CASL TechNotes

Consortium for Advanced Simulation of Light Water Reactors

April 2014
Volume 2, Issue 1

About Tech Notes

*CASL **Tech Notes** are produced by CASL's technical staff to spotlight areas of particular technical interest and to discuss CASL science and engineering innovation, accomplishments, activities, and plans. **Tech Notes** is produced on a quarterly basis.*

In this Issue...

- Westinghouse Completes its AP1000 Test Stand
- VERA-CS Coupled Multi-physics Capability demonstrated in a Full-core Simulation
- TVA provides Watts Bar Operating Data to Support Validation of CASL Tools
- University of Michigan adds Depletion Capability to MPACT for VERA-CS
- VERA Data Transfer Processes Provided by DTK
- Alternate Coupling Algorithm Investigations Completed
- Role of HPC in the Nuclear Power Industry
- TVA Test Stand Launched

Westinghouse Completes its AP1000® Test Stand

Last June, CASL launched its first Virtual Environment for Reactor Applications (VERA) Test Stand on the Westinghouse engineering computer cluster. Westinghouse selected the AP1000® pressurized water reactor (PWR) first core for the Test Stand simulation because its advanced design and operational features, and imminent deployment in units under construction, make it a particularly relevant and engaging application of CASL technology.

The Westinghouse Test Stand used VERA release 3.2 and the AP1000 PWR models were built using the VERA common input. Westinghouse engineers commented on the intuitive nature of the VERA common input, noting its simplicity is deceptive - it enabled set up of a fairly advanced core design in a relatively compact ASCII file of less than 1,000 lines. Only minor modifications to the input are required to run variations, making the VERA common input practical and effective.

Westinghouse used the VERA core simulator (VERA-CS) SP_N solver option; CASL has demonstrated the accuracy of this solver for start-up physics calculations in earlier validation simulations of Watts Bar Unit 1 [Ref. 3]. The SP_N method is a computational option to the discrete ordinates (SN) method also available in Insilico. It is known that SP_N does not converge to the true transport solution as its expansion order is increased; however, it has good computational efficiency and can achieve improved accuracy over traditional nodal diffusion core simulators.

(p.4)

VERA-CS Coupled Multi-physics Capability demonstrated in a Full Core Simulation

In December, CASL reported on the latest results from its Watts Bar reactor progression problem modeling. The most recent simulation includes quarter-core representation with coupled neutronics (including embedded cross section generation and neutron transport) and thermal - hydraulics. Last July, the team completed a demonstration of physics coupling that included heat generation and thermal-hydraulic feedback for a single fuel assembly; this latest demonstration extended the demonstration to quarter-core resolution and implemented the VERA critical boron search capabilities.

The neutronics solution was generated with Insilico using cross sections from XSPROC and the (p.6)

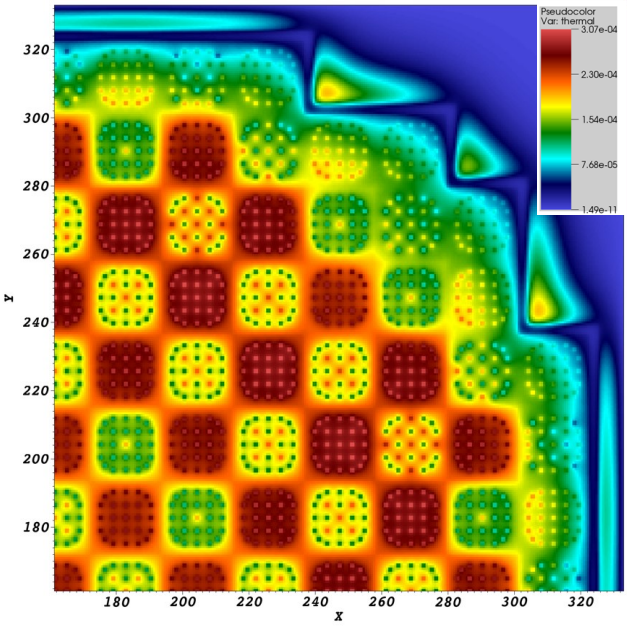

Figure 1 2D Slice of Thermal Flux Distribution near the Core Mid-plane

TVA Provides Watts Bar Operating Data to Support Development and Validation of CASL Tools

CASL partner Tennessee Valley Authority (TVA) provides more than 6,900 MW of electricity to the U.S. electrical grid through the operation of six commercial nuclear power reactors (3 PWRs and 3 BWRs), and will soon add another 1,180 MW with the startup of Watts Bar Unit 2. The PWR plants are of the same design (Westinghouse 4-loop Nuclear Steam Supply System) and together have data representative of a cumulative period of over 70 reactor years. These operating units provide invaluable data to the CASL Hub for code validation, and startup of Watts Bar Unit 2 provides a unique opportunity to acquire acceptance test data.

To support the development of CASL's VERA, TVA has provided information on past Watts Bar Unit 1 cycles, including detailed information such as core loading (p.7)

Coming in our next issue of TechNotes: Multi-physics predictions of CRUD and CIPS

University of Michigan adds Depletion Capability to MPACT for VERA-CS

MPACT is a three-dimensional (3-D) whole core transport code being developed by the University of Michigan and ORNL for the CASL VERA toolset. MPACT's integral transport solution is obtained by means of the method of characteristics (MOC) that employs discrete ray tracing for the radial direction and diffusion theory for the axial direction. In order to allow the analysis of commercial power reactors, the radiation transport code must be capable of describing the depletion of the nuclides in the reactor core; thus, a depletion module has been integrated into MPACT by University of Michigan researchers Ben Collins, Ang Zhu, Brendan Kochunas, and Tom Downar.

The numerical methods to implement nuclide point depletion and integrate a time dependent application within MPACT are based on those utilized by SCALE/ORIGEN-S [Ref. 1]. The module uses a hybrid approach that combines the matrix exponential method with Bateman equations for short-lived nuclides to obtain stable numerical solutions for all nuclides. This method eliminates the need to use linearized chains. The depletion of the nuclides is calculated for all regions and material combinations for each state point evaluated, and must therefore be extremely computationally efficient.

The exponential matrix method and the Bateman solutions complement each other in the approach used; the exponential matrix method is quite accurate when the transition coefficients are small but has problems when including large rate constants; conversely the Bateman solution has some numerical difficulties for extremely small rate constants, but is stable and accurate for large rate constants.

The point depletion calculation assumes a constant flux. To correct for this, MPACT adopts two commonly used techniques: the predictor-corrector and sub-step methods. The predictor-corrector method calculates the depletion at t_2 using a 1-group flux and cross section at t_1, produces a new predicted 1-group flux, cross section, and depletion using the predicted t_2 concentration, and then averages the two. Once the t_2 nuclides are obtained, a transport calculation is performed to obtain the steady state flux at t_2. The sub-step method is applied to perform multiple depletion calculations between transport calculations. Since the depletion calculation typically takes less time than the transport calculation this will often save computational time. Mathematically this is the equivalent of applying a time dependent normalization factor with a constant eigenvector representing the flux distribution to retain constant core power over a time step.

A separate depletion library was developed for MPACT based on the ORIGEN-S library which provides the basic decay constants and fission yields, as well as the 3-group cross-sections which are used for the depletion of the isotopes *not* contained in the MPACT multi-group transport library.

Prior to releasing the code update, the team ran several international benchmark problems to verify and validate the depletion routines implemented in MPACT. Additionally, the team ran several of CASL's Progression Problems. Figure 2 illustrates the 2D mid-core planar power distribution associated with the initial core loading and the depleted core for Watts Bar unit 1 Cycle 1. Comparisons with continuous energy Monte Carlo [Ref. 2] results indicate that MPACT provides an accurate result of the 2D power distribution with an eigenvalue difference of 133 pcm, an RMS pin difference of 0.295%, and a maximum pin difference of 0.894%. The 2D core was depleted at the average temperatures and densities used in the 3D model and the rated core thermal power was divided by the axial height of the fuel to obtain the power for the 2D core depletion of 9.33 MW. The expected radial shifts in the power distribution are observed in the depletion of the 2D core; the higher power locations in the core observed at the beginning of the cycle are reduced by the end of the cycle, with the power distribution flattening.

Currently, the depletion subroutine in MPACT assumes symmetric radial depletion of a fuel rod. The modification to azimuthally vary the depletion is expected to be relatively simple, but will considerably increase memory requirements. This and other capability enhancements are being considered for the next release of MPACT, scheduled for later this year. For more information on this work, see CASL-U-2013-0276-001.

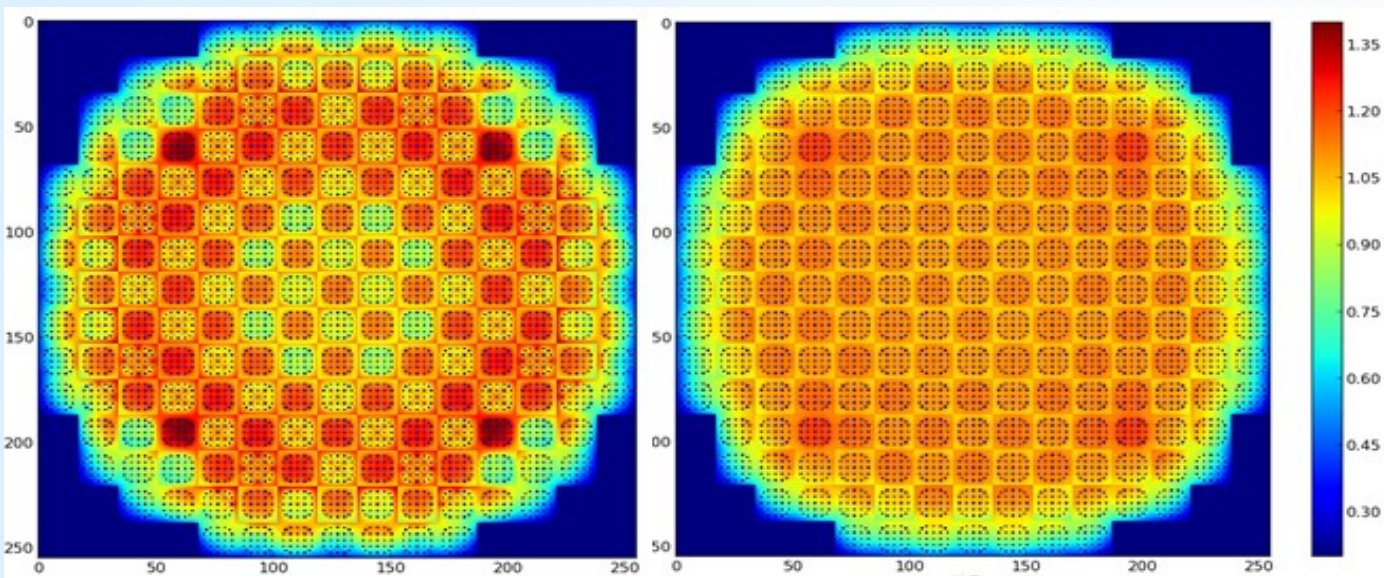

Figure 2 2D Core Power Distribution, Beginning (left) and End (right) of Cycle

Several works are referenced in the full CASL report; those works referred to in this article include:

[1] Y. Gokhan, K Clarno, I. Gauld, J. Galloway. *Modular ORIGEN-S for Multi-physics Code Systems*. International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (M&C 2011), Rio de Janeiro, RJ, Brazil, May 8-12, 2011, American Nuclear Society.

[2] SCALE: A Modular Code System for Performing Standardized ComputerAnalyses for Licensing Evaluation,"ORNL/TM-2005039, Version 6, Vols. I - III; see also CCC-750, Radiation Safety Information Computational Center, Oak Ridge National Laboratory , Jan. 2009

VERA Data Transfer Services Provided by DTK

In order to provide an infrastructure for data transfer operations between CASL physics codes, VERA includes the Data Transfer Kit (DTK). DTK is a software library developed for VERA, designed to provide parallel services for mesh and geometry searching and data transfer, providing a set of interfaces and tools used between geometric domains that may or may not conform in physical space.

Most of VERA's physics components utilize different mesh topologies to solve their part of the coupled problem. In a parallel formulation, such as that used by the VERA codes, if the geometries are arbitrarily decomposed, geometric alignment is not likely. To expedite parallel data transfer, it is necessary to relate the non-conformal meshes to each other. DTK provides a geometric rendezvous that manipulates the source and target geometries such that all geometric operations and data evaluation operations have a local formulation while data transfer occurs globally.

The algorithms implemented in DTK are based on the concept of geometric rendezvous as developed by Plimpton, Hendrickson, and Stewart [Ref. 7] originally implemented as part of the SIERRA framework [Ref. 8]. Their work has been extended to move towards a design for use with arbitrary physics codes such that varying representations of mesh, geometry, and fields are able to access these services [Ref. 9]. In addition, the original mesh-based rendezvous algorithms have been expanded to be used with both mesh and geometry representations of the geometric domain.

In DTK, the rendezvous algorithm developed by Plimpton et. al. behaves as a hierarchical parallel and geometric search tree. Using this algorithm, a secondary decomposition of a subset of the source mesh that will participate in data transfer is generated, forming the rendezvous decomposition. It can be viewed as a temporary copy of the source mesh subset that intersects the target geometry, generating a secondary decomposition of the geometric structures in the problem by imposing a geometry-based repartitioning, as illustrated in Figure 3. To transfer data between these meshes, each partition in each mesh needs to communicate data to each partition in other meshes, due to their geometric overlap. The rendezvous decomposition is a geometrically-balanced repartitioning of the source mesh in the transfer problem with the partitioning information shared amongst all meshes.

Once points have been accumulated in the rendezvous decomposition, a subset of the mesh that is in the vicinity of each target point is generated. This subset, which is typically much smaller than the mesh owned by a particular rendezvous process, is then searched with a more expensive point-in-element operation that transforms the point into the reference frame of each mesh element in the subset with a Newton iteration strategy. This mapped point is then checked against the canonical reference cell of that mesh element's topology to determine if the point is contained within.

A set of mapping algorithms based on geometric rendezvous are implemented within DTK that apply specifically to shared domain problems. A shared domain problem is one in which the geometric domains of the source and target intersect over all dimensions of the problem. Figure 4 gives an example of a shared domain problem in 3 dimensions. Here, Ω(S) (yellow) is the source geometry, Ω(T) (blue) is the target geometry, and Ω

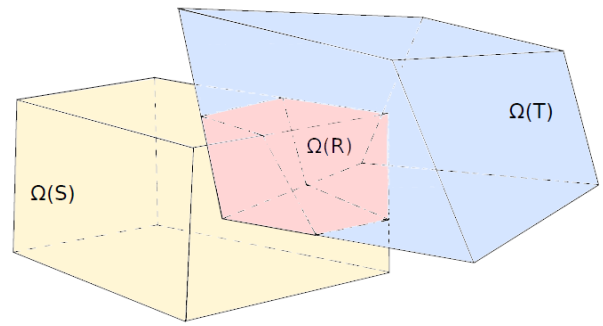

Figure 4 Shared domain example. Ω(S) (yellow) is the source geometry, Ω(T) (blue) is the target geometry, and Ω(R) (red) is the shared domain.

(R) (red) is their intersection and the shared domain over which mapping and the rendezvous decomposition will be generated. The purpose of these mapping algorithms is to efficiently generate a parallel topology map and the associated parallel communication plan that can carry out the data transfer repeatedly with the minimum required number of parallel messages and data.

A parallel topology map is an operator, M, that defines the translation of a field, $F(s) : R^D \rightarrow R^N$, from a source spatial and parallel domain, Ω_S, to a field, $G(t) : R^D \rightarrow R^N$, in the target spatial and parallel domain Ω_T, such that $G(t) \leftarrow M(F(s))$ and $M(t) : R^D \rightarrow R^N$; $\forall r \in [\Omega_S \cap \Omega_T]$, where N is the dimensionality of the field and D is the dimensionality of the spatial domain. It then follows that the geometric rendezvous is defined as a geometric-based parallel redistribution of the original source and target geometries defined over the region $\Omega_R = \Omega_S \cap \Omega_T$.

These maps are generated by creating source/target pairs found by searching the rendezvous decomposition. The rendezvous decomposition is searched for each target object to find the corresponding source object. In the case of finite element interpolation, the target object is a quadrature point in the target finite element mesh and the source object is a source element in the source finite element mesh that contains the target point. The map drives the field evaluation, $G(t) \leftarrow M(F(s))$, for all source/target pairs. Embedded within the map is a communication plan that describes the communication sequence for transferring the data from the source geometry to the target geometry. Once the field evaluations are complete, the communication sequence moves that data from the source geometry decomposition to the target geometry decomposition to complete the data transfer and the application to the map operator.

Figure 11 provides an illustration of the overall process. A typical use case of DTK in these cases is searching a mesh with a set of points and applying field data to those points through function evaluations (i.e., interpolation). For this use case, a scaling study of the mesh-based DTK implementation of the rendezvous (p.9)

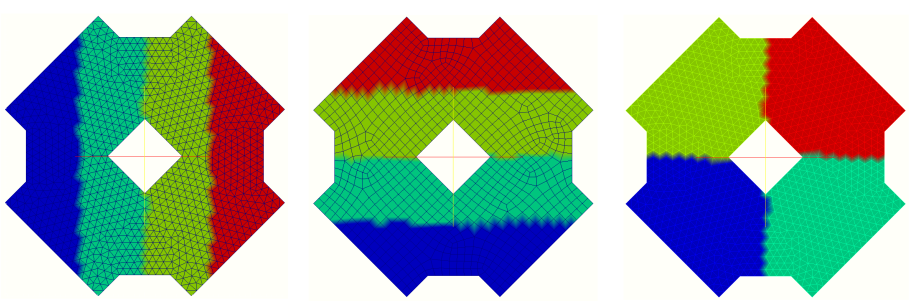

Figure 3 Rendezvous decomposition example. A triangle mesh (left) and a quadrilateral mesh (center) are partitioned into 4 parallel domains as indicated by color. The rendezvous decomposition (right) is generated as a geometric-based repartitioning of the source mesh that permits load balanced geometric operations.

Westinghouse Test Stand completed … (continued from p. 1)

VERA is known to be computationally demanding compared to typical industrial simulators, and in order to execute the Test Stand simulations Westinghouse purchased a dedicated computer cluster. The Westinghouse VERA cluster has 576 cores distributed on 48 nodes with a total memory of 4.6 TB (96 GB/node). Deployment to the Westinghouse cluster utilized the initial CASL VERA code build system, and provided ample opportunity to exercise the VERA code installation, build, and update process. Westinghouse did encounter several challenges associated with the first CASL Test Stand installation and resolution was attained thanks to a dedicated effort from CASL and Westinghouse personnel. The Westinghouse Test Stand deployment was instrumental in establishing a more streamlined code installation approach.

Since the initial Test Stand deployment, Westinghouse has successfully installed updates on two separate Westinghouse test stand systems. Additionally, Westinghouse worked closely with the CASL team to develop a VERA installation guide aimed at providing users with specific directions for configuring their computing cluster for VERA, installing the prerequisite utilities and third-party libraries, and obtaining later code updates. This guide has already been useful in deploying VERA to the EPRI Test Stand.

Calculations performed for the AP1000 PWR Test Stand focused on zero power physics test simulations, including predictions of 3D core hot zero power critical boron concentration, control rod worth, boron worth, Isothermal Temperature Coefficient (ITC), Moderator Temperature Coefficient (MTC), and Doppler Temperature Coefficient (DTC)). As illustrated in Figure 5, the results show excellent agreement with continuous energy Monte Carlo [Ref. 4] (CE MC) results, which are the numerical reference for the simulations performed.

Highlights from the reported results are provided in Tables 1 and 2.

- The Hot Zero Power (HZP) critical boron concentration predicted by VERA is within 3 ppm of the reference.
- The differential from the reference in predicted boron worth is 0.2 pcm/ppm.
- The root mean square (RMS) difference of the control rod worth prediction is 4 pcm, with a maximum difference of 9 pcm, across the eleven AP1000 control banks.
- The DTC, MTC and ITC, are respectively within 0.2, 0.4 and 0.6 pcm/F of the reference, with VERA consistently predicting more negative values; Westinghouse recommended that this be investigated further by the CASL team, as the differences observed in the temperature reactivity coefficients, while small, are not insignificant.

The results of the simulations are also close to the Westinghouse predictions using in-house core physics tools and licensed methods, which reinforces the confidence in Westinghouse's predictions for this advanced first core start-up, scheduled to occur within 2015.

Detailed power distribution analyses have also been completed using the CASL Test Stand. Several configurations were simulated, including an entire radial slice of the core with reflector and 3D multiple assemblies with partial control rod insertion, demonstrating the functionality of the AP1000 Axial Offset bank used for axial power distribution control. Comparison with CE MC simulations indicate a good agreement for VERA also from the power prediction standpoint, with a root mean square of the difference in pin power below 1.0% for all cases.

CASL is also working to incorporate its own CE MC capability, Shift, to VERA. The Westinghouse Test Stand applied Shift to the simulation of an entire AP1000 PWR 3D core model with multiple control banks inserted, as shown in Figure 6.

A primary objective of the CASL Test Stand was to obtain objective feedback on the applicability of the VERA tools to industrial applications. Westinghouse Fellow Engineer Fausto Franceschini led the Test Stand activities and performed the VERA simulations. Per Fausto, "VERA has demonstrated a potential to go one step beyond the current industry tools. While more extensive comparisons need to be performed, the results so far are very promising. VERA-CS could allow reliable high-fidelity predictions for a range of commercial reactors and operating conditions that could be used to anticipate and mitigate operational issues, resolve discrepancies with plant measurements, benchmark new methods, and corroborate the design of advanced cores." However, Fausto cautioned that the current simulations have already stretched the resources of Westinghouse's new industry-class engineering computing cluster. As the CASL team further implements depletion, thermal-hydraulic feedback, and fuel shuffling within VERA-CS, the computational needs are expected to multiply, and will likely exceed the capacity of the Westinghouse cluster. Westinghouse notes, "One of the challenges when developing a core simulator is to meet the computational resources available with an implementation of the underlying physics that leads to the appropriate accuracy for the foreseen range of applications and user base. (p.5)

Table 1 Hot Zero Power Reactivity Results

	CE MC	VERA	differential
Startup critical boron	**1313**	**1310**	**-3 ppm**
Boron Worth, pcm/ppm	-9.6	-9.4	+0.2
DTC, pcm/F	-1.54	-1.72	-0.18
Moderator Temperature Coefficient (MTC) pcm/F	-1.12	-1.50	-0.38
Isothermal Temperature Coefficient (ITC) pcm/F	-2.66	-3.22	-0.56

Note: CE MC temperature coefficient uncertainties estimated to be <0.1 pcm/F

Table 2 Control Bank Worth Results

		CE MC	VERA-CS	
Bank	**Material**	**Worth (pcm)**	**ΔWorth (pcm)**	**ΔWorth (%)**
MA	Tungsten	258	-1	-0.5
MB	Tungsten	217	-5	-2.1
MC	Tungsten	188	-2	-1.1
MD	Tungsten	234	0	0.0
M1	Ag-In-Cd	651	-4	-0.6
M2	Ag-In-Cd	887	3	0.4
AO	Ag-In-Cd	1635	-4	-0.3
S1	Ag-In-Cd	1079	0	0.0
S2	Ag-In-Cd	1096	-9	-0.8
S3	Ag-In-Cd	1124	0	0.0
S4	Ag-In-Cd	580	-3	-0.4
		RMS	**4**	**0.8**
		Max	**9**	**2.1**

Westinghouse Test Stand completed . . . (continued from p. 4)

The computational resources available now are much higher than when first-generation industry core simulators were conceived, but there are limits. It is acceptable to utilize leadership-class HPC such as ORNL's Titan for some applications, but it is also important to ensure that the VERA development path leads to capabilities that are usable by a broader engineering community."

The Westinghouse VERA Test Stand also identified other development items to address - thermal expansion capabilities to account for dimensional changes as the reactor heats up; provisions for advanced reflector and fuel designs; and the development of better documentation for VERA, though these are not cited as urgent development priorities.

All in all, the CASL-Westinghouse Test Stand was a resounding success. The code deployment methods were exercised, the codes were tested, and important experience feedback was obtained. Westinghouse PWR Core Methods Manager Bob Oelrich noted, "The benefits of the Test Stand go well beyond those of the technical analyses performed. The direct engagement of our engineering team with the CASL team and VERA-CS has raised much interest within Westinghouse. The enthusiasm and dedication offered to this project by the many people involved, inside and outside Westinghouse, are a testimony of their engagement and key contributions to the success of the Westinghouse VERA Test Stand".

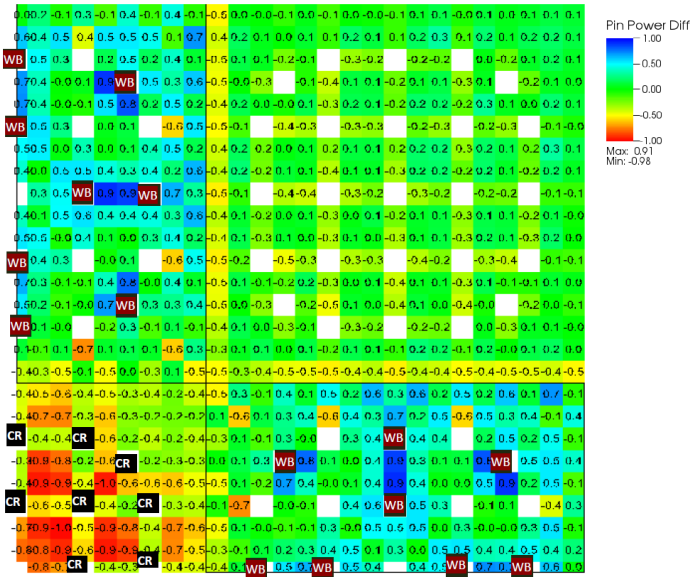

Figure 5 Pin-by-pin differential power distribution x100 for Ag-In-Cd Rodded Multi-lattice simulation (VERA versus CE MC [Ref. 4])

Westinghouse will continue to obtain VERA code updates as they become available, and plans to extend the AP1000 reactor simulations to cycle depletion with coupled thermal-hydraulics. For more information, see the Westinghouse VERA Test Stand report posted on the CASL website at http://www.casl.gov/docs/CASL-U-2014-0012-001.pdf.

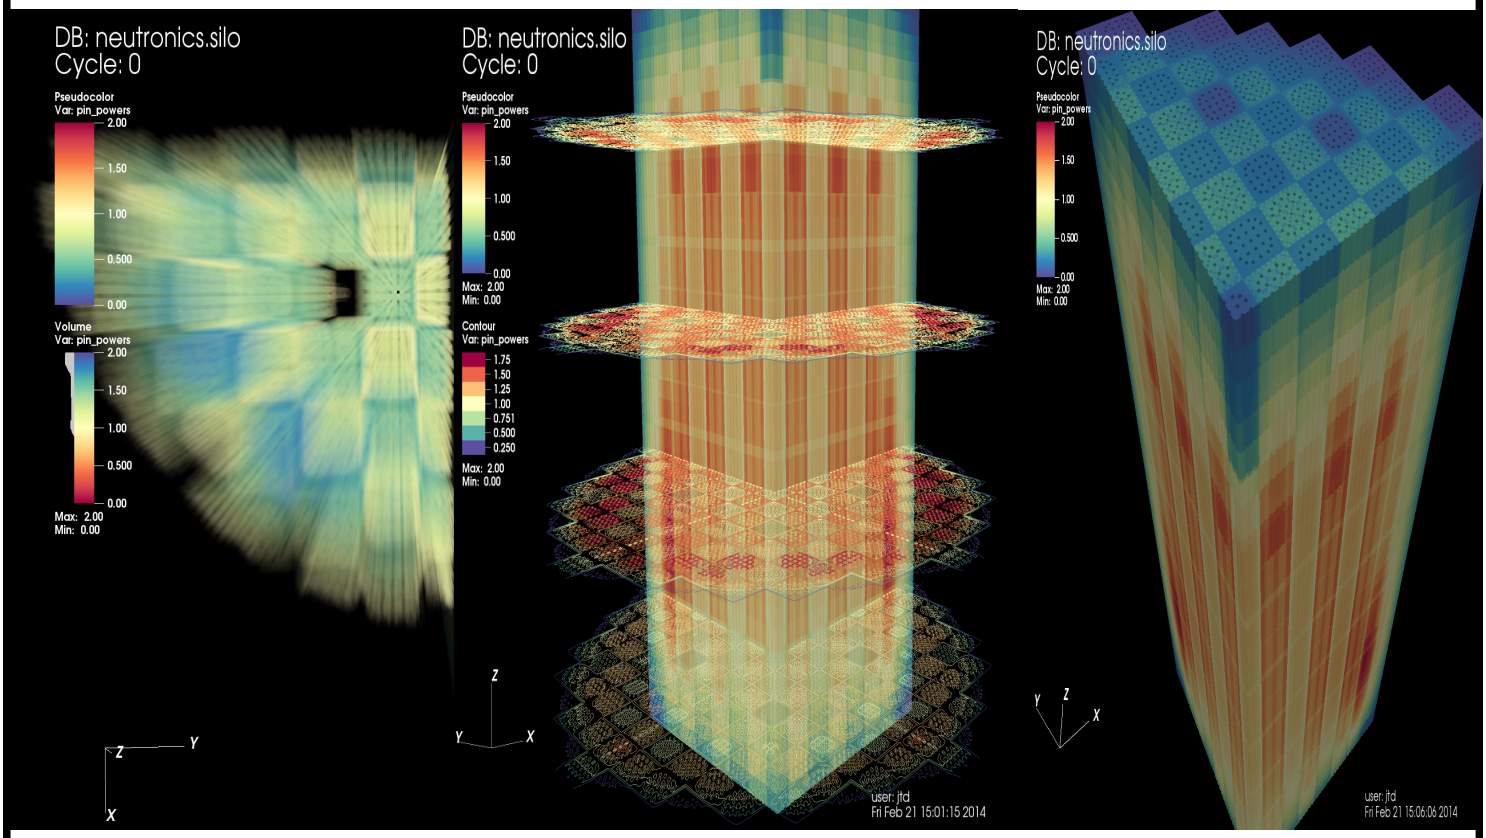

Figure 6 Power distribution from SHIFT simulation with simultaneous insertion of control banks to emulate MSHIM™ operation

Several works are referenced in the full CASL report; those works referred to in this article include:

[3] Gehin, J., A. Godfrey, F. Franceschini, T. Evans, B. Collins, S. Hamilton, Operational Reactor Model Demonstration with VERA: Watts Bar Unit 1 Cycle 1 Zero Power Physics Test, CASL Technical Report: CASL-U-2013-0105-001, June 2013.

[4] SCALE: A Modular Code System for Performing Standardized ComputerAnalyses for Licensing Evaluation,"ORNL/TM-2005039, Version 6, Vols. I - III; see also CCC-750, Radiation Safety Information Computational Center, Oak Ridge National Laboratory , Jan. 2009

AP1000® is a registered trademark of Westinghouse Electric Company LLC in the United States and may be registered in other countries throughout the world. All rights reserved. Unauthorized use is strictly prohibited. MSHIM™ is a trademark of Westinghouse Electric Company LLC in the United States and may be registered in other countries throughout the world. All rights reserved. Unauthorized use is strictly prohibited.

VERA-CS Coupled Multi-physics Capability Demonstrated... (continued from p. 1)

SP_N 3D transport solver. XSPROC generates cross sections for each unique region by running a 56-group pin cell calculation (BONAMI and XSDRN) for each region and spatially averaging the cross sections over the pin cell, then collapsing to 11 energy groups for the 3D transport solve. Because the cross-section calculation represents a significant portion of the computing time, the team refactored the code to distribute it over multiple processors. In a coupled calculation, the number of unique cross-sections that must be calculated by the neutronics codes quickly escalates. For a true full-core problem run, the number of unique cross-section regions that must be processed for the mesh employed is 3,705,625; taking advantage of quarter-core symmetry reduces the number of unique regions to 933,156.

The neutronics solution used the SP_N method with a 2x2 mesh in each pin cell and an axial mesh of 2.54 cm. The neutron flux is calculated from the bottom of the lower core plate to the top of the upper core plate in order to capture the axial leakage effects. The coolant flow solution is generated with a subchannel flow method (VERA Cobra-TF) using 49 axial levels in the active fuel region. The axial levels were defined to explicitly include the spacer grid heights and to use uniform mesh spacing between the spacer grids, for a maximum axial mesh length of about 7 cm. The fuel rod heat conduction model used 3 radial rings in each fuel rod.

Progression Problem 7 marks an important transition from smaller-scale coupled problems to a sizable industry-scale application, with massive amounts of data being handled by the code infrastructure. Data transfer between the neutronics code and the subchannel flow code, as illustrated by Figure 8, was completed at each pin cell using the same mesh defined for the channel flow calculation (the "coupling mesh"). The progression problem was run on the ORNL Titan supercomputer with 18,769 cores, and required 10 iterations and 17.5 hours to achieve convergence between the neutronics code and the channel flow code.

Figure 1 shows a 2D slice of the calculated thermal flux near the core mid-plane. The resolution of the solution allows a more accurate calculation of the flux variation from rod to rod, and features such as the Pyrex rods are discernable. This figure also shows the thermal flux leakage in the area outside of the active fuel region. Figure 7 shows the calculated coolant enthalpy in the active fuel region, and indicates that the enthalpy rise in the core is closely related to the assembly power, indicative of limited radial coolant mixing. Note the large enthalpy gradients at the core periphery.

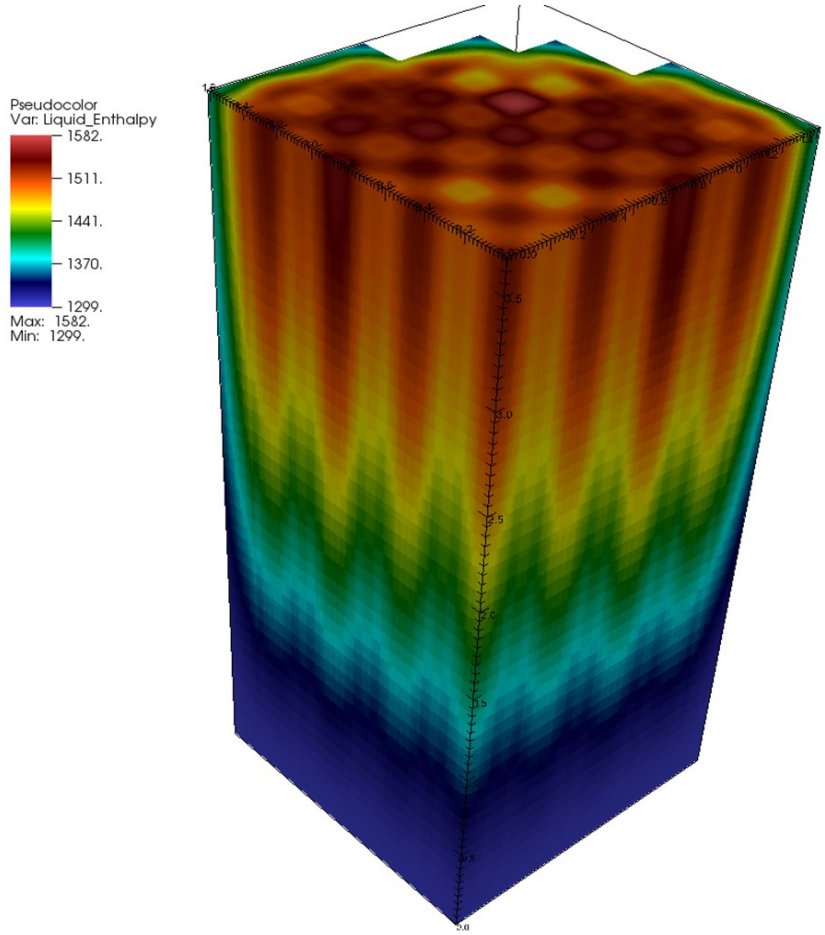

Figure 7 3D Coolant Enthalpy Distribution in the active fuel region .

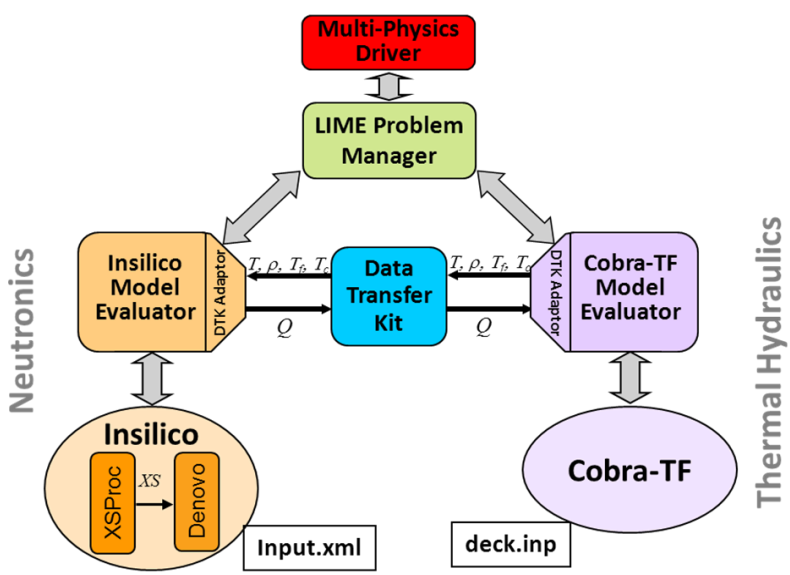

Figure 8 Key components of a coupled Insilico-CTF application created to solve the example problem.

While this work successfully demonstrates the VERA-CS coupled multi-physics capability, the team plans to work towards several potential improvements that were identified by this study. First, the team recommends convergence studies to develop an optimized set of modeling options for radial and axial mesh density (including fuel conduction mesh), SP_N order, energy group structure, scattering order, convergence criteria and the associated effects on run-time, along with the optimal number of cores. Additionally, the team recommended comparisons between the subchannel flow simplified fuel rod conduction model and the VERA fuel performance code predictions to ensure the simplified model is adequate. Finally, the team recommended investigating a dynamic convergence criteria strategy as a means of improving run time. With the upcoming addition of incore detectors and transient fission product models, VERA-CS will be ready for validation against Watts Bar Unit 1 power escalation measurements.

For more information, see CASL-U-2013-0196-000 and CASL-U-2013-0150-001.

TVA Provides Watts Bar Operating Data … (continued from p. 1)

patterns, local and total reactor power, axial power distributions, measured coolant flow rates and temperatures, coolant chemistry (including boron concentration), measured flux, control rod bank positions, and fuel inspection videos. Watts Bar Unit 1 operational data has been provided to CASL, from the 1996 startup tests though the end of the most recent cycle, for use in VERA validation.

The information collected from TVA's database is organized into three reports for each of the 12 past Watts Bar Unit 1 reactor cycles: core cycle design, core performance, and reported instrumentation data. The instrumentation data provides the measured data for selected instruments on a frequency of every five (5) days (averaged). This data will be supplemented as additional CASL simulations identify the need. TVA has also offered the possibility of augmenting plant instrumentation to obtain additional validation data, including additional startup data at Watts Bar Unit 2.

The CASL team has already taken advantage of the TVA database for comparison with the VERA-CS progression problem results. VERA-CS progression problems have used Watts Bar Unit 1 cycle 1 data to demonstrate capability; past CASL simulations include reactor startup testing, calculated control rod worths and critical boron, and hot full power power/coolant temperature distributions, as illustrated in Figure 9.

TVA is enthusiastic about the near-term benefits of participating in the CASL Hub, as well as the longer-term expected returns on the CASL investment. "This is the first higher-fidelity closely coupled multi-physics tool to tackle tough operating issues for the benefit of the industry at large. We believe that uncovering a deeper understanding of the feedback effects between the fuel performance, coolant conditions, and neutronics will enhance TVA's operating margins," said TVA Senior Vice President Joe Hoagland. "We plan to run the CASL codes to more fully understand operating issues like CRUD and coolant flow distribution effects at Watts Bar and our other reactors." CASL expects to utilize TVA's operating data to validate not only VERA-CS, but also to provide validation and corroborating information for many of the Challenge Problems.

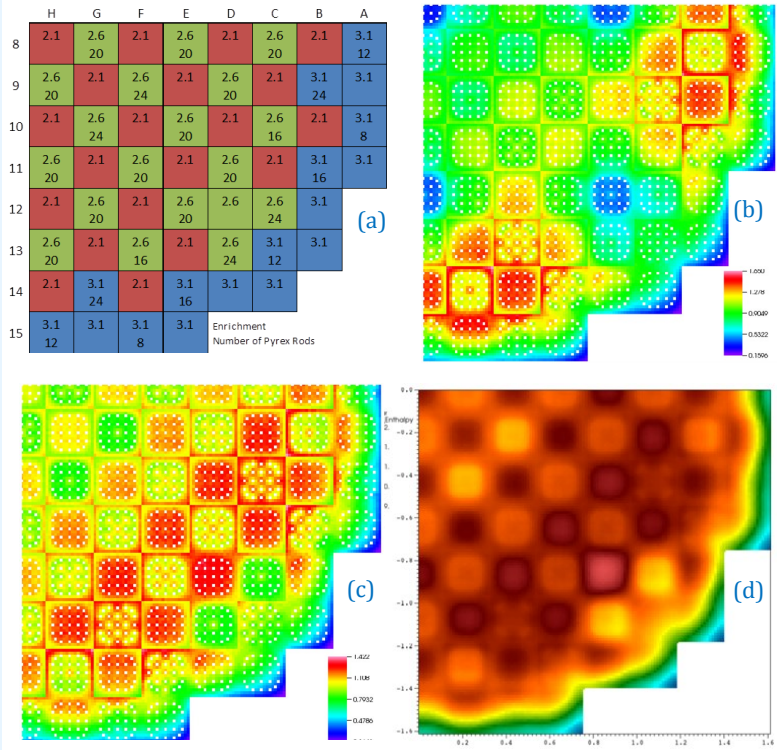

Figure 9 Watts Bar Cycle 1 (a) Core loading plan; (b) calculated radial fissions with control rod bank D fully inserted; (c) calculated radial fissions at initial criticality; and (d) calculated exit coolant enthalpy taken from CASL-U-2013-0196-000 and Virtual Environment for Reactor Applications (VERA) Physics Integration Status Update (Gehin, 2/25/2014).

Managing CASL Intellectual Property

As a consortium focused on applications to industry, the CASL Hub routinely handles Intellectual Property (IP); IP includes technical information, inventions, trade secrets, discoveries, know-how, and other proprietary ideas (whether or not it is patentable or copyrightable) belonging to CASL and others. This also includes derivative IP; that is, an expressive creation that includes major IP elements from a previous work. Also, because of the applied nature of CASL's mission, the distribution of some information must be restricted, in compliance with U.S. export control policy, regulation and law.

One of CASL's primary goals is to rapidly and successfully transfer nuclear reactor simulation and modeling technologies to the U.S. nuclear industry, yet this goal must be tempered by the need to protect the consortium's and the United States' intellectual property. In order to ensure that CASL information is handle appropriately, the Hub implemented the CASL Intellectual Property Management Plan (IPMP). The IPMP plays an essential role in the delivery, storage, handling and dispositioning of all CASL data, and addresses information sharing, conflicts of interest, ownership of inventions and copyrights, and establishes the CASL Master Non-Disclosure Agreement (NDA) as a vehicle to promote the sharing of information within the CASL team.

A non-trivial subset of CASL's IP considerations is the software being developed by the team. Both legacy software (i.e. software that pre-dates CASL's existence) and CASL funded software (modified legacy software and virgin software) are being integrated to CASL's VERA, and this makes deployment of VERA a complex task.

To support the future broad deployment of VERA and in the interim to allow for partner testing of VERA, CASL has used government use authorization (GUA) licenses. The GUA allows CASL partners, within the constraints of the agreement, to have access to, use, and modify certain VERA components in support of CASL research and development in the overall VERA development effort. A CASL Testing & Evaluation (T&E) license has also been recently established to allow a broader range of testing by non-CASL participants. The T&E license is expected to evolve to commercialization and non-commercialization licenses in support of the longer-term mission of broad deployment.

VERA has been released externally for testing and evaluation; because VERA is undergoing active development, only a subset of the VERA components is included. Those interested in this release should send an email to info@casl.gov and include the following information: Name and citizenship of the requestor, company/institution, intended application, brief description of the intended tests and evaluations to be performed, and full contact information.

Alternate Coupling Algorithm Investigation Completed

At the end of March, the CASL Physics Integration team (PHI) delivered a report investigating alternate coupling approaches in terms of efficiency, accuracy and robustness. VERA-CS currently relies on Picard iteration to provide coupling between the neutronics, fuel performance, and subchannel flow codes. Mark Berrill (ORNL), Kevin Clarno (ORNL), Steve Hamilton (ORNL), and Roger Pawlowski (SNL) looked at four alternatives to the currently used damped Picard Iteration method:

- Picard Iteration with Anderson Acceleration (Anderson),
- Preconditioned Jacobian-Free Newton-Krylov (JFNK),
- Preconditioned JFNK with an approximate Jacobian vector product that neglects cross section processing at linear iterations (MJFNK1), and
- Preconditioned JFNK with an approximate Jacobian vector product that uses a linear approximation of the cross sections at linear iterations (MJFNK2).

To test the behavior of these variants, the team used AMA Progression Problem 6 [Ref. 5] as a benchmark. Problem 6 consists of a single 17x17 PWR fuel assembly at hot full power, beginning of cycle. For the base configuration, Problem 6 was modeled using a single assembly power of 17.67 MW and 1300 ppm soluble boron concentration with a 56-energy group cross section library. All runs were executed in parallel on 289 processing cores (1 unit cell per core). The physics codes used in the study included XSProc (cross-section processing), Insilico radiation transport calculations (SP_3/P1 scattering), and AMP [Ref. 6] (linear continuous finite element discretization) .

Figure 10 shows the convergence behavior for Picard iteration and each of the JFNK variants. Picard iteration demonstrates the expected linear convergence rate and convergence for the full JFNK is quadratic. MJFNK1 follows the convergence behavior of the full JFNK for the first few iterations before separating and ultimately converging linearly, while MJFNK2 follows the convergence rate of full JFNK even more closely.

The investigation included parameterization studies on power, coolant boron concentration, the number of cross section energy groups specified in the problem, and the SP_N order. The work also broke down the solution time as a function of time spent in each physics component and as a function of the number of times each physics code was applied. Table 3 provides the time required to achieve a relative convergence tolerance of 10^{-4} for each solution strategy at four different power levels. Note that the damping factor used for the Picard Iteration variant was 0.45, which appears to be nearly optimal for a wide range of problems.

Table 3 Solution Time (seconds) as a function of Power Level

Method	Power Level 60%	80%	100%	120%
Picard	1296	1150	1353	1523
Anderson	1305	-	-	-
JFNK	6168	5958	8908	8844
MJFNK1	734	1035	1693	3164
MJFNK2	695	985	971	961

Table 4 shows the overall timing results for three different combinations of cross section processing and transport energy group structures. To evaluate the range of performance, the study used a range of cross section energy groups. As the number of energy groups is increased, the time spent in cross section processing quickly dominates the overall runtime, especially for the JFNK approach, since this approach requires full cross section processing at every linear iteration.

Table 4 Solution Time (seconds) for different XSProc Energy Group Structures

Method	XSProc Groups/SP_N Groups 8/8	56/23	252/23
Picard	980	1353	3868
JFNK	1685	8908	51727
MJFNK1	1070	1693	3793
MJFNK2	636	971	2013

For the large numbers of energy groups typically required for accurate reactor physics calculations it is apparent that native application of JFNK will result in a much higher computational expense than the currently used Picard iteration. However, a reduction in computational expense may be achieved (about half) for MJFNK2 in the 252 group XSProc case. Unfortunately, many of the CASL legacy codes require major rework to implement MJFNK2.

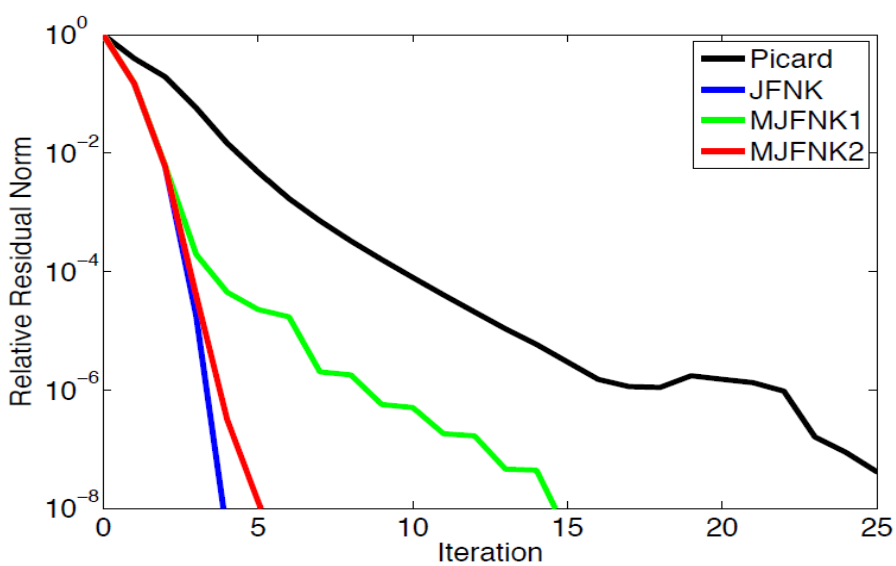

Figure 10 Nonlinear Convergence Behavior for Picard and JFNK Variants

Overall, the results of the study on Progression Problem 6 indicate that:

- Approximations to the nonlinear operator that avoid cross section processing can largely preserve the fast convergence rate of JFNK without the overhead,
- Computational savings can be achieved relative to Picard iteration with Newton-like methods, and
- The Anderson acceleration method was found to be 40% faster on simple test problems, but at realistic reactor conditions for full assemblies, it failed to converge. Further research on globalization strategies to achieve a contraction mapping will be needed to fully assess the method.

For more information, see L3 : PHI.CMD.P8.01.

Several works are referenced in the full CASL report; those works referred to in this article include:

[5] S. Palmtag. Coupled single assembly solution with VERA (Problem 6). CASL-U-2013-015-000.

[6] K. T. Clarno et al., "The AMP (Advanced MultiPhysics) Nuclear Fuel Performance Code,"*Nucl. Eng. Design* **252**(1), 108-120, November 2012.

VERA Data Transfer Services … (continued from pg. 3)

algorithm for data transfer was performed on the Titan Cray XK7 system at the Oak Ridge Leadership Computing Facility to assess recent performance improvements to the library. For each calculation, a tri-linear hexahedron mesh was generated to represent the source geometry with cube-shaped partitions on each parallel process. The target mesh was represented by random interpolation points located within the bounding box of the local target partition. In addition, the source and target mesh partitions were constructed with a reverse ordering such that the physical domains of no two partitions reside on the same parallel process.

In both the strong and weak scaling studies, two cases were analyzed: a best case scenario and a typical case scenario. For the former, there was no geometric overlap between the source and target partitions; only the source partitions were required to transfer data to the corresponding target partition. For the latter, partitions were allowed to overlap such that a source partition was required to transfer data to up to 27 target partitions, creating a greater communication load on the algorithm and representing a more typical physics-based mesh partitioning.

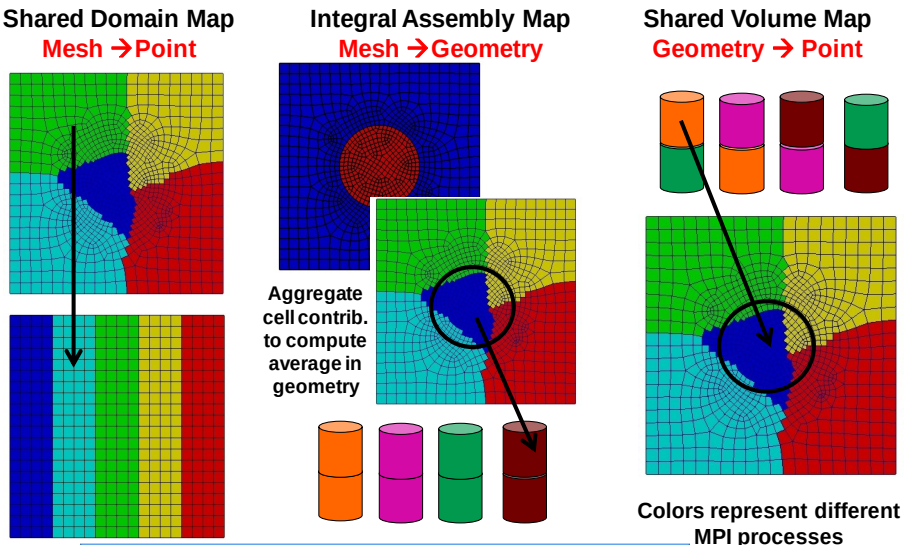

Figure 11 DTK Rendezvous Process

As illustrated in Figure 12, the best-case and typical-case scenarios indicate good scaling for the setup calculation and excellent scaling for in the apply calculation. In addition, the apply calculation has a significantly smaller runtime than the setup calculation. This is important for data transfer problems where the underlying mesh or geometry does not change, as the setup calculation will be performed once while the apply calculation will be performed each time data is transferred. In addition to parallel performance improvements, performance of the local computational kernel was also improved, as shown in Figure 12c. Although the time complexity of the implementation remains the same, performance improvements were made by updating the mesh-searching framework and the construction of the mesh in the rendezvous decomposition. Current DTK development work is focues on providing surface transfer capability for anticipated FY15 problems; look for new codes releases by the end of September. For more information see S.R. Slattery, P.P.H.Wilson, and R.P. Pawlowski. The Data Transfer Kit: A Geometric Rendezvous-Based Tool for Multiphysics Data Transfer, International Conference on Mathematics and Computational Methods Applied to Nuclear Science & Engineering (M&C 2013), Sun Valley, Idaho, 2013.

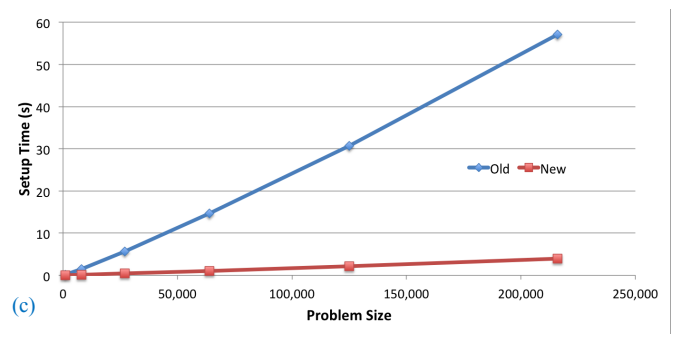

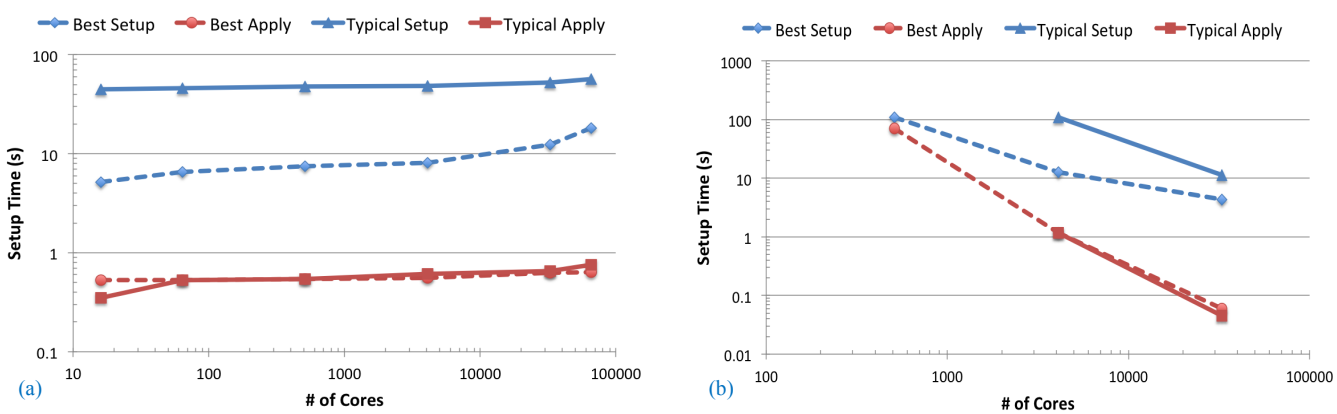

Figure 12 (a) Weak scaling study results. The blue curves report setup times and the red curves report apply times with dashed lines for the best-case scenario and solid lines for the typical case scenario. **(b)** Strong scaling study results. The blue curves report setup times and the red curves report apply times with dashed lines for the best-case scenario and solid lines for the typical case scenario. **(c)** Performance of the setup operation on a single process as a function of problem size (a problem size of 100 indicates 100 hexahedron elements mapped to 100 target points).

Several works are referenced in the full CASL report; those works referred to in this article include:

[7] S. Plimpton, B. Hendrickson, and J. Stewart, "A parallel rendezvous algorithm for interpolation between multiple grids," Journal of Parallel and Distributed Computing, vol. 64, pp. 266–276, 2004.

[5] J. Stewart and H. Edwards, "A framework approach for developing parallel adaptive multiphysics applications," Finite Elements in Analysis and Design, vol. 40, pp. 1599–1617, 2004.

[6] K. Chand, L. Diachin, C. Ollivier-Gooch, E. Seol, M. Shephard, T. Tautges, and H. Trease, "Toward interoperable mesh, geometry and field components for pde simulation development," Engineering with Computers, vol. 24, pp. 165–182, 2008.

Role of HPC in the Nuclear Power Industry

CASL's Chief Scientist, Paul Turinsky

Simulations are vital to the nuclear power industry, since — in contrast to some other industries — experiments at representative operating conditions are not always feasible. In the dawn of the nuclear power industry, back in the 1960s and 70s, the industry employed the leading edge HPC capabilities available to complete reactor core and system simulations. These simulations, validated by experiments, served as the basis for design decisions and licensing approvals.

Since that time, growth in computational power has outpaced all expectations. However, as opposed to other technology-based industries, the nuclear power industry did not expand its reliance on HPC. It is interesting to speculate on the reasons for this. Was it that modeling capabilities judged necessary could be realized on workstations? Was it that stagnation in the nuclear power marketplace discouraged investments in HPC capabilities? Was it the repetitive, labor-intensive nature of reload simulations that suggested the utilization of workstations over HPC? One could speculate endlessly on why HPC has not retained a more prominent role in the nuclear power industry in the past.

That said, a short survey of the industry indicates that lag may be swiftly changing. The CASL Industry Test Stand, of which two are noted in this issue of TechNotes, provides a demonstration of the industry's renewed interest in exploiting HPC.

CASL is demonstrating that HPC-enabled increased fidelity simulation capability can provide strong benefits to the nuclear power industry — the simulation results presented in this issue of TechNotes alone demonstrate that capacity. The benefits are derived from mesh refinements, the usage of models that more accurately capture the underlying physics, and through the capability to capture feedback effects. Some of these simulations can be completed on industry-class HPCs (hundreds to a few thousand cores), but others require leadership class HPCs (hundreds of thousands cores).

Of course the HPC of today will become the workstation of tomorrow. For more repetitive simulations, such as associated with core follow, industry-class HPC resources would seem appropriate; whereas, for non-repetitive simulations, such as designing a new spacer grid, leading edge HPC resources may be necessary and practical. There has been an implicit assumption that simulation model fidelity is enhanced (through finer simulation granularity) when greater computer power is available; this remains to be broadly confirmed. Another theory involves the use of HPC-based simulations to inform simulations that can be performed on lower power HPCs and workstations.

These notions have yet to be substantiated; however, what is clear is that CASL, by example, needs to demonstrate benefits to the nuclear power industry that will support the business case of introducing CASL software and associated needed HPC resources into the nuclear power industry workplace.

TVA Test Stand Launched

Over the last year, TVA has discussed with CASL the potential to simulate what appears to be a highly coupled phenomenon observed in several PWRs called Lower Plenum Flow Anomaly (LPFA). LPFA is characterized by reactor power fluctuations during operation and is thought to be associated with asymmetric thermal-hydraulic conditions originating in the lower plenum of the reactor. The formation and dissipation of vortices in the coolant may produce coolant flow rate and density differentials in the reactor core, resulting in a change in local power. The highly coupled nature of the effect will require a core simulator capable of reflecting the flow asymmetries and physics feedback effects.

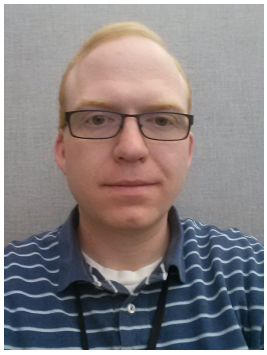

TVA Test Stand lead, Bill Bird

Early in April, TVA will launch a Test Stand aimed at simulating the flow distributions in the Watts Bar Unit 1 reactor using CASL's Hydra-TH computational fluid dynamics (CFD) code. TVA Test Stand lead Bill Bird explained that this is just the first step in simulating the LPFA phenomena; the results from the CFD simulation will be used as inputs to a VERA-CS simulation.

TVA's Nuclear Reactor Engineering and Fuels General Manager, Jim Lemons, suggested the simulation, noting that the explicit flow distribution and feedback to the neutronics and core power distribution are not currently modeled by the current industry reactor simulators, and this type of simulation should provide a much clearer picture of the power distribution mis-predictions that are routinely observed at Watts Bar.

TVA Test Stand lead Bill Bird observed that the project is extremely challenging and will require high performance computing capability to complete. TVA doesn't own a HPC system capable of this simulation; thus, the team will apply for an industry allocation on ORNL's Titan supercomputer.

CASL's series of test stands are aimed at implementing the CASL codes in industrial settings to solve industrial problems. The test stands are integral to CASL's long-term goal of producing tools that can be used to inform industry decisions regarding extending the life of existing commercial power reactors.

TVA CASL Project Manager Rose Montgomery

About CASL

CASL is a DOE Energy Innovation Hub focused on modeling & simulation of Commercial Light Water Reactors. CASL connects fundamental research and technology development through an integrated partnership of government, academia and industry that extends across the nuclear energy enterprise. Learn more at www.casl.gov.

U.S. DEPARTMENT OF ENERGY
Nuclear Energy

Acknowledgments

This document is sponsored by the U.S. Government under contract DE-AC05-00OR22725 with UT-Battelle, LLC, which manages the Oak Ridge National Laboratory (ORNL). Accordingly, the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce these documents, or to allow others to do so, for U.S. Government purposes. Unless otherwise noted, they have been placed in the public domain, although we request the following credit line be used when documents reproduced elsewhere: "Courtesy of CASL, Oak Ridge National Laboratory, U.S. Dept. of Energy.

Disclaimer

This document was prepared as an account of work sponsored by an agency of the U.S. Government. Neither the U.S. Government nor any agency thereof, or 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. Further, ORNL is not responsible for the contents of any off-site pages referenced. Reference 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 U.S. Government or any agency thereof. The views and opinions of document authors do not necessarily state or reflect those of the U.S. Government or any agency thereof.

Contact Us:

Technical Editor	Rose Montgomery	Oak Ridge National Laboratory
Distribution	Linda Weltman	PO Box 2008, MS6003
Email	casl-info@casl.gov	Oak Ridge, TN 37831-6003

OAK RIDGE NATIONAL LABORATORY
MANAGED BY UT-BATTELLE FOR THE DEPARTMENT OF ENERGY