 calculations	39

ABBREVIATIONS

AME	absolute mean error
ANTS	Alternative Nonlinear Two-phase Subchannel Solver
BiCGStab	biconjugate gradient stabilized
BOC	beginning of cycle
BWR	boiling water reactor
CASL	the Consortium for Advanced Simulation of LWRs
CMFD	coarse mesh finite difference
EOC	end of cycle
GNF	Global Nuclear Fuels
H1C1	Hatch Unit 1 Cycle 1
HEM	homogeneous equilibrium mixture
HPC	high performance computing
iFOA	industry funding opportunity announcement
INL	Idaho National Laboratory
LOCA	loss of coolant accident
LWR	light-water reactor
ME	mean error
MOC	middle of cycle
PB3C17	Peach Bottom Unit 3 Cycle 17
PB3C18	Peach Bottom Unit 3 Cycle 18
PETSc	Portable Extensible Toolkit for Scientific Applications
PLR	partial length rod
PWR	pressurized water reactor
RMS	root mean square
RMSE	root mean square of error
rRMSE	relative root mean square of error
STH	simplified TH
TH	thermal hydraulics
TIP	transverse in-core probe
VERA	Virtual Environment for Reactor Applications

1. INTRODUCTION

The Virtual Environment for Reactor Applications (VERA) [1] has been developed by the Consortium for Advanced Simulation of LWRs (CASL) over the past 10 years to address difficult problems facing commercial light-water reactors (LWRs) [2,3]. VERA comprises several high-fidelity physics packages, each of which models a different aspect of the reactor. These individual pieces are then linked to provide a fully coupled multiphysics solution for various LWR analyses.

The core simulator under VERA is composed of the neutron transport code MPACT [4,5], the sub-channel thermal hydraulics (TH) code CTF [6], and the depletion code ORIGEN from the SCALE package [7]. In addition to the core simulator, other codes are available for more specific analyses, such as the fuel performance code Bison [8] and the coolant chemistry code MAMBA [9].

The initial focus of VERA development was on modeling pressurized water reactors (PWRs), but during FY20 and FY21, significant efforts were made to extend VERA to cover boiling water reactors (BWRs). These efforts focused primarily on (1) extending the geometry capabilities to capture things like control blades and channel boxes and (2) improvements to the thermal hydraulics models and coupling for two-phase flow. Although much progress was made, challenges persisted in coupled BWR calculations by the end of this effort.

To complete the BWR capability, two main areas were identified for improvement during FY22. First, neglecting the heating in the bypass flow was identified as a source of error in the solution. This report presents how the bypass flow was initially treated and outlines the improved treatments added to VERA to improve the accuracy of the calculations. Second, the TH calculations were previously identified as a major source of computational expense. A new TH solution method called ANTS was added to CTF, which significantly reduces the TH solve time. ANTS provides a more efficient, stable solution for two-phase flow compared to the standard CTF methods while retaining the same spatial resolution; however, numerical stability issues were frequently encountered when using ANTS for core-scale BWR depletion simulations. Significant efforts were put into improving the numerical stability of ANTS in FY22 to allow for more robust and efficient coupled BWR calculations. These improvements are also presented in this report.

2. BYPASS HEATING MODELS

During testing of BWR models with VERA, it was discovered that the treatment of the bypass flow could have a significant impact on the axial power levels. Each of the attempted models are discussed in this section. The first set of models simply used fixed values for the temperature and density of the bypass flow. The next model applied an axial distribution of temperature and density, which was calculated using a homogeneous equilibrium mixture (HEM) model. Finally, the most accurate of the models was a spatially dependent model built on MPACT's simplified TH (STH) solver.

2.1 SIMPLE MODELS

The initial bypass treatment was simply to leave the bypass flow at inlet conditions. This had the effect of overestimating the power in the top portion of the core. To determine the extent of this effect, two additional options were implemented. The first option was to apply outlet conditions. In this case, the coolant TH calculation was performed, the results at the top of all fuel bundles were averaged together, and the resulting temperature and density were applied to the bypass flow everywhere. The second option was to apply saturation conditions. For this case, the temperature and density were calculated for saturation at the core outlet pressure and applied to the bypass flow. This results in a decreased water density compared to inlet conditions, but with no voiding. This was expected to be the closest to reality compared to the inlet and outlet options.

Initial testing showed that each of these models produced dramatically different axial power shapes in the core. This will be demonstrated later in Section 4.1. This is because in an operating BWR, the bypass flow does in fact experience some voiding, which is a significant feedback effect on the neutronics. Thus, a more advanced model than these three fixed-value treatments was needed.

2.2 HEM MODEL

The first physics-based bypass flow model that was attempted was an HEM model. This model calculates an axial distribution of power and density based on the axial power shape of the core. Because the power distribution is calculated at the core level, the result is a single axial shape for the coolant that is then applied to every bypass flow region. The axial power shape is calculated as follows:

$$Q(z) = f_{bypass} \int_x \int_y P'''(x, y, z) dy dx, \quad (1)$$

where f_{bypass} is the fraction of power deposited in the bypass region, and P''' is the power density distribution throughout the core. Currently, VERA requires f_{bypass} to be input by the user as a fixed value. Based on measurements or calculations of the amount of voiding in the bypass flow, this value is expected to be in the 2–4% range, but it varies from plant to plant.

Once the axial power is calculated, the enthalpy h at each axial level can be calculated:

$$h_{bypass}(z) = h_{inlet} + \int_0^z \frac{Q(z)}{\dot{m}_{bypass}} dz, \quad (2)$$

where \dot{m}_{bypass} denotes the mass flow rate through the bypass region. The enthalpy distribution is then used to calculate temperature and density in the bypass.

Initial testing of this model showed substantial improvement in the core axial power shape. However, the axial power shape in each fuel bundle still showed unacceptable errors. This occurs because the voiding in the bypass region will be driven primarily by the neighboring fuel bundles, causing some bypass regions to experience a greater degree of voiding than others. Thus, the core-wide HEM model addressed the core-wide axial power distribution but does not address the radial power distribution. Results for this model will not be shown since it has already been replaced by a more advanced model.

2.3 STH BYPASS MODEL

To solve the issues with the radial power distribution, a spatially dependent bypass flow model is required. To accomplish this, the HEM model discussed in the previous section was discarded in favor of using MPACT's new BWR STH solver. This solver is significantly more advanced than the HEM model above and inherently assumes spatial resolution. This section briefly summarizes the STH solver physics and describes the details of the solver's application to the bypass region.

2.3.1 STH SOLVER

The BWR STH model is a void drift formulation based on the conservation of energy, mass, and momentum equations. Several approximations are made to the equations to simplify the solution procedure for two-phase flow, and a number of closure relationships are applied for the mixture density, mixture quality, void fraction, two-phase friction, and wall friction. The details of the solver are presented more thoroughly in the paper by Graham et al. [10]. The parameters required to set up the solver are summarized in Table 1.

Table 1. Description of STH input parameters

Input Parameter	Units	Description
Axial Mesh Thicknesses	cm	The thickness of each axial node for the fluid flow
Flow Areas	cm ²	The total flow area in each channel at each axial level
Wetted Perimeters	cm	The total wetted perimeter in each channel at each axial level
Heated Perimeters	cm	The total heated perimeter in each channel at each axial level
Inlet Orifice Areas	cm ²	The area of the inlet orifice for each channel; defaults to 10 in ²
Inlet Orifice Loss Coefficients	Normalized	Normalized inlet orifice loss coefficient for each channel; applied by dividing by the square of the inlet orifice areas
Forms Loss Areas	cm ²	The area for the forms loss of each channel at each axial level; defaults to 10 in ²
Forms Loss Coefficients	Normalized	Normalized forms loss coefficients for each channel at each axial level; applied by dividing by the square of the forms loss areas

The two most notable differences between this solver and the HEM solver regarding bypass flow modeling is that (1) it has both radial and axial spatial resolution, and (2) the STH solver can predict the voiding that occurs in the flow. Additional effects that factor into the amount of void produced are not captured by a simplistic calculation of density from enthalpy.

2.3.2 STH APPLIED TO BYPASS FLOW

The STH solver can be applied to the bypass flows like it already is for the channel. A single channel is set up for each assembly's bypass flow. The geometry of the channel accounts for the detailed outer edge of the channel box—including thick-thin channel boxes—and the presence of control blades. The flow areas increase and the perimeters decrease above the tips of the control blades. The heated perimeters are currently assumed to be equal to the wetted perimeters. This is known to be an approximation since some of the energy is deposited in the channel box and control blade, but some of it is deposited directly into the bypass flow. However, for the purposes of modeling the bypass flow, treating all the heat as coming from solid surfaces in contact with the channel is expected to be a reasonable approximation.

Given these channel descriptions, the STH solver calculates axial temperature and density distributions for the bypass flow around each assembly. As with the in-bundle coolant solve, the bypass flow around each assembly is rebalanced to obtain the same pressure drop in all bypass flow channels, which results in a radial flow distribution for the bypass as well. This information has a significant impact on the neutronics calculation for coupled simulations.

2.3.3 BYPASS HEATING PARAMETERS

Three major parameters control the behavior of the bypass STH model. These parameters are presently user inputs to VERA, and therefore reasonable values must be chosen. Each of these values and the value selection process is discussed in the following sections.

2.3.3.1 Bypass Flow Fraction

The bypass flow fraction is the most well understood of the three parameters at this time. This value is a state-dependent value that must be input by the user. For full-core models, the bypass fraction is typically input as a power- and flow-dependent table based on plant data. This allows the bypass flow to properly match the state-dependent power and flow without the user entering the bypass every time.

2.3.3.2 Bypass Heating Fraction

The second parameter is the bypass heating fraction, which is specified as a single value for the entire calculation. This controls the fraction of the power from each assembly that is deposited into the bypass flow and surrounding structures rather than in the coolant and fuel. Power from neighboring assemblies is not accounted for. This value is typically expected to be in the 2–5% range but could be highly problem dependent. Eventually, this could be replaced by a direct calculation of the energy deposition in the various regions of the core.

2.3.3.3 Bypass Inlet Orifice Loss

The final parameter is the inlet orifice loss for the bypass flow, which is specified as an assembly-wise map. For the channel TH calculation, the inlet orifice loss coefficients may be on the order of 1.0 or less; typically, the map has one value for the peripheral assemblies and a different value for the interior assemblies. However, for the bypass flow, this parameter is about two orders of magnitude higher, and there are typically no data available for the actual loss coefficient for the bypass flow. To avoid forcing the user to manually tune the distribution for the bypass flow, a search capability was implemented. If the distribution is not input, then the code will assume (1) that all assemblies have the same inlet orifice loss for the bypass flow and (2) that the pressure drop in the bypass regions must be the same as inside the bundle. Using these two assumptions, the bypass flow solver can iterate on the inlet orifice loss until the bypass pressure drop matches the in-bundle pressure drop. This iteration is implemented as a Newton iteration, which in practice takes only about three Newton iterations to sufficiently converge. The bypass flow solver is efficient enough that the increased runtime from the additional STH solves to obtain the loss coefficient is negligible.

2.3.4 COUPLED CALCULATION FLOW

The calculation flow for a coupled neutronics–TH solve with the bypass flow heating and inlet orifice loss coefficient search enabled is shown in Figure 1. The coolant TH is solved first, iterating until the pressure drop in all fuel bundles is the same. The bypass flow TH solve occurs next. Again, the solver iterates on the flow until the pressure drop in each bypass region is the same. Next, the bypass pressure drop is compared to the coolant pressure drop. If they are not equal, then the inlet orifice loss coefficients are updated, and the bypass flow TH calculation is repeated. This results in bypass flow temperature and density distributions that have a pressure drop consistent with the rest of the coolant.

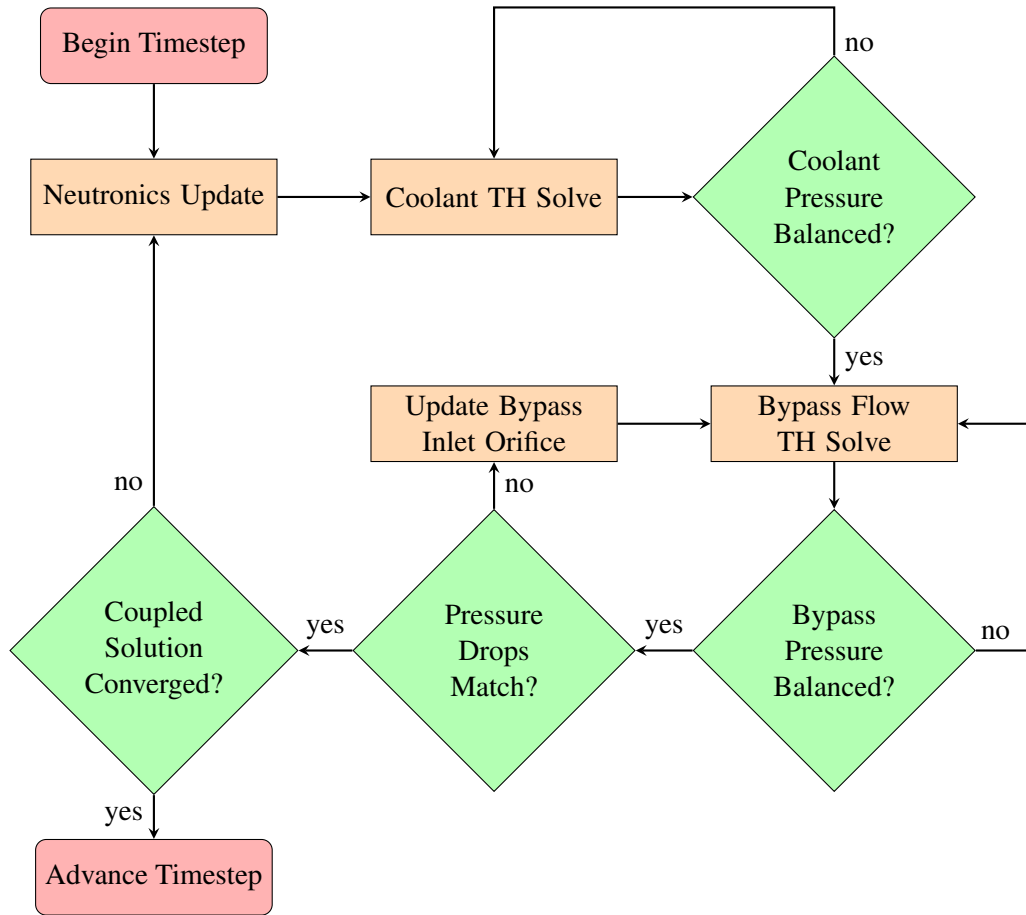


Figure 1. Calculation flow when using bypass heating with inlet orifice loss coefficient iteration

3. ANTS MODEL IMPROVEMENTS

ANTS is the computationally inexpensive subchannel capability developed during the BWR industry funding opportunity announcement (iFOA) project for providing neutronic thermal feedback in VERA BWR models [11–13]. The motivation for developing ANTS was to provide a more simplified solver than CTF that is still accurate for two-phase, pin cell resolution modeling of nominal BWR conditions, but which runs in a fraction of the time required by a CTF simulation. This goal is achieved with ANTS by using a four-equation model and by solving the model level-by-level, sweeping from inlet to outlet. This results in substantially less time to solve the system matrix but still allows for the calculation of pressure-directed crossflows between subchannels. Closure models for turbulent mixing, void drift, two-phase pressure drop, and phase slip are all included to accurately capture nominal BWR two-phase flow behavior.

The solution approach limits the solver’s application to upward flow, pre-dryout conditions with all boundary conditions supplied at the inlet, which is sufficient for modeling of a BWR depletion simulation. Other advanced features offered by CTF, such as droplet modeling and loss of coolant accident (LOCA)–specific models, are not included in ANTS. To facilitate integration into VERA, ANTS was integrated into CTF as an option accessible by changing an input flag. This provides ANTS access to the user interface, coupling interface, the fuel rod solver, and testing tools developed for CTF. ANTS was briefly introduced in the BWR iFOA final report [11]. A more detailed description of the code solution method is documented in Kropaczek [12]. Additionally, a comprehensive assessment including verification and validation testing was performed and documented in Salko [13].

3.1 NUMERICAL IMPROVEMENTS

Although ANTS was demonstrated to be fast-running for full-scale models and accurate for verification and validation cases, numerical instabilities were often encountered in certain geometries and operating conditions. For example, the disappearance of partial length rods (PLRs) in Global Nuclear Fuels (GNF) geometry often causes a large crossflow that can destabilize the numerical solution and certain low-flow and -power conditions that occur during normal BWR operation can lead to instabilities in some assemblies, causing liquid depletion in certain subchannels and subsequent code failure. To address these issues, a review of the numerical solution approach was performed and several improvements were implemented to make the solver more robust and capable of modeling the previously mentioned geometries and conditions. Three simulations were used for testing ANTS:

- Peach Bottom Unit 3 Cycle 17 (PB3C17),
- Peach Bottom Unit 3 Cycle 18 (PB3C18),
- Hatch Unit 1 Cycle 1 (H1C1).

The simulations were performed using VERA, which uses MPACT for the neutronics and the hybrid combination of STH and ANTS for the TH portion of the solution. Note that in this approach, the first five iterations of each statepoint are solved using the STH to get the power distribution solution closer to its final value and then ANTS is used for the remainder of the solves until the statepoint is considered converged. When starting this project, all three of these simulations were failing in the ANTS solve due to numerical stability issues. For the two Peach Bottom cases, the failure was traced back to liquid being depleted in one or more of the computational cells during the Newton solution process (mass flux crossflow being driven too high during initial iterations). This leads to a flow reversal, which quickly crashes the solution. The H1C1 case was run after some numerical improvements were made to ANTS and it was found that, while

Table 2. Summary of numerical controls that were developed and tested for ANTS.

ID	Description
1	Total crossflow change ceiling
2	Turbulent mixing crossflow ceiling
3	Under-relaxation of crossflow terms
4	Disabling crossflow for liquid deficient cells
5	Phase change limiter
6	Floor on axial vapor and liquid mass flow rates
7	Outer iteration controls

the ANTS solve was not failing, it was leading to a poor pressure distribution solution, which prevents the BWR pressure balance loop from converging and crashes the VERA simulation.

The ANTS numerical solution works by solving one level at a time and sweeping from inlet to outlet. Upward flow is assumed, and the solution is propagated from the upwind level to the downwind level. At each level, the lateral momentum equation is solved for each connection (gap) between subchannels so that the pressure drop in all subchannels is balanced. The crossflows are adjusted using a Newton method so that the axial mass flow rate distribution satisfies the balanced pressure drop. The adjustment of the crossflow terms were originally unbounded, meaning the solution of the Newton iteration for crossflow adjustment was directly applied before moving onto the next outer iteration or next level if the iterative approach was satisfied. Each assembly is solved independently and then the BWR pressure balance iteration loop is performed on the CTF side to adjust inlet flows to the assemblies so that a single, uniform core pressure drop is obtained for the TH solve. Note that this requires running ANTS every time the inlet mass flows are updated. In the early outer iterations for an axial plane, it is possible for the crossflow adjustments to be unrealistically large enough to fully deplete the liquid in a cell. The resulting axial mass flow rate update then results in a flow reversal, which is not realistic or compatible with the assumptions made in ANTS and, therefore, quickly results in the solution diverging or an arithmetic error being encountered. A collection of different limiters were tested to address the stability issues that were encountered for the VERA test cases. These controls are summarized in Table 2.

Not every one of the above-mentioned limiters were retained in the final version of ANTS, but the techniques and their effectiveness are documented here.

3.1.1 TOTAL CROSSFLOW CHANGE CEILING

A simple solution for controlling this behavior was to add a maximum allowable gap mass flux change. When updating the gap crossflow, the magnitude of the crossflow change is limited to a ceiling value. However, it is important to limit all crossflow mass fluxes by an equal ratio if any are limited so that the axial plane pressure drop response in all subchannels are consistent with the solution that was obtained from the Newton iteration loop. Eq. (3) shows how the limiter ratio is calculated.

$$\beta_x = \min \left[\frac{\Delta G_{\max}}{\max [|\Delta G_{\text{solved}}|]}, 1.0 \right] \quad (3)$$

The β_x value is the ratio that will be multiplied by all crossflow update values, ΔG_{solved} , which are solved by the Newton iteration. In the equation, ΔG_{\max} is the maximum allowable change to crossflow mass flux in one Newton iteration. Setting this value required some testing, but a final default value of $10 \text{ kg m}^{-2} \text{ s}^{-1}$ was set to ensure good convergence of the majority of assemblies and levels in the core simulation. Note

that the limiter has a ceiling of 1.0 to prevent from crossflow changes being larger than what was found in the Newton iteration. This limiter may impact the solution during the early iterations when solution changes are large, but should have no effect as the gap mass fluxes arrive at their steady-state values that are required for satisfying the lateral momentum equations.

3.1.2 TURBULENT MIXING CROSSFLOW CEILING

To prevent excessive turbulent mixing from depleting cell liquid, a ceiling was placed on this value relative to the adjacent channels that are promoting the mixing. The liquid turbulent mixing is calculated using the expression found in Eq. (4), where β is the turbulent mixing coefficient, \bar{G} is the average axial mass flux of the two connected channels, ii and jj , $\bar{\theta}$ is a two-phase multiplier, α is the volume fraction, and ρ_ℓ is the liquid density.

$$G_\ell^{TM} = \frac{\beta \bar{G}}{\bar{\rho}} \bar{\theta} \left((1 - \alpha_{ii}) \rho_{\ell,ii} - (1 - \alpha_{jj}) \rho_{\ell,jj} \right) \quad (4)$$

A limiter is defined to be relative to the average axial flow rate of the two connected subchannels. In equation form, the magnitude of the liquid turbulent mixing is limited as $|G_\ell^{TM}| < \epsilon_{TM} \bar{G}$, where ϵ_{TM} is a user defined value that should be between 0.0 and 1.0. After testing of the PB3C17 model, it was found that a limiter of 4e-4 (0.04%) resulted in good stability while not impacting the quality of the ANTS validation cases significantly.

3.1.3 UNDER-RELAXATION OF CROSSFLOW TERMS

Another limit on crossflow that was briefly tested was to under-relax the crossflow terms for the current level by pulling in the solution of the upstream level. A linear under-relaxation is done as follows:

$$M_j = (1 - \beta_{cf}) M_{j-1} + \beta_{cf} M_j. \quad (5)$$

Here, M is the crossflow term, which can be turbulent mixing or void drift of liquid or vapor mass, energy, or momentum. Note that these are the mixing terms that are calculated via the turbulent mixing and void drift models and not the pressure-directed crossflow term that is solved in the lateral momentum equation outer iteration loop. The β_{cf} is the under-relaxation coefficient, which shall be between 0.0 and 1.0 and the j index represents the axial level number. While this term did help to stabilize the solve for some conditions, it was found to significantly degrade the two-phase validation tests that were used to assess ANTS, so this control was ultimately not retained.

3.1.4 DISABLING CROSSFLOW FOR LIQUID DEFICIENT CELLS

Despite controlling potential large oscillations in the solution and magnitude of the liquid turbulent mixing with the limiter applied in Section 3.1.2, it was found that it was still possible for cell outflows to exceed inflow as liquid mass flow rate becomes small for certain sets of conditions. To robustly prevent liquid mass depletion and reversed flow in these cases, a more aggressive limiter was added on the crossflow. In the outer iteration loop, after the update of the crossflow term, a calculation of the updated axial mass flow rate distribution is performed, which takes into account the newly updated pressure directed cross flow, turbulent mixing, and void drift. If the axial flow of any subchannel becomes negative, then all of the pressure-directed crossflow terms for that subchannel will be reset to their old iteration values, and the turbulent mixing and void drift terms will be set to zero. An important note about this update check is that only the outflow terms for each subchannel are considered. It was found that if net flow is considered instead, then the crossflow

reset performed to prevent liquid depletion in one cell may ultimately result in depletion in the adjacent cell because the crossflow that was increasing mass in the cell was just cut off.

Although this last control is more aggressive, ANTS is designed for nominal BWR operating conditions where the final solution will not involve any liquid depleted cells; rather, the maximum volume fraction is anticipated to be less than 95% in the model. These numerical controls were implemented to prevent temporary liquid depletion that may occur as the iterative process approaches the final solution.

3.1.5 PHASE CHANGE LIMITER

The phase change rate is calculated using the Lahey [14] mechanistic model, which accounts for direct heating, evaporation, and condensation effects. This model does not account for the fact that the calculated phase change rate may exceed the amount of liquid available in the cell. It is also possible that, while the phase change rate does not exceed liquid mass at time of calculation, the crossflow effects will combine with phase change to deplete the cell. To prevent the phase change rate from depleting the liquid in the cell, a check is performed for each cell in the plane to first see if the cell will be depleted. This check is only performed if the phase change rate is positive (meaning it leads to evaporation). If a depletion will occur, the phase change rate is limited to 90% of the previous level liquid mass flow rate. Note that the new level liquid mass flow rate is calculated by summing the effects of crossflow and phase change to the previous level value.

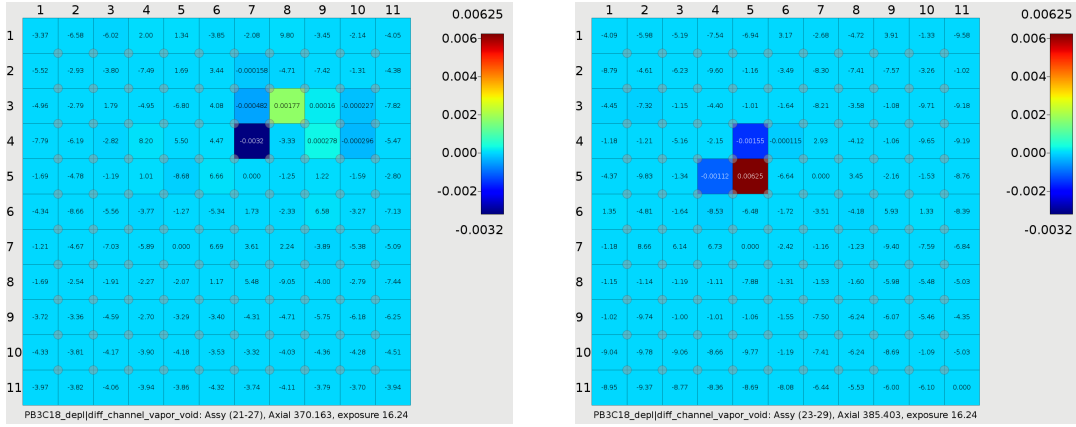
3.1.6 FLOOR ON AXIAL LIQUID AND VAPOR MASS FLOW RATES

It was found that, even for the more aggressive control of completely disabling crossflow as described in Section 3.1.4, some cell liquid depletion was still sometimes encountered. An even more aggressive limiter was applied by forcing a floor of zero for both the liquid and vapor axial mass flow rates. This limiter can cause mass balance issues because resetting the flow from a negative number to zero means that the magnitude of the mass flow rate was physically removed from the system. While this control presents a risk for convergence of the system, it is very effective at preventing unrealistic flow reversal scenarios. Again, because ANTS is not intended to be used for dryout situations, it is anticipated that this limiter will only be used during the iterative process and will not impact the final solution. To track the impact of this control, an internal flag in ANTS is used to track when the limiter is used anywhere in the plane. If this flag is still enabled after reaching convergence in the plane, a warning is printed to the screen output and log file alerting the user of the limiter as well as what assemblies and channels are being affected.

3.1.7 OUTER ITERATION CONTROLS

The numerical controls that were implemented into ANTS were tested for the three coupled simulations that were identified at the beginning of this subsection. It was found that the PB3C17 cases will run successfully with the implementation controls. As previously mentioned, it was found that the H1C1 case still encountered difficulty due to a poorly converged pressure distribution. It was believed this was caused by the intra-assembly pressure distribution not being sufficiently converged, which can lead to spurious oscillations that will make the BWR pressure balance loop struggle. This was addressed by further assessing the outer iteration tolerance and maximum number of outer iterations to ensure that the levels were being suitably converged.

As a first step, the convergence edit information was greatly expanded as later discussed in Section 3.3.2. After this, it was found that several assemblies and planes were not reaching their tolerance. It was found that increasing the maximum number of outer iterations from 10 to 50 leads to a much larger portion of the assemblies converging. Despite this, there are still often levels of some assemblies that will not reach their tolerance due to the solution starting to oscillate; however, void and pressure metrics that were added indicate that these oscillations are rather small in terms of the impact on the solution.



(a) Assembly with minimum void difference.

(b) Assembly with maximum void difference.

Figure 2. Peach Bottom Unit 3, Cycle 18, beginning of cycle void difference when ANTS outer iteration stopping criteria is decreased from $1 \cdot 10^{-5}$ to $1 \cdot 10^{-6}$.

In addition to reviewing the number of outer iterations, the solver tolerance, which is defined for the lateral momentum equation residual, was also reviewed. While making changes to the numerical controls, it was found that modifications that were not expected to have an impact on the final solution were leading to seemingly random differences in the solution throughout the core. The randomness indicated that the solution was perhaps not as tightly converged as needed. To test this, the PB3C18 case was run for a single state point at the beginning of the cycle using two different configurations of ANTS. In the first configuration, the outer iteration residual stopping criteria was set to $1 \cdot 10^{-5}$ and in the second, it was set to $1 \cdot 10^{-6}$. Note that it was observed that a larger number of planes failed to converge with the decreased tolerance, but residuals overall were driven to smaller numbers. Fig. 2 shows the two assemblies where the minimum and maximum differences in void occur between these two cases. As the figure indicates, void differences are near zero for most of the subchannels (light blue cells are of the order of 10^{-5}); however, a small set of subchannels will see differences as high as 0.6% or as low as 0.3% void. For a second state point in the depletion simulation, the differences grow larger to as much as 1.6% void difference, which indicates that a tighter outer iteration stopping criteria may be required.

3.1.8 ASSESSMENT OF CHANGES

After implementing all changes that were found to improve ANTS stability, an assessment was done to ensure the new changes did not degrade predictions. The numerical controls were summarized in Table 2. Recall that Item 3 was not used in the final version of ANTS because it was previously found to cause a significant degradation of the validation cases. Item 4 was also not used because it was found to not be sufficient in preventing liquid depletion. Item 7 was already assessed in Section 3.1.7.

Item 1 was assessed by running a standalone ANTS simulation of PB3C18 with the 3D power distribution from a VERA simulation read as input. This was run once with the mass flux limiter set to $100 \text{ kg m}^{-2} \text{ s}^{-1}$ and again with the limiter dropped to $10 \text{ kg m}^{-2} \text{ s}^{-1}$, which is the default value. The root mean square (RMS), minimum, and maximum differences in void fraction between the two cases are $7.3 \cdot 10^{-6}$, -0.2% , and 0.2% , respectively. Note that loosening this limiter is most likely going to impact the ability of ANTS to converge as tightly. It was observed that the same assemblies had convergence issues in both cases, but that the case with the loose tolerance had more oscillation in the solution. Overall, the change in results is insignificant for this limiter.

Item 2 was assessed by running the PB3C18 case multiple times using different values for the crossflow

Table 3. Results sensitivity study to the crossflow turbulent mixing limiter for Peach Bottom Unit 3, Cycle 18

Limiters compared	RMS	State 1		RMS	State 2	
		Min	Max		Min	Max
0.001–0.01	4.1%	-31.1%	14.2%	4.1%	-32.6%	14.8%
0.01–0.1	1.3%	-12.5%	11.3%	1.3%	-15.1%	11.4%
0.1–0.5	0.02%	-3.4%	2.6%	0.12%	-23.6%	8.8%

Table 4. Results for the ANTS validation cases for different values of the crossflow turbulent mixing limiter. Only bundle cases measuring void are considered and all results shown are in units of percent void.

Case	Metric	Limiter=1.0	Limiter=0.1
BFBT	ME	1.92	1.92
BFBT	RMSE	5.75	5.75
PSBT Series 5	ME	-0.09	-0.09
PSBT Series 6	ME	0.04	0.03
PSBT Series 7	ME	4.2	4.2
FRIGG	ME	2.6	2.6
FRIGG	RMSE	3.6	3.6

turbulent mixing limiter for each case. These limiter values included: 0.001, 0.01, 0.1, 0.5, and 1.0. Table 3 presents statistics for the differences in volume fraction prediction between the different cases. For each comparison, the RMS, minimum, and maximum differences are shown for each of the states. Note that the 1.0 case did not converge and is therefore not included in the comparison. The results show that differences in results are very large when the limiter is very small and changed. Approaching a limiter of 50% of the axial flow leads to much smaller changes in the results with changes in the limiter.

In addition to the PB3C18 case, the ANTS validation cases were also run with different values of the turbulent mixing limiter. The first set of results were generated with the limiter set to 1.0 and the second set was generated with it set to 0.1. Note that there were no convergence issues with the limiter set to 1.0 for the validation cases. Results are summarized in Table 4. Only bundle cases are compared because the mixing limiter will have no impact on single flow channel geometry. The statistics shown in the results include mean error (ME) and root mean square of error (RMSE). The results indicate that the solution is almost identical for the two different limiters. Based on these results, it is recommended that the limiter be set to 0.1 because RMS changes are well below 1% void for the PB3C18 case when further increasing to 0.5 and this value should also provide sufficient margin to the case that fails to converge. Furthermore, a value of 0.1 has no significant impact over using a less restrictive value of 1.0 on validation cases .

Items 5 and 6 were assessed together. Results were compared against a case with the Item 1 limiter applied. Note that the simulation for Items 5 and 6 used the tighter outer iteration controls that were discussed in Section 3.1.7, while the base case using the Item 1 limiter used the looser tolerance. Table 5 shows the differences from the Item 1 assessment results. The liquid crossflow presented in the table is the net liquid mass entering the subchannel due to crossflow effects. For reference, the axial liquid mass flow rate in the locations where the minimum and maximum crossflow differences occur are on the order of 0.05–0.1 kg s⁻¹. These results show that the impact of the evaporation and axial mass flow rate controls are negligible for the

Table 5. Change in results after adding the crossflow ceiling numerical control.

Dataset	State 1		State 2	
	Min	Max	Min	Max
Void	$-1.3 \cdot 10^{-3}$	$1.6 \cdot 10^{-3}$	$-1.7 \cdot 10^{-4}$	$1.8 \cdot 10^{-4}$
Pressure [Pa]	$-1.8e-1$	$6 \cdot 10^{-2}$	$-1.2 \cdot 10^{-2}$	$4.1 \cdot 10^{-2}$
Liquid crossflow [kg/s]	$-3.0 \cdot 10^{-4}$	$3.1 \cdot 10^{-4}$	$-7.0 \cdot 10^{-6}$	$7.3 \cdot 10^{-6}$

Table 6. Summary of statistics comparing ANTS to experimental results before and after numerical controls (Items 1, 5, and 6) are added

Test	Quantity	Baseline	New controls	Metric	Units
FRIGG	Void	2.4	2.4	ME	%
FRIGG	Void	3.5	3.5	RMSE	%
BFBT	Void	2.1	2.1	ME	%
BFBT	Void	5.5	5.5	RMSE	%
PSBT Series 1	Void	5.8	5.8	AME	%
PSBT Series 2	Void	3.9	3.9	AME	%
PSBT Series 3	Void	5.9	5.9	AME	%
PSBT Series 4	Void	2.6	2.6	AME	%
PSBT Series 5	Void	5.0	5.0	ME	%
PSBT Series 6	Void	6.0	6.0	ME	%
PSBT Series 7	Void	8.4	8.4	ME	%
Riso Series 100	ΔP	14.6	14.6	rRMSE	%
Riso Series 300	ΔP	17.4	17.4	rRMSE	%

conditions tested.

Finally, an assessment is performed using the ANTS validation cases to check the impact of the remaining numerical controls on the quality of results. Table 6 presents the summary of changes in statistics when considering Items 1, 5, and 6 from Table 2 for the validation tests that have been used for ANTS assessment. Additional metrics presented in this table include absolute mean error (AME) and relative root mean square of error (rRMSE). Items 1, 5, and 6 are the controls that are expected to have no impact on the final solution. As the results indicate, no significant change to the statistics were observed when implementing these controls.

No assessment was done on Item 7 for the validation tests, as the impact of tightening the tolerance was already shown in Section 3.1.7.

3.2 PERFORMANCE IMPROVEMENTS

The ANTS solve time was relatively fast in comparison to the MPACT portion of the VERA solve time; however, a collection of performance bottlenecks were identified that were straightforward to address, leading to a significant improvement in the runtime. Furthermore, the added numerical constraints may require the ANTS solution to take more iterations than it previously had, so investing some time in these performance improvements helps to limit the TH solve time to a small percentage of total VERA runtime. The following changes were made to improve ANTS solve time:

- Add the Portable Extensible Toolkit for Scientific Applications (PETSc) as an option for solving the crossflow Jacobian matrix at each level
- Switch from directly accessing CTF solution data (stored in derived type objects stored in CTF) to simple arrays defined in the ANTS derived type
- Perform updates of solution terms at select points in the solution instead of on the fly to reduce computational expense
- Perform matrix-math operations where possible instead of element-by-element calculations

These performance improvements were identified by inserting timing calls into the ANTS solution and by using a profiling tool to target specific lines of code that were consuming a significant amount of solve time. Using PETSc for the matrix solve is one of the more significant improvements. The biconjugate gradient stabilized (BiCGStab) solver was used for performing the matrix solve. Despite the matrix appearing to be sparse, it was found that setting it up so that each row contains seven terms (the main gap plus the gaps that touch each of the channels connected to the main gap) leads to the solution becoming unstable. For this reason, the matrix is built as if it were dense, which still leads to a significant speedup over the native direct inversion that was being done previously.

The performance of ANTS was tested by running a single state point of PB3C18. Note that the following timing results were produced with a version of ANTS that *does not* have the outer iteration controls applied (Item 7 from Table 2. This is done to produce a baseline set of timing data. The state point has a depletion of about 16.7 GWd/MTU, a relative core power of 82%, and a relative core flow of 81%. The 3D power distribution was generated from a coupled simulation, which was then used to run ANTS as a standalone solve. Note that CTF drives the solution and performs a pressure balance iteration to adjust assembly inlet flows to balance core pressure drop. A total of seven pressure balance iterations are performed, meaning that ANTS is solved to steady state seven times in this simulation. The simulation was performed on the Sawtooth high performance computing (HPC) system using one processor per assembly. The original ANTS simulation takes 145 s, whereas the one with performance improvements takes 50 s: a 2.9×speedup. A breakdown of time spent in different parts of the solution is shown in Fig. 3 for the case after performance improvements were made.

Whereas the pressure matrix solve was close to 50% of the simulation time prior to the performance improvements, it now comprises 14% of the solve time. 17% of the solve time is spent in CTF, which includes performing the fuel rod solution, calculating heat transfer coefficients, and driving the BWR pressure balance iteration loop. Note that the initialization time was removed from the timing results because that price is paid only during the first solve of CTF/ANTS in a coupled simulation, which will ultimately involve hundreds of solves.

To further test the performance, a fully coupled simulation was performed of three depletion points of PB3C18 using the improved version of ANTS. The solver used hybrid STH, meaning the first 5 TH solutions were done by STH before switching to ANTS. CTF was called a total of 33 times, and a total of 32 ANTS solutions were performed. Recall that some calls to CTF require multiple ANTS solves due to the BWR pressure balance iteration loop. Furthermore, a feature was employed that shuts off the TH solve when power is changing less than a threshold amount (1 W cm^{-1} in the case of this simulation), which means that some calls to CTF will involve no ANTS solves and will exit immediately. Total wall time spent in the calls to CTF for this case comprised 3.4% of the total VERA walltime, making the TH solve portion a fairly insignificant computational burden.

The PB3C18 case was run again with the standard $1 \cdot 10^{-5}$ outer iteration tolerance and another time with the tighter $1 \cdot 10^{-6}$ outer iteration tolerance. Only two states were run for these cases. The intention with this set




Figure 3. Profile of time spent in different sections of CTF/ANTS for a full-core, steady-state solution of a depletion point of Peach Bottom Unit 3, Cycle 18 with 3D power distribution after performance ANTS improvements.

of runs was to determine the effect of tightening the iteration controls. The loose tolerance case performs 17 coupled iterations in each state and overall, 4.1% of the runtime is spent in ANTS. The tight tolerance case performs 17 coupled iterations in the first state and 16 coupled iterations in the second state. While the number of coupled iterations did decrease in the second state, the portion of runtime spent in ANTS increases to 7.4% of the total runtime. Note that the hybrid coupling was used for both of these cases, meaning that the first five coupled iteration of each state was performed using STH, and the solve time spent in STH is not included in the ANTS runtimes given above.

3.3 USER INTERFACE IMPROVEMENTS

3.3.1 OUTPUT EDITS

The following output edits have been added to the HDF5 file to provide more simulation data to the user:

- liquid velocity
- vapor velocity
- liquid mass flow rate
- vapor mass flow rate
- liquid enthalpy
- liquid mass crossflow
- vapor mass crossflow

The liquid and vapor velocity terms are calculated using the relationship $\dot{m}_k = \alpha_k \rho_k V_k A$, where \dot{m} is the mass flow rate and A is the subchannel flow area. The k subscript denotes the phase, which will be either liquid or vapor. The liquid mass flow rate, vapor mass flow rate, and liquid enthalpy are all solved quantities, so they are simply printed out after converting to the VERA output file units. The liquid mass crossflow and vapor mass crossflow terms are defined as the net liquid and vapor mass flow rate entering the subchannel due to all crossflow effects in each of the gaps that connect to that subchannel. This term may be negative if there is a net removal of mass from the subchannel due to crossflow. These additional metrics can be useful for analyzing BWR simulation results and can be enabled via an input file option.

3.3.2 CONVERGENCE EDITS

Previously, only the root processor of the parallel simulation was writing its convergence information to the log file, meaning that only one assembly in the core was being printed. To help in diagnosing convergence issues, this capability was expanded to provide more information to the user. This first required doing a reduction of selected convergence information to the root process for each level of the model. When this reduction is complete, the root process has information on all assemblies for writing output to the log file. Upon completion of a planar solve, the root process will write to the log file that the planar solve is completed and that either all assemblies have reached their solver tolerances or there have been some failures. If there are failures, the following information will be written for each failing assembly:

- the assembly experiencing the failure to meet its tolerance,
- the final residual,
- the solver tolerance,
- the maximum change in volume fraction between the last two successive outer iterations,
- the maximum change in pressure between the last two successive outer iterations,

- the maximum difference between minimum and maximum pressure in the level.

Note that if a particular assembly does not reach its solver tolerance by the user-specified number of maximum outer iterations, ANTS will still move on to the next level. The new metrics that are printed to the log file will give the user an indication of how much the volume fraction and pressure are oscillating. If these values are relatively small, it may be suitable to accept the solution. The final metric is printed to indicate the uniformity of the planar pressure distribution.

In addition to this output, a more detailed option is added to write a separate log file for every assembly in the model. The name of each log file will include the assembly number to help with identifying the source of the information. These log files will include:

- the final level residual for every outer iteration that was taken,
- the individual residuals for every gap in the level at the last outer iteration,
- the crossflow mass flux change for every gap in the level at every outer iteration.

This information can help to show if the solution is oscillating and the severity of the oscillation if it exists. Note that this information is only printed if there is a convergence failure for that level. Also, a unique solve ID is printed to both the main log file and the assembly log file so that the user can find assembly-specific information based on reduced information that can be found in the main log file.

4. RESULTS

To test the new capabilities, H1C1 was used [15]. This builds on a previous assessment of VERA BWR simulation capabilities performed in 2021. H1C1 provides transverse in-core probe (TIP) detector data to compare against, enabling quantification of the accuracy improvements for the various models discussed in this report. The cycle has three fuel types, all using 7×7 fuel without water rods. These fuel types and a quarter-core 2D slice from VERA are shown in Figure 4. The model was approximated to be quarter symmetric, and the operational history was approximated with 40 state points. Twenty of those state points have TIP detector measurements; the parameter associated with the bypass flow are summarized in 7 for states with TIP measurements.

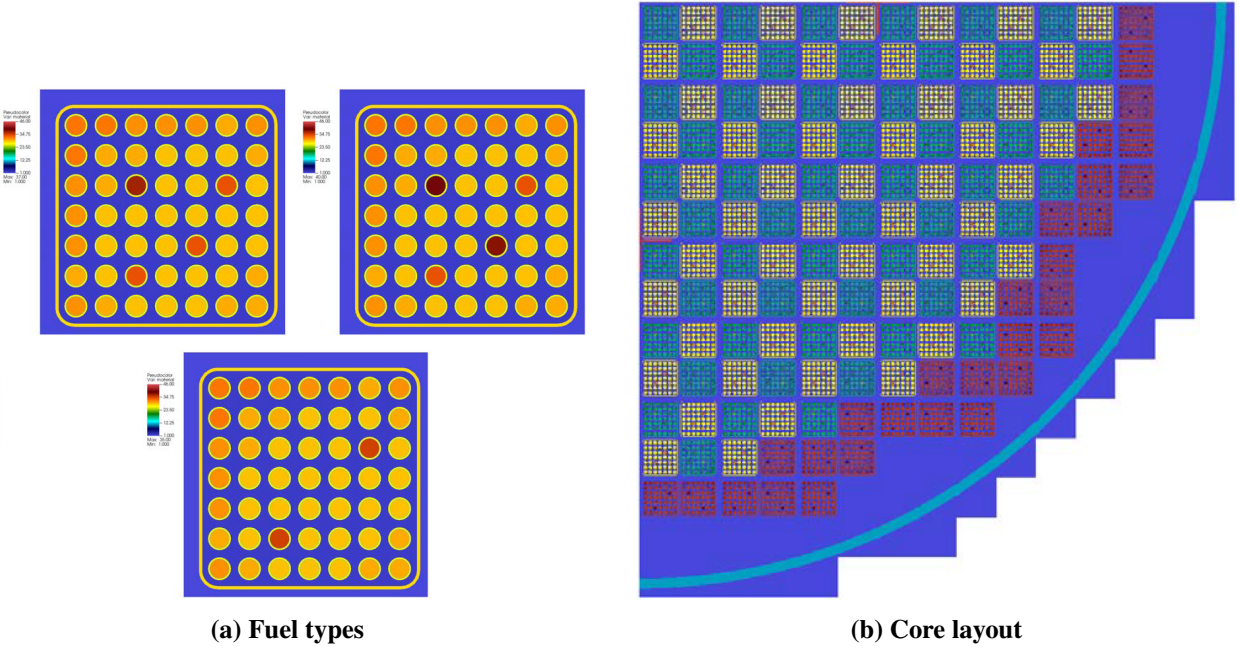


Figure 4. VERA rendering of H1C1 geometry

This section presents results for this cycle at various levels of fidelity. First, comparisons are briefly shown using fixed TH conditions in the bypass. Inlet conditions, saturated conditions, and outlet conditions are compared. This comparison shows that although some improvements can be achieved by selecting an improved set of values, none of these options resolve the radial and axial effects of the bypass heating on the power. Section 4.2 presents more detailed results of calculations using the bypass heating model. These results should be significantly better since a physics-based model is being used to capture the bypass flow feedback instead of the ad hoc conditions discussed in Section 4.1.

Table 7. Summary of key state points for H1C1

State Index	Exposure GWD/MTU	Power (% Rated)	Flow (% Rated)	Bypass (% Flow)
5	0.702176	50	43.95	3.27
7	1.28199	89.86	86.62	8.01
10	2.58383	95.69	100.38	9.54
11	3.11625	86.12	84.97	7.82
12	3.64647	85.84	77.58	7
14	4.15794	79.93	81.91	3.9265
16	4.31998	76.07	100	4.992
18	4.68595	78.57	94.39	4.66157
20	5.02436	80.54	90.96	4.45954
22	5.79378	86.37	93.63	4.61681
24	6.59296	82.96	79.49	3.78396
25	6.97987	93.14	95.92	4.75169
26	7.16175	93.14	95.92	4.75169
29	8.06014	87.4	96.56	4.78938
31	8.93869	83.62	100	4.992
33	9.40277	91.58	99.62	4.96962
34	9.82936	87.48	100.25	5.00672
35	9.88338	88.38	90.83	4.45189
37	10.1193	90.64	100.13	4.99966
39	10.3111	86.78	98.34	4.89423

4.1 FIXED TH CONDITIONS

Table 8 summarizes the TIP comparisons for each of the three fixed bypass TH conditions, while Figure 5 shows the errors across all the state points. The saturated conditions clearly performs the best, which is expected. The outlet conditions performed quite poorly, which is not surprising since the bypass flow should only void near the top of the core, not along the entire height of the core. Additionally, the calculation was far less stable when using outlet conditions, so the calculation made it just over halfway through the cycle. Because of the poor performance and accuracy, the outlet conditions were abandoned instead of spending computational resources to finish the calculation.

Table 8. RMS of TIP comparisons for all state points for fixed bypass flow TH conditions and STH

Condition	Radial		Axial		3D	
	RMS	Max	RMS	Max	RMS	Max
Inlet	4.96%	11.61%	19.18%	35.1%	20.58%	57.21%
Saturated	4.87%	15.17%	16.92%	11.83%	26.38%	49.23%
Outlet (11/20 TIP states)	3.69%	40.72%	41.62%	8.78%	104.08%	127.70%

While the saturated conditions showed some improvements over the inlet conditions, the errors are still higher than they should be for a high fidelity multiphysics solution. Specifically, we see improvement in the axial and 3D maximum errors for the saturated conditions compared to the inlet conditions. However, it is clear that the axial power shape is not properly resolved by the saturated conditions, necessitating the use of the bypass heating model and ANTS TH solver for this case to achieve greater accuracy. This is discussed




Figure 5. TIP comparisons for all state points for fixed bypass flow TH conditions with STH

in the next section.

4.2 BYPASS HEATING MODEL

This section compares the fixed saturated conditions with the bypass heating model. It also compares the heating model with STH to the heating model with ANTS. Figure 6 shows k_{eff} for every state. For the first third of the cycle, the trends are very similar between the three treatments. Later in the cycle, the saturated conditions show improvement and are generally within 500 pcm of critical. This is not surprising since saturated conditions are likely a decent average of the bypass conditions and k_{eff} is a global quantity. The bypass heating underpredicts k_{eff} for both STH and ANTS, indicating that some improvements may be needed to more accurately predict the reactivity.




Figure 6. k_{eff} for each TH treatment for H1C1

The core pressure drop is shown in Figure 7. There is a significant difference between STH and ANTS for most state points, often around 5 psi. This is not surprising because of the increased fidelity and greater number of closures in ANTS, but it does demonstrate a notable shortcoming of STH. Figure 8 shows the converged bypass inlet orifice loss coefficient and bypass flow for ANTS. As the bypass flow increases, the loss coefficient decreases accordingly to ensure pressure balance. The trends for the STH were identical, just with different magnitudes since the STH core pressure drop was calculated to be lower than for ANTS. Finally, Figure 9 shows the maximum bypass void for each state point. The values are nearly identical between STH and ANTS and range from just above 1% to about 30%. The higher values are likely too high for this core and represent room for improvement in the bypass heating model, as evidenced by the k_{eff} results being too low. Specifically, it is expected that a user-specified heating fraction is the most significant limitation; direct calculation of the power deposition in the bypass flow should be performed instead and would likely reduce the bypass void for many of these state points.

Figure 10 shows the exit void distribution for the bypass flow regions at 5.79378 GWD/MTU. The distribution is shown in quarter symmetry since that is how H1C1 is modeled. In the core interior, the lower exit voids correspond to locations that more fully inserted control blades. The highest exit voids occur away from




Figure 7. Pressure drop for bypass heating treatments for H1C1

deeply inserted control blades and away from the core periphery, in the highest power fuel bundles. The core periphery has the lowest bypass void, which is expected since peripheral fuel bundles are typically the lowest power. All features of the exit void generally follow the expected behavior.




Figure 8. Converged bypass loss coefficients for bypass heating treatments for H1C1




Figure 9. Maximum bypass void fraction for bypass heating treatments for for H1C1

0.216407	0.160648	0.166862	0.229293	0.23025	0.171161	0.173568	0.249593	0.24928	0.186217	0.154145	0.158326	0.0798189
0.176408	0.21461	0.226751	0.216021	0.224155	0.235256	0.246613	0.244884	0.245903	0.241776	0.213927	0.170293	0.0904776
0.175847	0.217615	0.223326	0.215641	0.2179	0.241357	0.244545	0.244965	0.244214	0.247395	0.222337	0.172754	0.0878771
0.22085	0.136296	0.144384	0.23526	0.241374	0.155875	0.164748	0.250519	0.25904	0.212074	0.1819	0.174647	0.0743654
0.220686	0.139415	0.150267	0.243871	0.239443	0.158922	0.168898	0.25778	0.261559	0.207685	0.171993	0.14534	0.0480931
0.188139	0.225856	0.234095	0.232596	0.235577	0.255785	0.260757	0.260633	0.248877	0.234755	0.183527	0.0418322	
0.193681	0.231121	0.240532	0.237479	0.243915	0.26626	0.266575	0.261057	0.249304	0.219933	0.165478		
0.236063	0.171818	0.185449	0.256802	0.263962	0.229654	0.231945	0.252225	0.245555	0.20763	0.119938		
0.236874	0.172168	0.183936	0.256071	0.266166	0.221573	0.21258	0.237795	0.216651	0.169048	0.0727022		
0.196973	0.230555	0.240854	0.250298	0.251295	0.236414	0.218544	0.207555	0.16567	0.0260041			
0.175015	0.223432	0.234516	0.234048	0.227163	0.204613	0.163024	0.117675	0.0731787				
0.162409	0.174989	0.191371	0.194466	0.168583	0.06261							
0.0858552	0.0945649	0.101064	0.0893715	0.0711031								

Figure 10. Bypass exit void distribution at MOC for H1C1

4.3 TIP COMPARISONS

This section shows comparisons to the plant TIP data for the saturated conditions and the bypass heating models with STH and ANTS. Figures 11 and 12 show the RMS and max differences, respectively. As expected, the bypass heating model shows improvement over the saturated conditions for most state points since it is resolving both radial and axial effects in the bypass region. There is generally little difference between STH and ANTS with the bypass model on. The TIP detectors are in the narrow-narrow gaps and most influenced by the 4 fuel bundles surrounding that location; if the assembly powers are relatively accurate with STH then the pin-wise TH resolution of ANTS is not expected to significantly change the results.




Figure 11. RMS differences in TIP measurement comparisons for each TH treatment for H1C1




Figure 12. Max differences in TIP measurement comparisons for each TH treatment for H1C1

4.3.1 TIP MEASUREMENT 1 – 0.702176 GWD/MTU

The first TIP measurement occurs at 0.702176 GWD/MTU, close to beginning of cycle (BOC). In 11 this measurement looks an outlier, as the saturated conditions performed far better than the bypass heating model. That measurement corresponds to state 5 in Table 7, which shows that state 5 has the lowest bypass flow of any state with TIP data. Additionally, the state is only 50% power. The bypass model predicts void up to 24% in the bypass flow; it is likely that for these conditions the bypass model is overpredicting the power deposition in the bypass flow, and thus overpredicting the bypass void. This is confirmed by Figures 13 through 15, where the bypass heating model shifts far too much of the power toward the bottom of the core. This represents room for improvement to the bypass model.

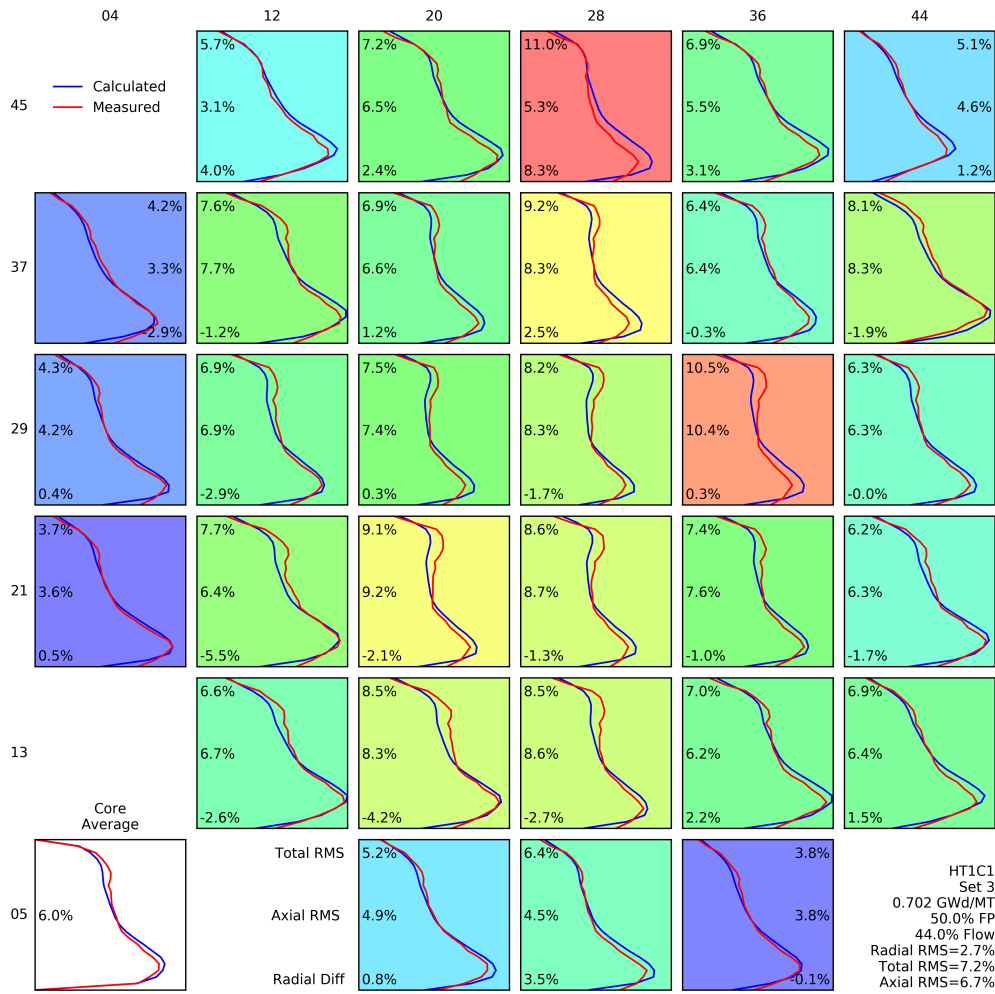


Figure 13. TIP comparison map for saturated conditions at 0.702176 GWD/MTU

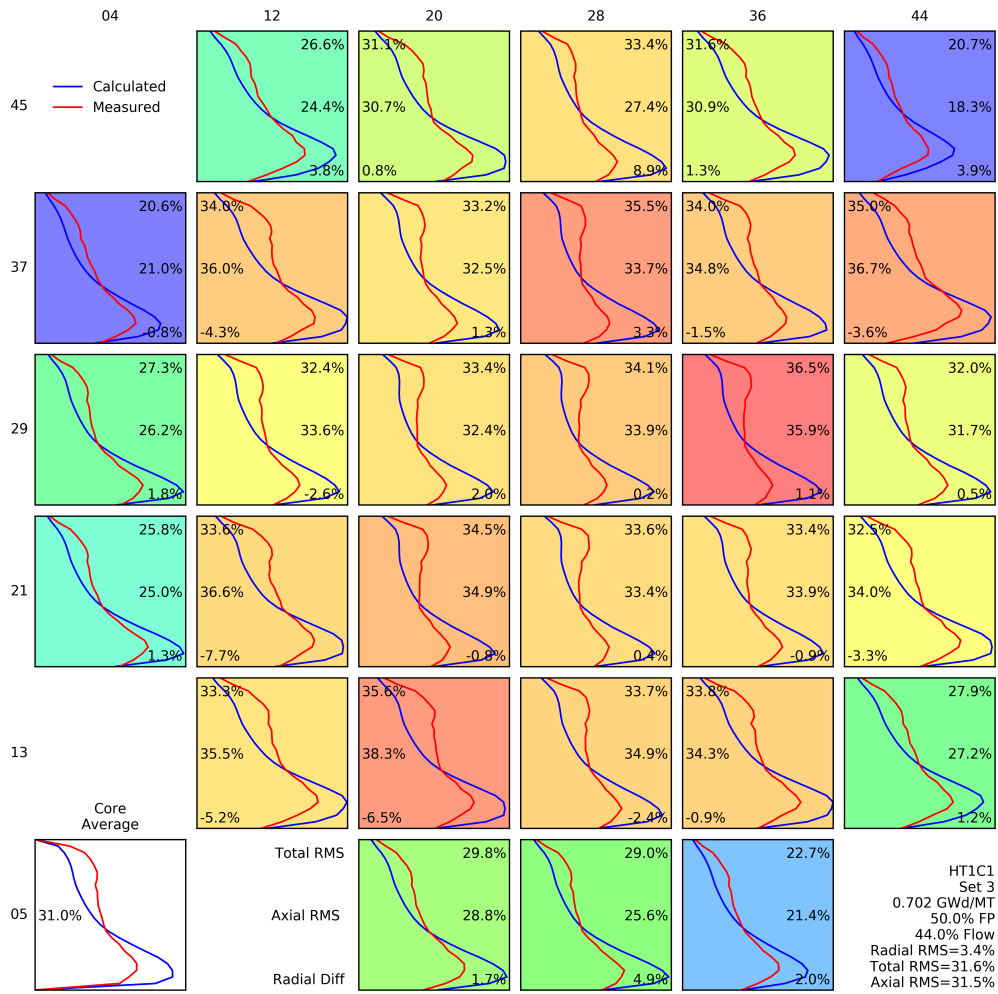


Figure 14. TIP comparison map for STH+heating conditions at 0.702176 GWD/MTU

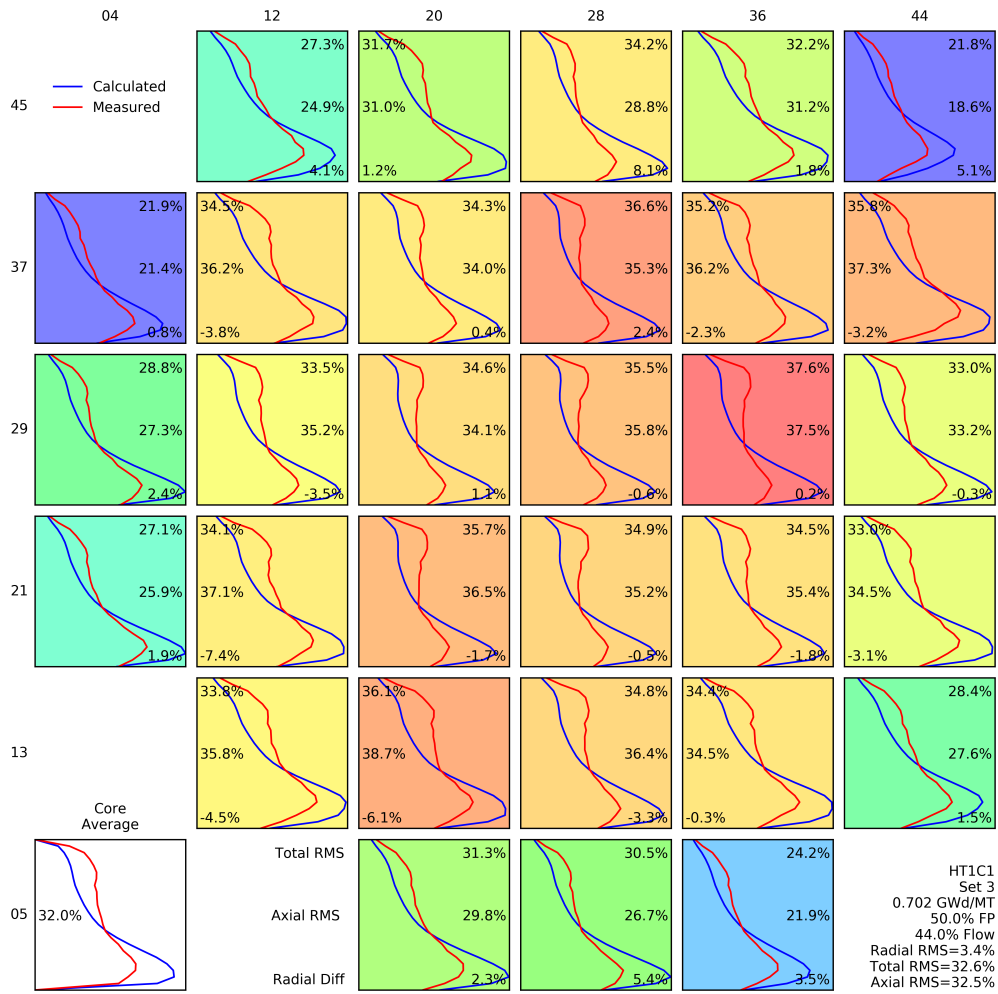


Figure 15. TIP comparison map for ANTS+heating conditions at 0.702176 GWD/MTU

4.3.2 TIP MEASUREMENT 2 – 1.28199 GWD/MTU

The second TIP measurement occurs at 1.28199 GWD/MTU, which is still fairly close to BOC. This state point is much closer to typical operation, with about 90% power, 87% flow, and 8% bypass. Figures 11 and 12 showed significantly better results for this state point for all the models. The saturated conditions were already good, so the bypass heating model shows only minor improvements at the core level. Figures 16 through 18 show the TIP maps for this state point. In contrast with the previous state, the bypass heating model predicts the power shape well for this state. Several detector locations show that the heating model shifts a little more power toward the bottom of the core than the saturated conditions, which in this case is an improvement on the results compared to the TIP measurements.

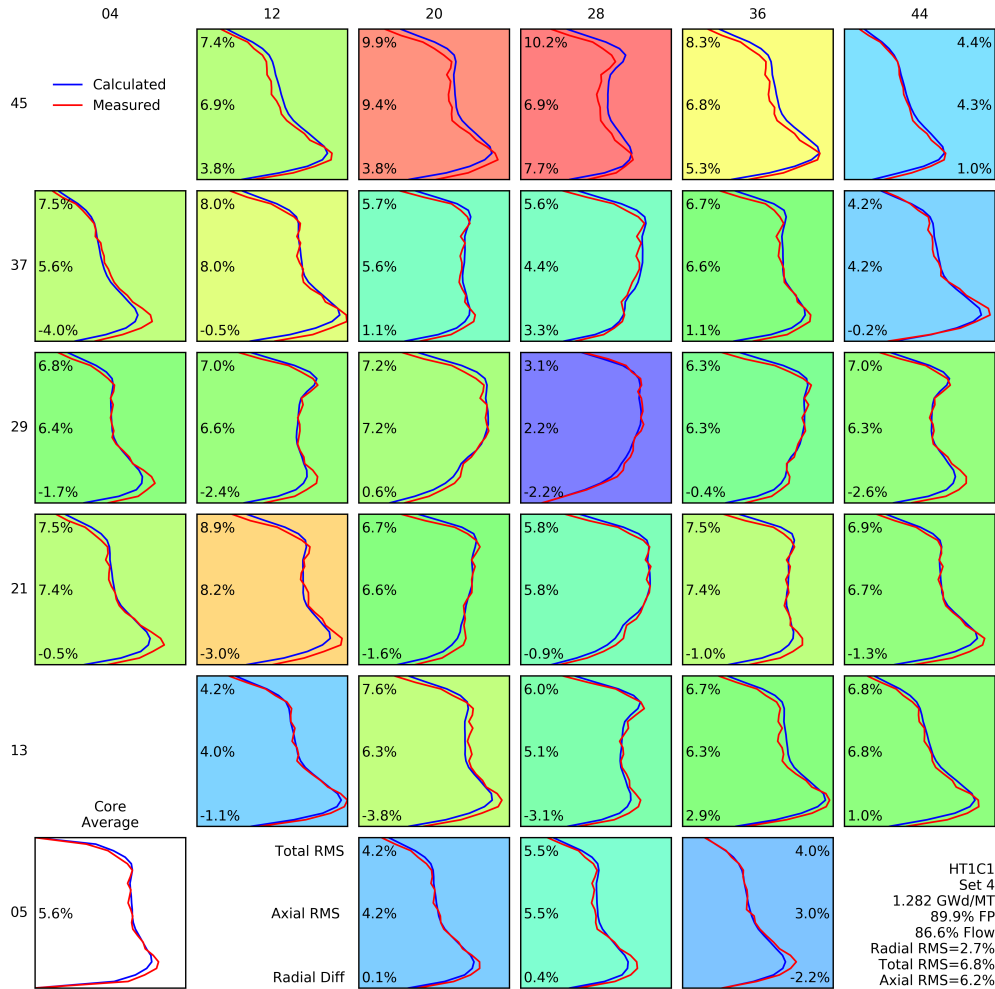


Figure 16. TIP comparison map for saturated conditions at 1.28199 GWD/MTU

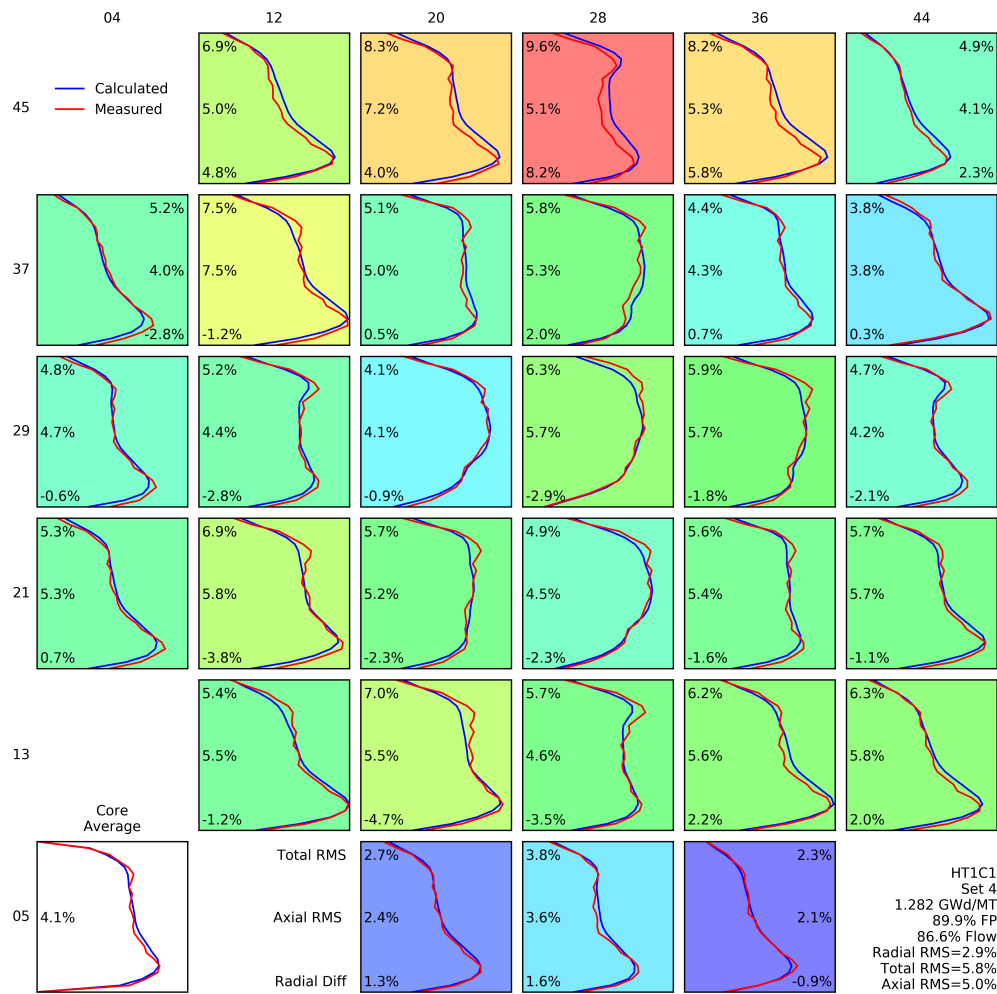


Figure 17. TIP comparison map for STH+heating conditions at 1.28199 GWD/MTU

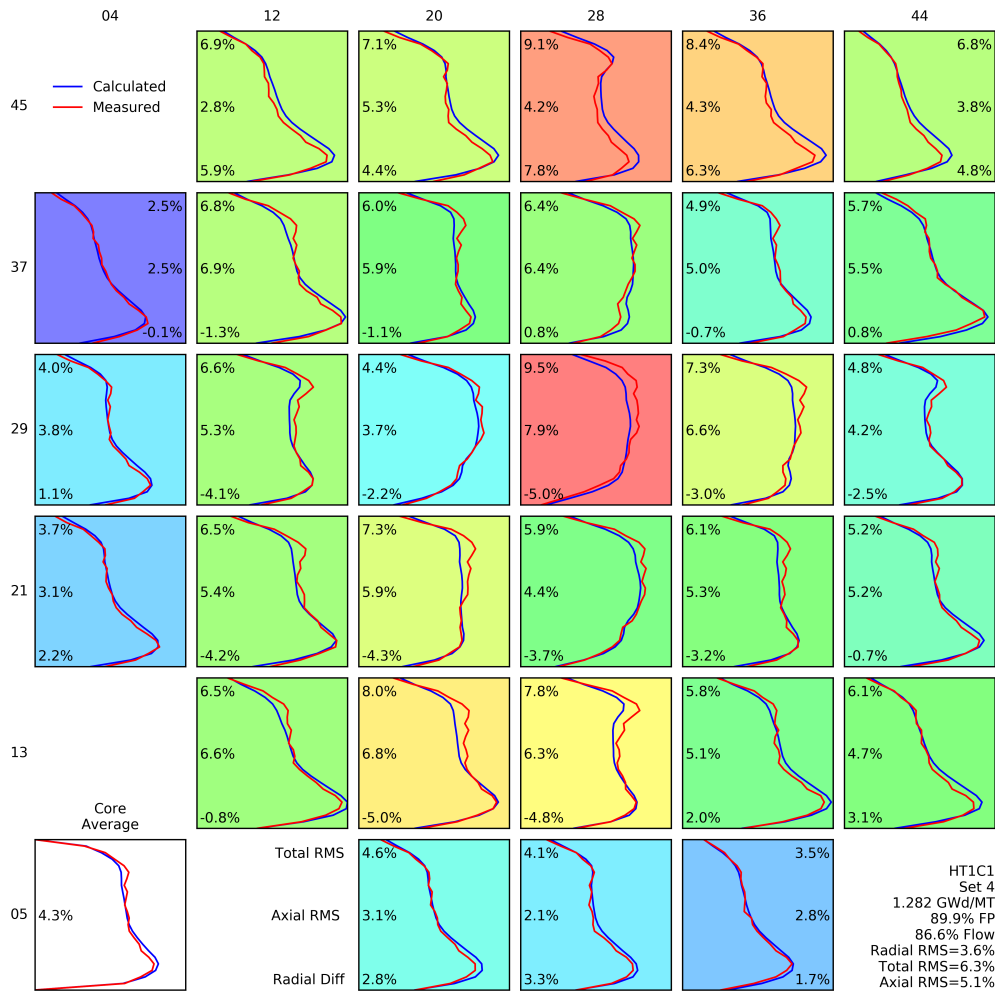


Figure 18. TIP comparison map for ANTS+heating conditions at 1.28199 GWD/MTU

4.3.3 TIP MEASUREMENT 10 – 5.79378 GWD/MTU

The next TIP measurement presented here occurred at 5.79378 GWD/MTU. This is approximately middle of cycle (MOC). For this case, as with several others around MOC, the saturated conditions do not perform as well as at BOC and MOC. The bypass heating, on the other hand, performs fairly consistently throughout the cycle. Figures 19 through 21 show the detailed TIP maps for this state point. There are a few locations where the bypass model is still slightly underpredicting the power in the lower half of the core, but it is far more accurate than the cosine-like power profile predicted by the saturated conditions. This illustrates the type of improvement expected from the bypass heating model.

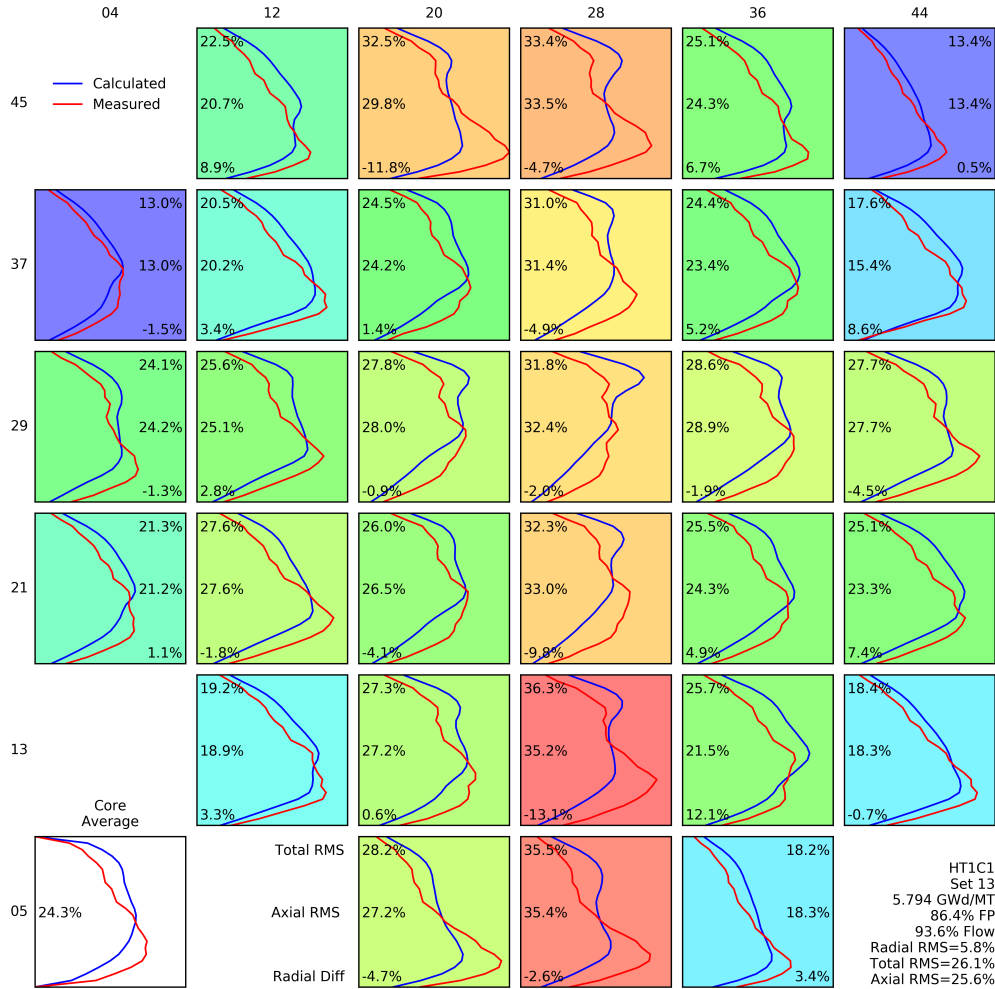


Figure 19. TIP comparison map for saturated conditions at 5.79378 GWD/MTU

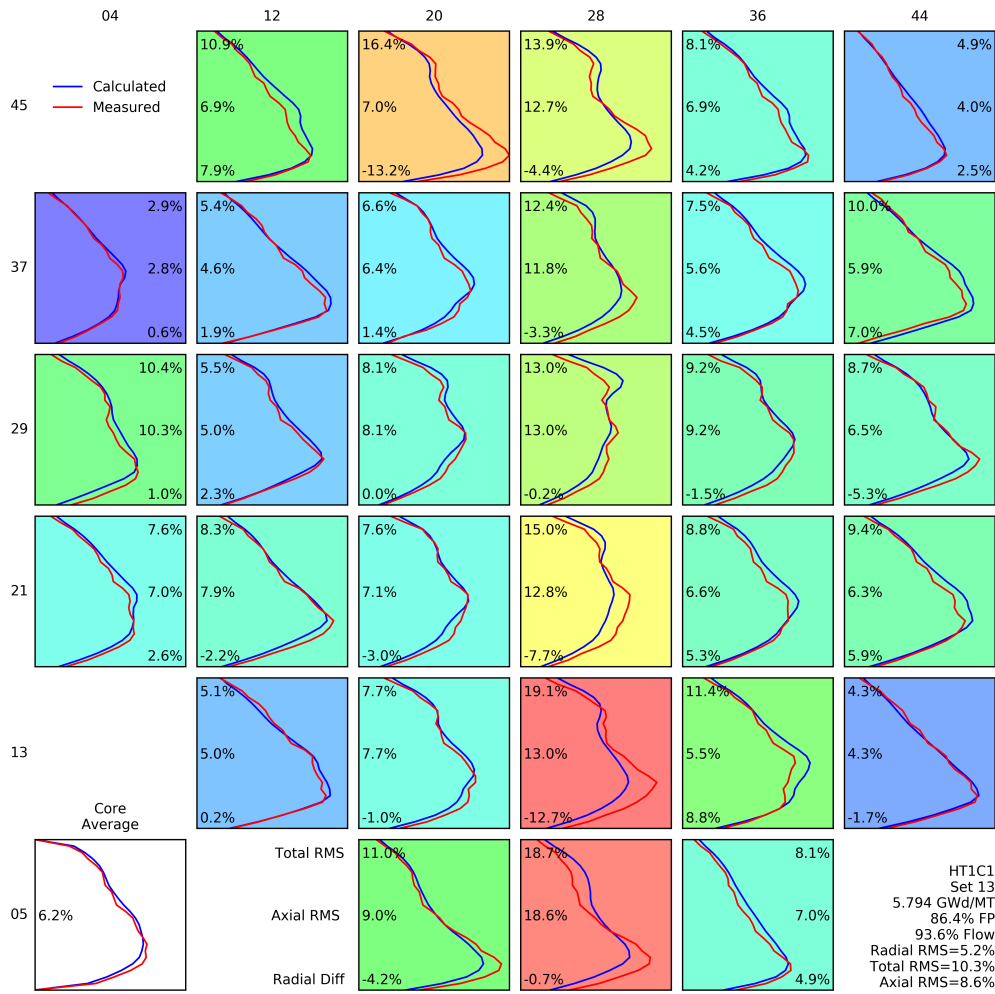


Figure 20. TIP comparison map for STH+heating conditions at 5.79378 GWD/MTU

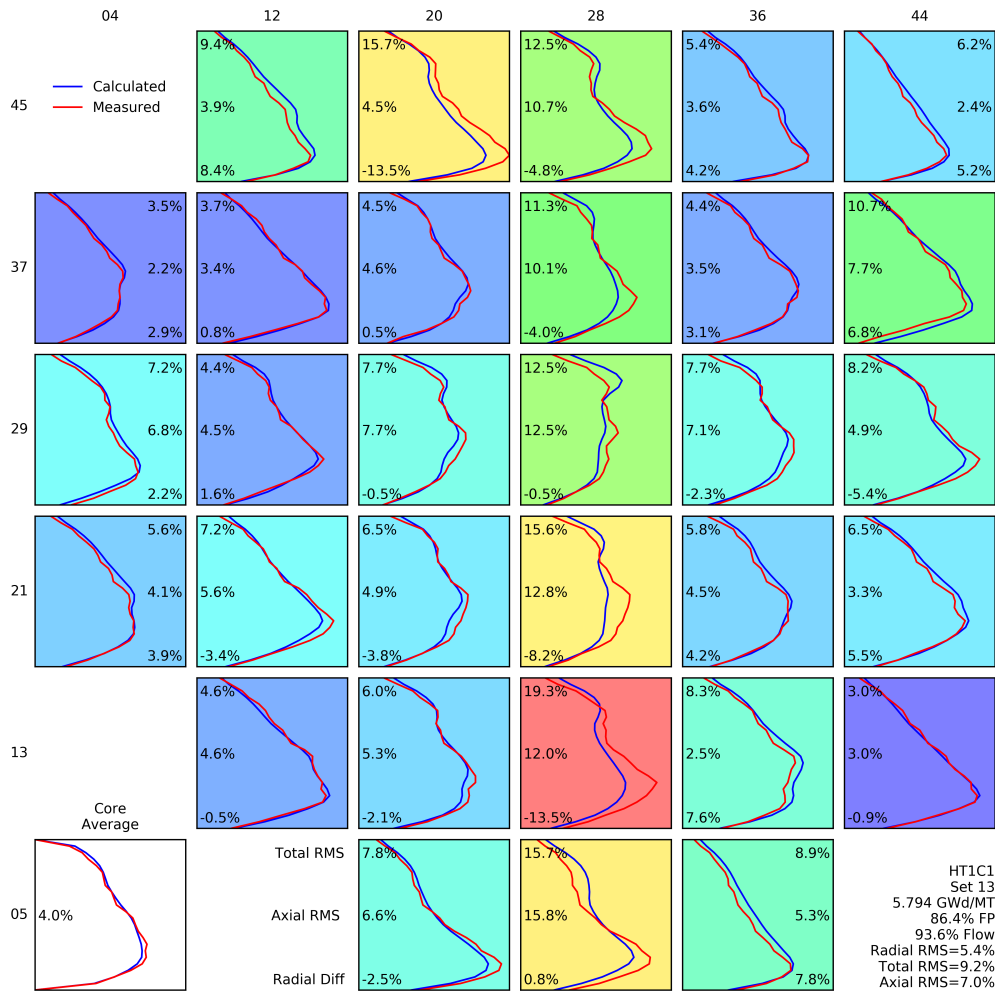


Figure 21. TIP comparison map for ANTS+heating conditions at 5.79378 GWD/MTU

4.3.4 TIP MEASUREMENT 20 – 10.3111 GWD/MTU

The final TIP comparisons are for end of cycle (EOC) and are shown in Figures 22–24. The improvements for the bypass heating model are not as large as at MOC because the saturated conditions performed better. This is not surprising because in many cases fuel depletion throughout a cycle has the effect of averaging away errors that occurred early in the cycle. However, the bypass heating model still shows improvements of several percent at every location.

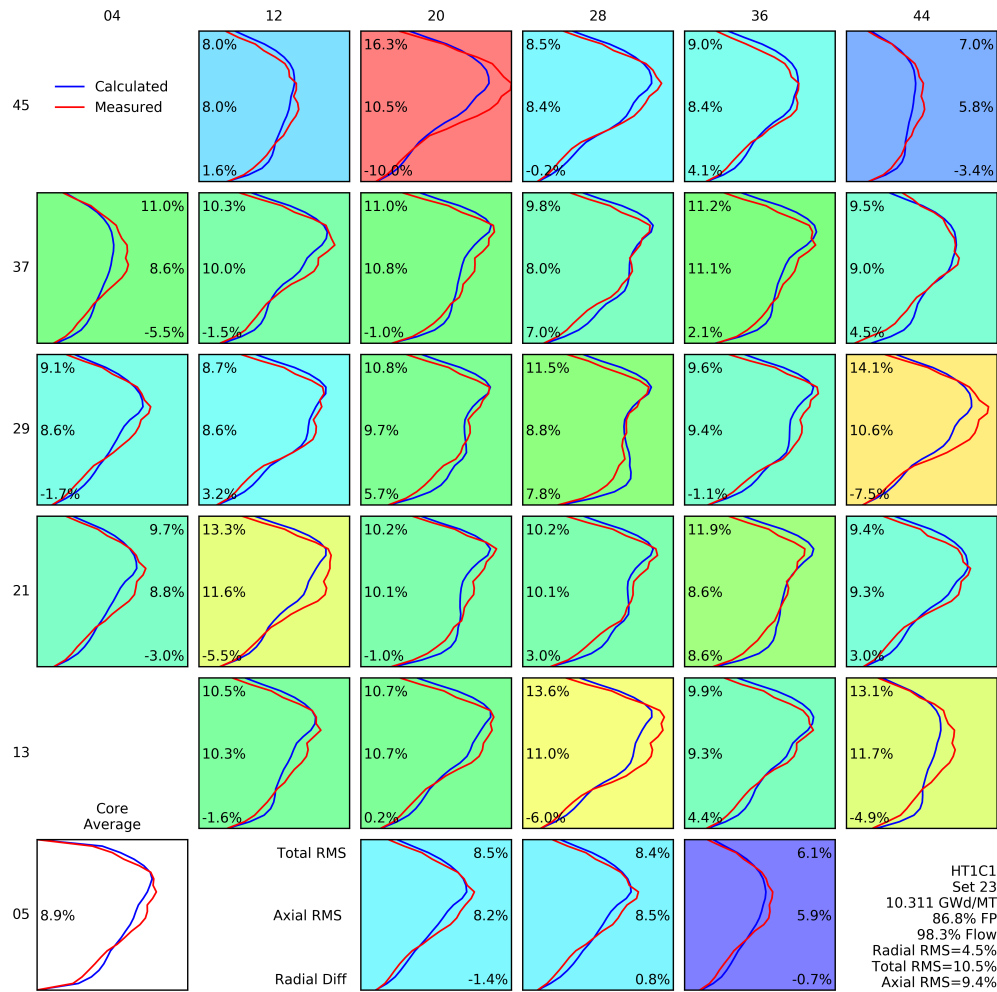


Figure 22. TIP comparison map for saturated conditions at 10.3111 GWD/MTU

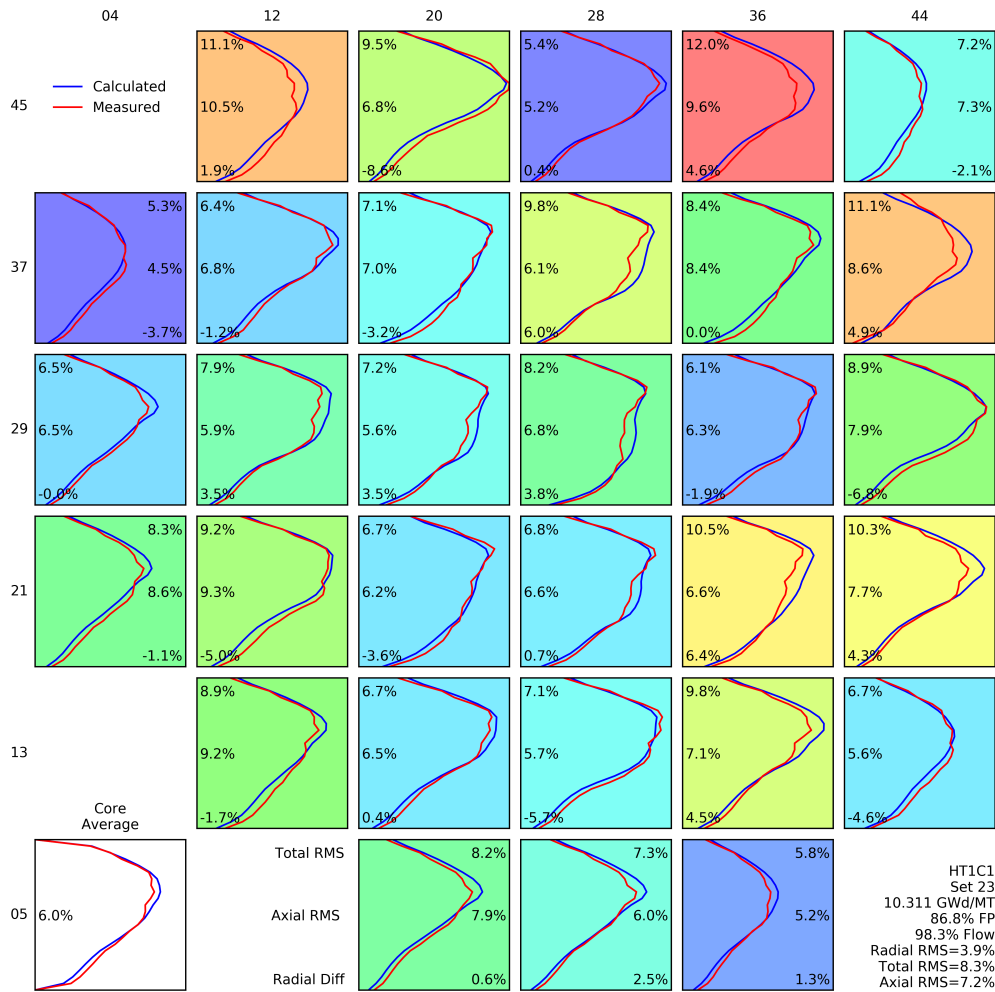


Figure 23. TIP comparison map for STH+heating conditions at 10.3111 GWd/MTU

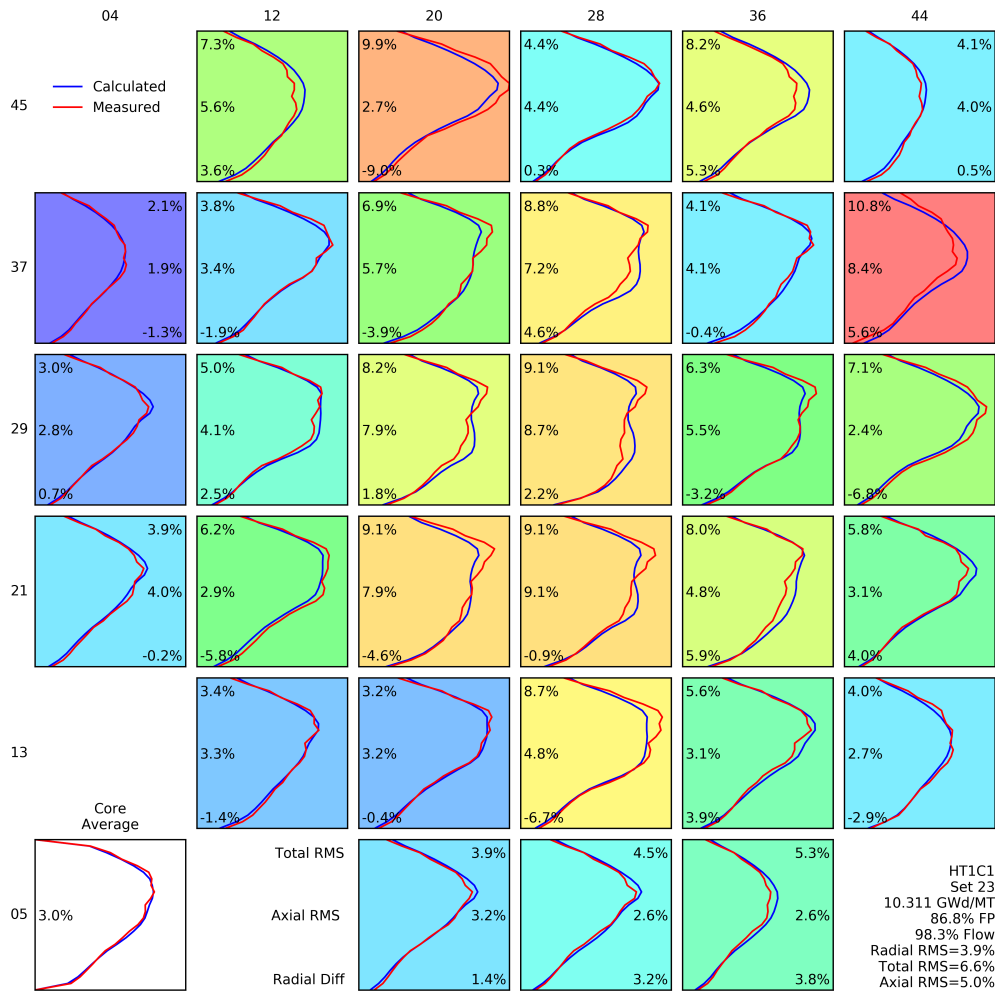


Figure 24. TIP comparison map for ANTS+heating conditions at 10.3111 GWd/MTU

4.4 PERFORMANCE

Each calculation was run on the Idaho National Laboratory (INL) Sawtooth cluster [16] using 960 compute cores. Table 9 summarizes each calculation’s performance. There is a clear trend of increased coupled iterations when enabling the bypass heating option. Similar trends have been observed with VERA previously with all TH coupling disabled and a user-specified axial void distribution. The void distribution makes the 2D/1D calculation in MPACT harder to converge, especially for the 3D coarse mesh finite difference (CMFD) and 1D axial transport calculations. By enabling the bypass heating model, an axial void distribution is introduced in the bypass regions as well, causing a similar effect and slowing convergence. Surprisingly, the number of coupled iterations with ANTS is lower than for STH when the bypass heating model is enabled. However, the expense of each iteration is much higher and therefore the total runtime for H1C1 was still about 33% greater for ANTS than STH. Table 9 also shows the memory usage. Memory usage is not determined until the end of the calculation, so only 4 of the 5 models have the data available. The ANTS model certainly took the most memory; at a maximum of 2.36 GB on a core and only 1.707 GB average, there are no memory concerns since most HPC systems have at least 4 GB per core. The STH+heating model shows slightly less memory usage compared to the inlet and outlet conditions; this is likely because the case had to be run in multiple parts, so there was some small growth in memory usage throughout the longer fixed condition calculations. In general, the bypass heating model should be a very small memory increase over fixed TH conditions.

Table 9. Summary of performance data for each H1C1 simulation

Simulation	Coupled Iterations	Walltime (HH:MM:SS)	Memory per Core (GB)	
			Average	Maximum
Inlet + STH	994	22:27:37	1.411	1.875
Outlet + STH [†]	1546	39:53:02	n/a	n/a
Saturated + STH	1071	24:14:36	1.411	1.887
Bypass heating + STH	1634	61:37:19	1.394	1.831
Bypass heating + ANTS	1224	82:10:49	1.707	2.360

[†]Timing data for completion of 24/40 state points; memory data is unavailable for calculations that do not complete.

Figure 25 shows the number of iterations taken to solve each state point for each of the 5 models. Again, the outlet conditions performed quite poorly. Surprisingly, STH with the bypass heating enabled was the second worst on average, but it was still significantly faster than the ANTS calculation. Table 10 shows the average amount of time spent in each iteration for various parts of the solve for the three most accurate models; some portions of the calculation are excluded, such as depletion and the radial transport, because they require an approximately fixed amount of time per coupled iteration or state point. The two primary culprits for increased runtime are CMFD and the self-shielding calculations. The CMFD calculation becomes harder to converge as void gradients are introduced in the coolant. This has been observed with user-specified axial void distributions without TH coupling enabled. Thus, the bypass heating model increases the difficulty of the CMFD calculation compared to saturated conditions, and ANTS adds further difficulty converging the CMFD calculation. Significant research effort has been spent to reduce CMFD runtime using multigrid approaches; at this point, these approaches have sped up the calculation at the expense of stability and have not been robust enough for whole-core BWR calculations, but further research is underway. The other significant difference is the self-shielding calculation. This calculation is triggered each iteration depending on the maximum change in temperature or density due to TH feedback. The ANTS calculation solves a much higher resolution TH problem, and therefore triggers many more self-shielding calculations than the STH version of the problem. The STH+bypass solve took an average of 14 self-shielding solves per

state point; the other two options took an average of about 27 self-shielding solves per state point. This increased runtime could be implemented by switching to a new self-shielding method recently implemented in VERA [17].




Figure 25. Number of coupled neutronics/TH iterations required to solve each state point for each TH model.

Table 10. Time (s) per coupled iteration for each component of for H1C1 calculations

Solve Component	STH + Saturated	STH + bypass heating	ANTS + bypass heating
Total Time	81.1	134.0	240.0
TH	0.3	0.6	8.5
CMFD	27.0	86.8	147.0
Axial Transport	2.9	2.6	3.1
Self-Shielding	24.8	7.8	36.1
Macroscopic XS Update	4.8	4.8	4.9

5. CONCLUSIONS AND FUTURE WORK

This report discusses the development and implementation of a new physics model for the heating in the bypass flow region of BWRs. Results for this model are shown in comparison to simpler treatments. For most comparisons, the bypass heating model showed significant improvement over simpler models. However, there were a couple of exceptions, for which further improvements to the model have been identified. Additionally, these comparisons were also made between STH and ANTS. For comparisons with the TIP detectors, ANTS did not produce significantly different results from STH. However, it is still significant to have completed H1C1 with ANTS because it demonstrates the numerical improvements to ANTS that have improved its robustness. ANTS also provides a far greater degree of TH detail than STH, which may be necessary for certain applications even if it does not significantly impact the TIP calculations. Additionally, several other models were demonstrated with ANTS to further prove its increased robustness.

Several other improvements could be made to the multiphysics BWR modeling capability that were identified during the course of this work. They are listed in approximate order of priority.

- Currently, the amount of heat deposited in the bypass flow is determined partially by user input. Instead, the neutron and gamma fluxes calculated by VERA should be used to directly calculate the energy deposition in the bypass region. This will eliminate the need for a user to tune an input parameter. A spatial distribution of power deposition will also be more accurate than a fixed fraction.
- Improvements to neutronics portion of the solve is warranted. The neutronics solve for BWR models is consistently more difficult to converge than for PWR models. While ANTS presents a major performance improvement for VERA for BWR models, significant efforts are still needed to further reduce the computational burden of these calculations. Additionally, the more efficient self-shielding calculation that was recently implemented in VERA should be tested for whole-core BWR calculations.
- Heat conduction between the coolant and the bypass region through the channel box is currently neglected. Methods to model this conduction should be developed to further improve the accuracy of the bypass heating model is.
- Larger scale testing of the new ANTS numerical controls is needed, including analyzing full cycle depletions of multiple cores and cycles with the new options to ensure the solution remains robust and accurate.
- STH and CTF/ANTS currently use two different fuel models. The STH is a highly simplified 1D fuel conduction solver, whereas the CTF fuel solver *CTFFuel* accounts for a variety of burnup and irradiation effects that significantly impact fuel temperatures. Efforts are planned to modify MPACT to allow STH to be used for the fluid flow calculation while using *CTFFuel* for the conduction calculation. This will greatly improve the consistency of the two methods.
- The STH solver could be extended to refined spatial resolution, such as quarter assembly or subchannel resolution. Both of these resolutions have been implemented for the PWR version of STH, so extending them should prove trivial while offering an additional degree of accuracy for the feedback effects.

The first two bullet points above are already being pursued at the time of this writing. It is expected that they will result in significant improvements in both accuracy and performance.

6. ACKNOWLEDGMENTS

The authors wish to thank Exelon for initiating the VERA BWR project and providing the data necessary to build and validate the BWR models. The authors also wish to acknowledge that this project was funded through the US Department of Energy (DOE) Office of Nuclear Energy (NE).

This research made use of the resources of the Nuclear Computational Resource Center [16] at Idaho National Laboratory, which is supported by the Office of Nuclear Energy of the US Department of Energy and the Nuclear Science User Facilities under Contract No. DE-AC07-05ID14517.

REFERENCES

- [1] J. A. Turner et al. The virtual environment for reactor applications (VERA): Design and architecture. *Journal of Computational Physics*, 326(544), 2016.
- [2] Virtual Environment for Reactor Applications. Available at <https://vera.ornl.gov/>.
- [3] Consortium for Advanced Simulation of Lightwater Reactors. Available at <https://casl.gov>.
- [4] Brendan Kochunas, Benjamin Collins, Shane Stimpson, Robert Salko, Daniel Jabaay, Aaron Graham, Yuxuan Liu, Kang-Seog Kim, William Wieselquist, Andrew Godfrey, Kevin Clarno, Scott Palmtag, Thomas Downar, and Jess Gehin. VERA Core Simulator Methodology for PWR Cycle Depletion. *Nuclear Science and Engineering*, 185(1):217–231, 2017.
- [5] Benjamin S. Collins, Shane Stimpson, Blake W. Kelley, Mitchell T. H. Young, Brendan Kochunas, Aaron Graham, Edward W. Larsen, Thomas Downar, and Andrew Godfrey. Stability and Accuracy of 3D Neutron Transport Simulations Using the 2D/1D Method in MPACT. *Journal of Computational Physics*, 326:612–628, August 2016.
- [6] R. Salko, A. Wysocki, M. Avramova, A. Toptan, N. Porter, T. Blyth, C. Dances, A. Gomez, C. Jernigan, and J. Kelly. *CTF Theory Manual*. The North Carolina State University, 2017.
- [7] SCALE: A comprehensive modeling and simulation suite for nuclear safety analysis and design. Technical Report ORNL/TM-2005/39, Oak Ridge National Laboratory, 2011.
- [8] R. P. Pawlowski et al. Design of a high fidelity core simulator for analysis of pellet clad interaction. In *Proc. M&C 2015*, Nashville, TN, 2015.
- [9] B. Collins et al. CASL research and development activities and results for the CRUD induced power shift (CIPS) challenge problem. Technical Report CASL-U-2018-1703-000, Oak Ridge National Laboratory, September 2018.
- [10] Aaron Graham, Benjamin Collins, Bob Salko, and Mehdi Asgari. Multiscale Thermal Hydraulic Coupling Methods for Boiling Water Reactor Simulation. In *PHYSOR 2022: Making Virtual a Reality – Advancements in Reactor Physics to Leap Forward Reactor Operation and Deployment*, 2022.
- [11] Mehdi Asgari, Dave Kropaczek, Aaron Graham, Robert Salko, Kang Seog Kim, Kyle Gamble, Aysenur Toptan, Scott Palmtag, Thomas Downar, Brendan Kochunas, and Tomasz Kozlowski. Final report for “modeling and analysis of exelon bwrs for eigenvalue & thermal limits predictability” project. Technical Report ORNL/TM-2021-2349, Oak Ridge National Laboratory, 2021.
- [12] David J. Kropaczek, Robert Salko, Belgacem Hizoum, and Benjamin S. Collins. Two-phase subchannel method via non-linear iteration. *Submitted to: Nuclear Engineering and Design*, 2022.
- [13] Robert Salko, Belgacem Hizoum, David J. Kropaczek, and Mehdi Asgar. Verification and validation of the alternative nonlinear two-phase subchannel (ants) code. *Submitted to: Nuclear Engineering and Design*, 2022.
- [14] R.T. Lahey and F.J. Moody. *The Thermal-Hydraulics of a Boiling Water Nuclear Reactor, Second Edition*. American Nuclear Society, 1996.
- [15] N. H. Larsen and J. G. Goudey. Core Design and Operating Data for Cycle 1 of Hatch 1. Technical Report EPI-NP-562; TRN 79-007651, General Electric Co., January 1979.

- [16] Idaho National Laboratory. Nuclear Computational Resource Center. <https://inl.gov/ncrc>.
- [17] K. S. Kim, A. M. Graham, and M. A. Jessee. Dancoff-Based Wigner-Seitz Approximation for the Subgroup Resonance Self-Shielding in the VERA Neutronics Simulator MPACT. In *PHYSOR 2022: Making Virtual a Reality – Advancements in Reactor Physics to Leap Forward Reactor Operation and Deployment*, 2022.

