

NEAMS Milestone Report: M2MS–20OR030102 FW–CADIS PWR Ex-Core Analysis with Shift through VERA



Tara Pandya
Katherine Royston
Thomas Evans

July 1, 2020

DOCUMENT AVAILABILITY

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

Website: <http://www.osti.gov/scitech/>

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: <http://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: report@osti.gov
Website: <http://www.osti.gov/contact.html>

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.

Reactor and Nuclear Systems Division

**NEAMS Milestone Report: M2MS-20OR030102
FW-CADIS PWR Ex-Core Analysis with Shift through VERA**

Tara Pandya
Katherine Royston
Thomas Evans

Date Published: July 1 2020

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	v
LIST OF TABLES	vii
ACRONYMS	ix
ACKNOWLEDGMENTS	xi
EXECUTIVE SUMMARY	xiii
1. INTRODUCTION	1
2. OPTIMIZATION METHODS	2
2.1 MULTITHREADING	2
2.2 MULTIGROUP CROSS SECTION REDUCTION	3
2.3 DOMAIN DECOMPOSITION	3
2.3.1 Implementation	3
2.3.2 Guidance on Decomposing Domains	4
2.4 FW-CADIS VARIANCE REDUCTION	5
3. RESULTS	7
3.1 MULTITHREADING	7
3.1.1 5-mini Model	7
3.1.2 Small Modular Reactor Model	8
3.1.3 Full-Core Subcritical Watts Bar Nuclear Plant Unit 1 (WBN1) Model	11
3.2 MULTIGROUP CROSS SECTION REDUCTION	14
3.3 DOMAIN DECOMPOSITION	15
3.3.1 DD Memory Analysis	15
3.3.2 5-mini Model	18
3.3.3 SMR Model	21
3.4 FW-CADIS	25
4. SUMMARY	29
4.1 FUTURE WORK	29
REFERENCES	31

LIST OF FIGURES

1	Axial slice of 5-mini problem geometry at $z=50$ cm.	7
2	Axial slice of the quarter symmetry SMR model at an elevation of 200 cm with unique pins.	9
3	Quarter symmetry SMR ^{235}U fission source distribution	9
4	Total neutron flux and relative error in vessel and bioshield of quarter symmetry SMR at an elevation of 120 cm	10
5	Full-core SMR ^{235}U fission source spatial distribution.	12
6	Total neutron flux and relative error in vessel and bioshield of full core SMR at an elevation of 120 cm	13
7	Axial slice of the WBN1 model at an elevation of 200 cm.	14
8	Peak memory consumption after generation of the SMR model, metadata, and geometry.	16
9	Memory cost of the material compositions in the non-depleted fuel SMR problem.	17
10	Quarter symmetry SMR ^{235}U fission source distribution run with DD.	23
11	Total neutron flux and relative error in vessel and bioshield of quarter symmetry SMR at an elevation of 120 cm run with DD and CADIS.	24
12	Axial slices of the quarter symmetry SMR model without unique pins and ex-core detectors.	25
13	Quarter symmetry SMR Denovo total flux.	26
14	Quarter symmetry SMR CADIS adjoint source distributions.	27
15	Quarter symmetry SMR FW-CADIS adjoint source distributions.	27
16	Quarter symmetry SMR adjoint function spatial distributions.	28

LIST OF TABLES

1	5-mini peak memory usage during Shift transport with varying number of threads.	8
2	SMR computation times and memory usage with Shift on 4 nodes and varying number of MPI tasks and threads.	11
3	SMR computation times and memory usage with increasing number of threads.	11
4	Comparison of SMR computation times and memory usage with and without cross-section reduction for calculation of variance reduction parameters.	14
5	Domain-decomposed 5-mini forward case memory usage and computation times with isotopic composition transfer.	19
6	Domain-decomposed 5-mini forward case memory usage and computation times transferring all.	19
7	Load-balanced versus not load-balanced domain-decomposed 5-mini CADIS case memory usage and computation times with isotopic composition transfer.	20
8	Load-balanced domain-decomposed 5-mini CADIS case memory usage and computation times transferring all.	21
9	Domain-decomposed SMR forward case memory usage and computation times with isotopic composition transfer.	22
10	Domain-decomposed SMR forward case memory usage and computation times transferring all.	22
11	Domain-decomposed SMR CADIS case memory usage and computation times.	23
12	Detector flux tally results for SMR ex-core geometry.	27

ACRONYMS

API	application programming interface
CADES	Compute and Data Environment for Science
CADIS	Consistent Adjoint Driven Importance Sampling
CASL	Consortium for Advanced Simulation of Light Water Reactors
CE	continuous energy
CTF	COBRA-TF
DD	domain decomposition
DR	domain replication
FW-CADIS	Forward-Weighted CADIS
HPC	high performance computing
HZP	hot zero power
KBA	Koch-Baker-Alcouffe
MC	Monte Carlo
MPI	message passing interface
MSOD	multiple-set overlapping domain
ORNL	Oak Ridge National Laboratory
PWR	pressurized water reactor
SMR	small modular reactor
TVA	Tennessee Valley Authority
VERA	Virtual Environment for Reactor Applications
WBN1	Watts Bar Nuclear Plant Unit 1

ACKNOWLEDGMENTS

The authors would like to acknowledge the MPACT and SCALE developers for their assistance and collaboration in this milestone work, especially the efforts of the following individuals:

- Shane Stimpson
- Bob Salko
- Shane Henderson
- Mark Baird
- Eva Davidson
- Elliott Biondo

The authors would also like to acknowledge and thank Gary Wolfram and the Tennessee Valley Authority (TVA) for providing the data needed for the WBN1 model.

Preliminary work related to this research was supported by the Consortium for Advanced Simulation of Light Water Reactors(<http://www.casl.gov>), an Energy Innovation Hub (<http://www.energy.gov/hubs>) for Modeling and Simulation of Nuclear Reactors under US Department of Energy (DOE) Contract No. DE-AC05-00OR22725. This research used the resources of the Compute and Data Environment for Science (CADES) at the Oak Ridge National Laboratory, which is supported by the Office of Science of the US Department of Energy under Contract No. DE-AC05-00OR22725.

EXECUTIVE SUMMARY

This report presents the work completed for the NEAMS milestone M2MS–20OR030102 titled "FW–CADIS PWR Ex-Core Analysis with Shift through VERA." The work completed for this milestone includes the implementation, integration, and optimization of memory and performance improvement methods in Shift for fully coupled ex-core calculations through Virtual Environment for Reactor Applications (VERA). Fully coupled in this context means the transfer of moderator boron concentration, pin-wise fission source, depleted compositions, temperatures, and moderator densities from MPACT (with COBRA-TF (CTF)) to Shift. The ability to run ex-core calculations with VERA has been enabled and used for several years by Consortium for Advanced Simulation of Light Water Reactors (CASL) partners. However, this implementation was limited and potentially computationally burdensome. This work has enabled the ability to run higher-fidelity ex-core calculations on moderate computing clusters by focusing on multithreading, domain decomposition, and Forward-Weighted CADIS (FW-CADIS) variance reduction. Tests performed on small cores, a small modular reactor (SMR), and CASL progression problems show very promising memory reduction and computational performance. Recommendations for settings when running fully coupled high-fidelity ex-core calculations with VERA are documented. Without these optimization methods, many processors on a compute node would be left unused for the entire ex-core calculation. Therefore, these methods enable the user to better use the resources available and reduce computation time.

ABSTRACT

This report documents the integration, implementation, and optimization of fully coupled high fidelity ex-core analyses with Shift through Virtual Environment for Reactor Applications (VERA). Fully coupled in this context means the transfer of moderator boron concentration, pin-wise fission source, depleted compositions, temperatures, and moderator densities from MPACT (with COBRA-TF (CTF)) to Shift. The ability to run ex-core calculations with VERA has been enabled for several years, and this feature has been used recently by industry. This work has investigated the use of multithreading, domain decomposition, and Forward-Weighted CADIS (FW-CADIS) to reduce the memory footprint and enhance the computational performance of ex-core calculations with Shift through VERA. These new capabilities will enable ex-core calculations that could not previously be simulated on moderately sized high performance computing (HPC) machines.

1. INTRODUCTION

In order to run high-fidelity ex-core calculations, high performance computing (HPC) is needed. With this need comes limitations in modern day computing architectures, specifically lower on-node memory and more processors per node than previous high performance machines. Virtual Environment for Reactor Applications (VERA) can take advantage of the computing power of modern day moderate computing clusters to perform pressurized water reactor (PWR) analyses of practical interest [1].

The purpose of this milestone was to optimize and further develop the ability to perform fully coupled ex-core analysis as applied to PWRs using Shift [2] through VERA. A fully coupled problem is one that couples pin-wise depleted compositions, fission source, temperatures, boron concentration, and moderator densities for each state point between MPACT [3] (with COBRA-TF (CTF) [4]) and Shift [5]. To achieve full coupling, a large memory footprint is needed for the Monte Carlo (MC) calculation, including the deterministic forward and adjoint calculations used for variance reduction. The work completed under this milestone integrated and implemented memory-saving methods into Shift and VERA to make running these ex-core calculations possible.

When discussing the methods and results throughout the rest of this report, the term *unique pins* is used. This means that each pin location in the core geometry has unique materials. For example, for a fuel pin with a single ring, clad, moderator, and 50 axial levels (including spacer grids), Shift builds 50*3 unique materials for this fuel pin. Therefore, an assembly with a 17×17 array of these fuel pins would have 43,350 unique materials. Extrapolating to a quarter-core problem leads to a number of unique material regions in the MC calculation on the order of millions. The need to store this material information in the reactor metadata in Shift is what causes the large memory footprint.

This work builds upon the development of ex-core calculations in VERA under Consortium for Advanced Simulation of Light Water Reactors (CASL). This report is laid out in the following manner. Section 2. details the integration and implementation of the optimization methods in Shift and VERA to enable fully coupled ex-core calculations. Section 3. presents testing results with small ex-core problems including memory usage and computational performance. Section 4. wraps up the major accomplishments and planned future work.

2. OPTIMIZATION METHODS

This section presents the optimization methods integrated and implemented in Shift and VERA to enable fully coupled high-fidelity ex-core calculations.

2.1 MULTITHREADING

The basic implementation details of threaded Shift for ex-core calculations in VERA are discussed in this section. It builds upon the results and details in the CASL symposium paper [6].

The multithreading implementation that was previously implemented in Shift was integrated into VERA. This primarily entailed adding the *num_threads* option in the SHIFT block of the VERA input and ensuring the default was set to one thread per process (which means no multithreading). Assuming the executable was built with OpenMP enabled, the user only needs to specify more than one thread with this parameter to use multithreaded Shift with VERA. Only the MC transport solve of the Shift ex-core calculation is threaded; however, this will generally be the most time-consuming step of an MC calculation.

The threading implementation in Shift launches a set of message passing interface (MPI) tasks and then uses OpenMP to launch multiple threads of execution concurrently from each of these MPI tasks. Shift uses these threads to perform the transport of particles since each particle can be simulated independently. The threads launched under an MPI task use shared memory, enabling a simulation to use significantly less memory by launching fewer MPI tasks while running the transport solve on the same number of processors. This provides a method to reduce memory usage while keeping the MC transport time nearly constant.

There are a few caveats that should be kept in mind when using the multithreading implementation for ex-core calculations with VERA.

- Only the Shift transport solve part of the calculation is threaded; MPACT has multithreading implemented but does not scale well beyond 4 threads per processor, so parallel decomposition is the preferred method for computational performance; CTF is not threaded.
- Shift and MPACT are launched with the same MPI settings when running ex-core calculations in VERA.
- The user should take care to run calculations with the appropriate MPI options to see the full benefits of running with threads (e.g., ensure multiple threads are not running on the same core).
- Threading cannot be used simultaneously with domain decomposition in Shift as this is a limitation of the current methods implemented in Shift.

As mentioned in the above list, for ex-core calculations, the MPACT multithreading capability can be used to a limited capacity and CTF is not threaded, so the user must take this into account when running Shift with threading. However, threading can be used simultaneously with both MPACT and Shift. The performance of the Shift multithreading capability is demonstrated in Section 3.1.

2.2 MULTIGROUP CROSS SECTION REDUCTION

It was recognized during the performance optimization process that the time and memory to process the multigroup cross sections for the deterministic calculation were prohibitive when running large problems with unique pins. It should be noted that the multigroup cross sections generated for the MPACT solve are not used for the deterministic forward and adjoint calculations for several reasons: first, the MPACT and Shift geometries have unique meshes; second, MPACT does not generate multigroup cross sections for the ex-core regions.

The time to build the multigroup cross sections for calculating the deterministic solution can dominate the Shift calculation time for some problems. As mentioned in Section 1., when using unique pins, the multigroup cross sections for the deterministic calculation are typically generated for each unique material region in the core. However, since the deterministic solution is purely used to generate variance reduction parameters, it is expected to be low-fidelity and does not need to generate unique cross sections throughout the core to be effective.

The number of cross sections built for calculating variance reduction parameters was reduced by averaging the compositions, through volume fraction mixing, in three axial levels in each radial pin (bottom, fuel, and top). This substantially reduces the cross-section processing time. Note that this approximation is used only to generate forward or adjoint fluxes for purposes of hybrid variance reduction. These would not be the multigroup cross sections to use if a fully converged and accurate adjoint flux is needed. This feature is enabled by setting the *use_reduced_xs* parameter to **true** in the SHIFT block of the VERA input for Consistent Adjoint Driven Importance Sampling (CADIS) and Forward-Weighted CADIS (FW-CADIS) ex-core problems.

2.3 DOMAIN DECOMPOSITION

This section presents the implementation details and user guidance for domain-decomposed ex-core calculations with VERA. It builds upon the details in the CASL symposium paper [6].

2.3.1 Implementation

The first step in making VERA ex-core calculations domain decomposed was to incorporate the domain decomposition (DD) features already implemented in Shift. Shift's DD implementation uses multiple-set overlapping domain (MSOD) decomposition, which divides typical global quantities across geometric domains and then replicates these domains on sets for efficiency of the MC transport. Full details can be found in previously published articles [2][7]. Shift has the ability to decompose some of its source types, weight windows, some of its tally types, and particle transport, along with the typical decomposition of the deterministic forward and adjoint solutions.

After integrating the current Shift DD implementation into VERA, the pincell metadata decomposition when using unique pins was implemented in Shift and VERA. With this added feature, the following can be decomposed when running ex-core calculations:

- in-core pincell metadata (materials),
- weight windows,

- adjoint and forward deterministic calculations,
- mesh tallies, and
- particle transport.

The following are replicated (also called global quantities) when running ex-core calculations:

- geometry,
- ex-core materials,
- fission source for Shift fixed source calculation,
- cell tallies, and
- cylindrical mesh tallies.

All quantities transferred from MPACT to Shift are communicated between all MPACT processors and the Shift processors making up all blocks/domains on the first set. After transfer of these quantities, each domain communicates the transferred quantities to the corresponding domain on all sets.

One drawback of DD is the potential to impact computational performance due to increased communication and load imbalance among processors. To mitigate this issue, previous work under CASL investigated a strategy to determine the optimal domain decomposition for load imbalanced problems given a total number of processors [8]. In the near future, VERA could take advantage of this optimal decomposition algorithm that has been implemented in Shift.

The following are caveats for using the current domain-decomposed implementation in VERA.

1. The user must input the boundary mesh in the SHIFT block. This includes the following parameters: *x_bnd_mesh*, *y_bnd_mesh*, *z_bnd_mesh*, and *bc_bnd_mesh*. See the VERA user manual for further details about these parameters [5] and the example in Section 2.3.2.
2. The boundary mesh edges that intersect the active core should align with a pincell or assembly geometric boundary to best performance.
3. Unique pins must be enabled with the *create_unique_pins* set to **true** in the SHIFT block.

2.3.2 Guidance on Decomposing Domains

As mentioned in the previous section, to enable a domain-decomposed Shift ex-core calculation with VERA, the user must specify the boundary mesh and boundary conditions. As an example, to run a quarter-core problem on 16 Shift processors, the boundary mesh input might be the following.

```
bc_bnd_mesh  reflect vacuum reflect vacuum vacuum vacuum
x_bnd_mesh   50.0 80.0 125.0
y_bnd_mesh   50.0 80.0 125.0
z_bnd_mesh   -5.0 175.0 350.0
```

In this example, Shift is running on 8 blocks and 2 sets.

Domain decomposition could be used for any ex-core problem using unique pins. However, it is most useful for the following scenarios:

- coupling of fission source, isotopic compositions, and temperatures and densities;
- limited memory available on computing cluster (this will be true for most machines);
- full-core or quarter symmetry with unique pins.

To avoid gross load imbalance, the Shift domains should split the core with as close to an equal number of unique pincells per domain as possible. This includes the axial levels. As always, there is a balance that must be considered in giving up some computational performance to reduce memory per domain. In other words, the smaller each domain, the more communication between domains, and hence the more computational performance degradation.

2.4 FW-CADIS VARIANCE REDUCTION

The CADIS and FW-CADIS variance reduction methods [9, 10] are implemented in Shift, but until this work only CADIS was accessible through the VERA application programming interface (API) to Shift. Users can now also run hybrid MC transport calculations using the FW-CADIS method by setting *problem_mode* to **fwcadis** in the SHIFT block of the VERA XML file. Running in FW-CADIS mode requires the same hybrid calculation parameters that are needed for a CADIS run to be specified in the SHIFT block. These hybrid parameters are described in the ex-core VERA user manual [5].

The CADIS and FW-CADIS methods use the Denovo deterministic code [11] to generate variance reduction parameters. These parameters are then used to optimize the MC calculation for a response of interest, e.g., vessel fluence or detector response, resulting in a lower statistical uncertainty and significantly reducing MC computation time. The CADIS method uses an adjoint deterministic solve to generate parameters based upon optimization for a single detector or region, but the FW-CADIS method is needed to optimize for multiple regions or detectors. In other words, FW-CADIS is suited to simultaneously reducing variance in multiple regions, whereas CADIS is suited to reducing variance for a single region. The FW-CADIS method first uses a low-fidelity forward deterministic calculation to estimate the responses in the regions of interest indicated by the user. The following adjoint deterministic calculation is then run with the adjoint sources inversely scaled by the response estimates to generate the appropriate variance reduction parameters. This enables the FW-CADIS method to lower the statistical uncertainty for several different regions at the same time, whereas the CADIS method would disproportionately lower the uncertainty in regions with a higher response value.

In the current Shift implementation, the CADIS and FW-CADIS setup and solve sequences live in the general Shift frontend known as Omnibus [12]. These sequences use Denovo, which is where the deterministic solvers reside in Shift, for forward and adjoint deterministic calculations. This milestone has enabled VERA to use FW-CADIS for PWR ex-core calculations with multiple regions of interest which is especially needed for beyond the beltline calculations.

To enable FW-CADIS, a workflow change was implemented to take advantage of the Omnibus sequences through VERA. The primary modifications included the following.

- Integration of using the Omnibus front-end through VERA.

- A VERA model class was implemented in Shift and integrated into the process of running Shift through VERA. This was needed to align with the Omnibus sequences and has the additional benefit of enabling the VERA input geometry to be run directly through Omnibus for testing.
- The managers and sequences for running Shift through VERA have been modified to more closely resemble the Omnibus workflow, which reduces code duplication and simplifies the addition of features. This included pulling the building of the geometry and physics out of the VERA managers.
- Implementation of an Omnibus FW-CADIS sequence through VERA.
 - An FW-CADIS sequence for VERA was derived from the previously existing CADIS VERA sequence. It includes an additional initialization method where the forward Denovo manager is created and run to calculate a low-fidelity solution before calling the CADIS initialization to run the adjoint Denovo manager and generate the variance reduction parameters.
 - A Denovo source builder for FW-CADIS through VERA was implemented to take the solution from a forward Denovo calculation and create the adjoint source for the adjoint Denovo calculation. This is a modified version of the Omnibus Denovo source builder that has been updated to work with the VERA interface to Shift.

The CADIS and FW-CADIS implementations through VERA are both setup to optimize MC transport for cell tallies currently, but in the future this will be extended to allow the user to also optimize for mesAt this time, only using energy integrated responses to bias the adjoint source has been tested; however, the capability for the energy binned treatment is available in the code and can easily be exposed to the user.

The FW-CADIS implementation through VERA is demonstrated in Section 3.4 for an ex-core problem with multiple detectors.

3. RESULTS

This section presents the results of testing the optimization methods detailed in Section 2.. It is important to note that all of these results below transferred depleted isotopic compositions only tracked the short list of depleted nuclides in MPACT which are: ^{234}U , ^{235}U , ^{236}U , ^{238}U , ^{16}O , ^{135}Xe , and ^{10}B .

3.1 MULTITHREADING

The following sections present the memory savings and computational performance observed while running two small-core ex-core problems using multithreading, in Sections 3.1.1 and 3.1.2, as well as the results from a full-core subcritical Watts Bar Nuclear Plant Unit 1 (WBN1) problem, in Section 3.1.3.

3.1.1 5-mini Model

The performance testing with multithreading was first investigated using a "mini" version of CASL Progression Problem 5 [13] ("5-mini"). The 5-mini problem is composed of 7×7 assemblies laid out in a 9×9 core with a baffle, thin barrel, and a single vessel region. A raytrace of an axial slice through the geometry is given in Fig. 1. All of the simulations performed ran a full-core single state point at hot zero power (HZIP) without feedback.

To demonstrate Shift's multithreading capability when running through VERA, Shift was set to run the 5-mini problem in forward mode with 1×10^7 particle histories, non-unique pins, and the nuclide-based fission source provided by MPACT. Because feedback was not turned on and non-unique pins were used, the coupling between MPACT and Shift transferred only the pin-averaged fission source. This VERA ex-core calculation tallied flux and accumulated fluence in the reactor pressure vessel.




Figure 1. Axial slice of 5-mini problem geometry at $z=50$ cm. The colors in this figure represent the materials in each region.

To examine the memory advantages of multithreading, the 5-mini problem was run on a single node using a total of 32 processors with different numbers of threads on a small CASL Linux cluster with 36 2.10 GHz Intel Xeon CPUs and 126 GB of memory per node. MPACT uses 16 processors for all cases, while the number of Shift processes is varied with the number of threads to consistently run Shift on a total of 16 processors. For example, using 4 threads per process, Shift is run with 4 MPI tasks, MPACT uses 16 MPI tasks without multithreading, and MPI is invoked with `-np 20`. Table 1 gives the peak memory usage during the Shift transport calculation for each case as reported in the Shift output from a system call. The memory usage shows excellent scaling as the number of threads is increased up to 16 threads per Shift task.

Table 1. 5-mini peak memory usage during Shift transport with varying number of threads. All cases run MPACT and Shift on 16 processors each.

MPI tasks	Shift MPI tasks	# threads per Shift task	Memory [MB]	Ratio
32	16	1	9413	—
24	8	2	4759	0.51
20	4	4	2452	0.26
18	2	8	1284	0.14
17	1	16	686	0.07

3.1.2 Small Modular Reactor Model

VERA's performance on ex-core problems with multithreading was tested on a larger model representative of a small modular reactor (SMR). The SMR model has a 7×7 core made up of 17×17 assemblies with a baffle, barrel, vessel, and bioshield. The bioshield beyond the vessel is defined in a supplemental ex-core geometry file [5]. A raytrace of the geometry modeled with unique pins and quarter symmetry is given in Fig. 2.

VERA was used to run this SMR problem for a single state with thermal-hydraulic feedback on and a fully coupled ex-core calculation. The MPACT and CTF calculations converged with default settings, with MPACT using graph partitioning. Because the problem was fully coupled, the pin-wise compositions, temperatures, moderator densities, boron concentration, and fission source were communicated to Shift from MPACT once MPACT had completed the first state point calculation. The ^{235}U spatial source distribution transferred to Shift from MPACT, is shown in Fig. 3.

Shift ran with CADIS variance reduction to optimize for the bioshield, simulated 10^8 particle histories, and continuous energy (CE) pole data was used for cross-section broadening based on the transferred temperatures. Cylindrical tallies were set up in the supplemental ex-core file to tally neutron flux in the vessel and bioshield. A slice through the tally results at an elevation of 120 cm and overlaid on the problem geometry is shown in Fig. 4.

This quarter symmetry SMR problem was run with varying numbers of processors per node and Shift threads to examine the memory usage and computation time. All cases were run on the Compute and Data Environment for Science (CADES) Oak Ridge National Laboratory (ORNL) institutional cluster. CADES is a commodity Intel x86_64 cluster with 32 cores and 128 GB of RAM per node. Shift ran on 4 nodes for




Figure 2. Axial slice of the quarter symmetry SMR model at an elevation of 200 cm with unique pins. The colors represent the materials in each region. Because unique pins are enabled, the materials in each pin are defined as unique materials, which corresponds to each pin having different colors in this figure.




Figure 3. Quarter symmetry SMR ^{235}U fission source distribution. The NE quadrant of the geometry is shown here because Shift internally models this quadrant as a reflection of the SE quadrant.

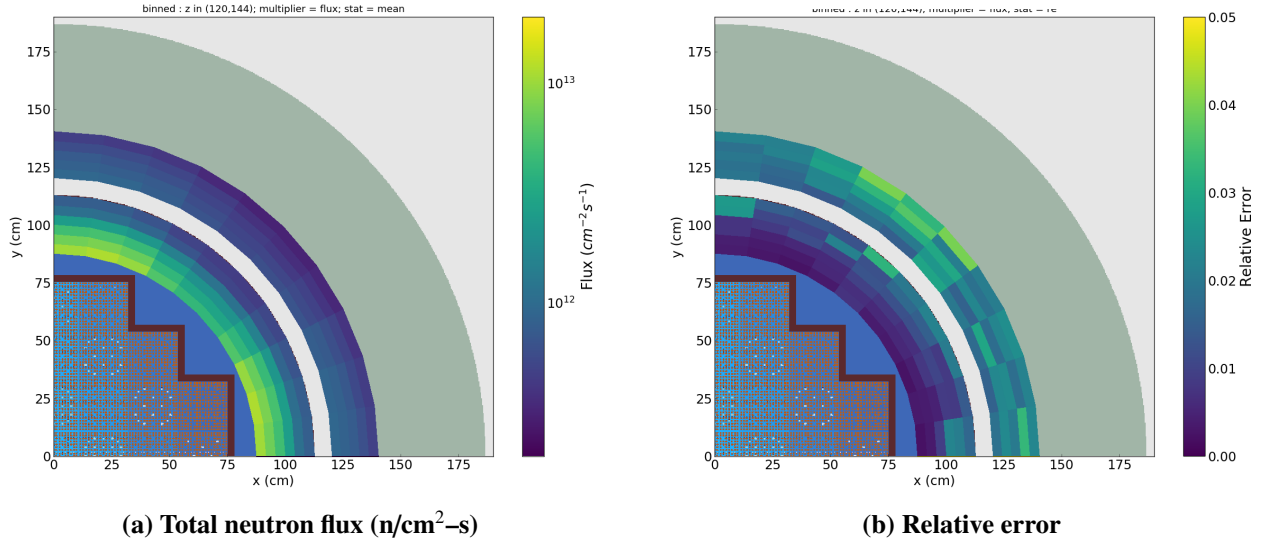


Figure 4. Total neutron flux and relative error in vessel and 20 cm of bioshield of quarter symmetry SMR at an elevation of 120 cm. The NE quadrant of the geometry is shown here because Shift internally models this quadrant as a reflection of the SE quadrant.

all cases, and the number of MPACT nodes was varied to maintain 16 MPACT processors across all cases. Note that unlike the smaller 5-mini problem in Section 3.1.1, the SMR model needed to be run across multiple nodes due to memory limitations on a single node. In the 5-mini problem, the number of MPACT tasks was kept constant while the number of Shift tasks and threads was varied on a single node. For this SMR problem, the number of MPI tasks per node was made consistent across all nodes, regardless of whether that node had MPACT or Shift tasks. This means that as the number of Shift threads was increased, the MPACT calculation ran on fewer processors per node and therefore used more nodes. This was necessary to account for the fact that Shift and MPACT are launched with the same MPI parameters when running through VERA.

Table 2 compares total computation time, MC transport time, and total Shift memory usage while running Shift on 4 nodes and reducing the number of MPI tasks while increasing the number of threads. MPACT is run with 16 processors for all cases, so the number of MPACT nodes is increased as the number of processors per node is decreased. The results in Table 2 demonstrate that running the MC calculation with more threads and fewer MPI tasks enables a significant reduction in memory while maintaining the MC computation time. Going from 2 threads per process to 4 reduced the Shift memory usage by a factor of 1.8, and going from 2 threads to 8 reduced memory usage by a factor of 3. The trade-off of running only Shift with threads is that MPACT will either take longer to run if keeping the number of nodes constant or will have to be run on additional nodes, as was done for this problem.

The impact of increasing the number of threads while maintaining a constant number of MPI tasks is examined in Table 3. Increasing threads demonstrates the improvement in computation time that can be gained using multithreading. This is particularly helpful for large problems that cannot otherwise fully use the processors on a node due to memory constraints. In Table 3, as the number of threads per process is doubled, the MC transport time is reduced by a factor of ~ 1.8 . This particular problem is too large to run

Table 2. SMR computation times and memory usage with Shift on 4 nodes and varying the number of MPI tasks and threads. All cases run MPACT on 16 cores.

MPI tasks	Proc. per node	Shift MPI tasks	Threads per task	Total time [s]	MC transport time [s]	Memory [GB]
80	16	64	2	5.44E+03	1.30E+03	375
48	8	32	4	6.42E+03	1.29E+03	203
32	4	16	8	8.52E+03	1.32E+03	123

Table 3. SMR computation times and memory usage with increasing the number of threads. All cases run MPACT and Shift with 16 MPI tasks.

MPI tasks	Proc. per node	Shift MPI tasks	Threads per task	Total time [s]	MC transport time [s]	Memory [GB]
32	4	16	1	1.51E+04	7.94E+03	145
32	4	16	2	1.11E+04	4.34E+03	130
32	4	16	4	9.62E+03	2.42E+03	127
32	4	16	8	8.52E+03	1.32E+03	123

with 32 processors per node, but running with 16 processors per node and 2 threads per process enables the MC transport to finish in nearly the same time as a calculation using 32 processors per node.

This SMR model was also run as a fully coupled, full-core problem. The multigroup cross-section reduction method described in Section 2.2 was used for this simulation to reduce the overall computation time. This problem had 1,283,735 material compositions in Shift and was run using 8 processors per node with MPACT and Shift each using 64 MPI tasks over 8 nodes. Shift ran with 4 threads per MPI task. The spatial distribution of the ^{235}U fission source is given in Fig. 5. A slice through the tally results in the vessel and first 20 cm of bioshield at an elevation of 120 cm and overlaid on the problem geometry is shown in Fig. 6.

The total Shift memory usage when running the full-core SMR model was 920 GB over 64 processors for an average of 14.375 GB per process or 115 GB per node. Note that this problem requires the use of fewer than 32 MPI tasks per node in order to fit in the 128 GB of memory on each node available on the CADES cluster. Without multithreading, three-quarters of the processors on each node would be left unused for the entire calculation. Multithreading enables the user to better use the resources available and reduce computation time. It should be noted that some computational platforms allow the user to supply an MPI machine file to specify configurations for each processing node. This would allow more flexibility in using threaded Shift but requires advanced user knowledge.

3.1.3 Full-Core Subcritical WBN1 Model

A subcritical full-core WBN1 ex-core model that was developed for validating source range detector responses during refueling as detailed in [14] was used with the new multithreading capability in VERA. The specific model run was for startup of cycle 8 during the last step of refueling and included a detailed






Figure 5. Full-core SMR ^{235}U fission source spatial distribution.

ex-core model.

A raytrace of the model geometry near the core mid-plane is given in Fig. 7. The flux and ^{235}U response were tallied in both source range detectors and all 8 power range detectors. The problem was run with CADIS variance reduction and optimized for the flux in the north and south source range detectors. VERA runs MPACT and then transfers the fission source distribution to Shift for MC transport. The problem was run on 115 nodes with 16 processors per node for a total of 1840 processors. The MPACT solution converged with default parameters on 90 nodes or 1440 processors. The Shift MC calculation ran on 25 nodes with 400 MPI tasks and 2 threads per task. This configuration enabled Shift to fully use the available processors on the node during the MC transport step of the calculation. Additional details about this case can be found in [15].



(a) Total neutron flux



(b) Relative error

Figure 6. Total neutron flux and relative error in vessel and 20 cm of bioshield of full-core SMR at an elevation of 120 cm.




Figure 7. Axial slice of the WBN1 model at an elevation of 200 cm. The colors represent the materials in each region.

3.2 MULTIGROUP CROSS SECTION REDUCTION

The cross-section reduction method was tested with the SMR ex-core problem described in Section 3.1.2 using CADIS variance reduction. This option is enabled by setting *use_reduced_xs* to **True** in the SHIFT block of the VERA XML file. Timing and memory results are compared against calculations without the cross-section reduction in Table 4. The cross-section reduction option reduces the Shift computation time by more than half for this case. A modest reduction in memory usage is also observed. A small difference in the adjoint flux is seen between using unique multigroup cross sections and this reduced option, but this does not seem to affect the ex-core quantities of interest. In other words, running with unique multigroup cross sections in each pin region is not needed to produce sufficient variance reduction parameters for ex-core quantities. Further studies are needed before this option is set as the default to use for hybrid ex-core calculations with VERA.

Table 4. Comparison of SMR computation times and memory usage with and without cross-section reduction for calculation of variance reduction parameters. All cases run MPACT on a single node and Shift on 4 nodes.

XS reduction	MPI tasks (MPACT/Shift)	Proc. per node	Shift threads per task	Total time [s]	Shift time [s]	Memory [GB]
Off	80 (16/64)	16	2	5.28E+03	4.38E+03	367
On	80 (16/64)	16	2	2.71E+03	1.81E+03	353
Off	40 (8/32)	8	4	7.20E+03	5.77E+03	203
On	40 (8/32)	8	4	3.33E+03	1.90E+03	183

3.3 DOMAIN DECOMPOSITION

The following test problems were run for DD testing. All were multi-state ex-core calculations. The purpose of these tests was to show the successful implementation and demonstration of domain-decomposed Shift ex-core calculations including computational performance. All of the tables shown in this section present the total Shift execution time, the Shift MC transport time, and the total memory usage over all Shift processors as reported by a system call on each node. Assessing memory usage accurately is difficult and future work will include investigating alternative ways to confirm the memory usage.

3.3.1 DD Memory Analysis

The purpose of DD MC is to reduce the memory usage per MPI process to enable the full problem model to fit in global memory. Domain replicated MC reproduces the entire problem domain including geometry, materials, tally regions, sources, and importance maps (for hybrid cases) on each processor and only decomposes the computational phase space across particle histories. In a hierarchical MPI/multi-threaded simulation, the problem can be decomposed into fewer MPI ranks per node, thus increasing the memory available per rank and preserving particle concurrency through the use of threads (this was shown in Sect. 3.1. However, for the largest, full core simulations with depleted fuel compositions at multiple state points, even this memory saving is unlikely to suffice. Shift’s spatial DD implementation in VERA decomposes materials, tally regions, and the importance maps such that larger spatial decompositions will ultimately result in a problem size per processor that fits on the available architecture. Because MC DD can be plagued by extreme load-balancing inefficiency [8], the objective is always to decompose spatially into as few domains as possible in order to fit on a given machine. Particle parallelization in DD is natural because particles are sourced uniquely on the domains where they are born within each set of decomposed blocks. The full DD capabilities in Shift are summarized in Refs. [2, 8].

To demonstrate the memory cost of the problem model in both domain replication (DR) and DD, a stand-alone unit-test was created that loads the problem model, metadata (reactor state description including materials), and geometry and implemented detailed memory diagnostics. For the full-core SMR model (see Sects. 3.1.2 and 3.3.3) the incremental, peak, memory usage is illustrated in Fig. 8. A quick examination of this figure shows that the peak in both DR and DD is nearly the same, around 2.2 GB. However, a closer look at the generator memory usage, where the material compositions are constructed, shows a decrease in memory use in the DD case. This is shown in more detail in Fig. 9. Examination of Fig. 9 indicates that the material compositions are clearly spatially decomposed. The reason the overall memory is not exactly 4× lower is due to the fact that some materials need to be replicated in multiple domains owing to the fact that the domain boundary splits geometric cells.

Shift’s current implementation of general geometry for ex-core problems requires unique cell constructions for each unique material region. In these problems, each material is unique in order to support multiple state-point calculations; thus, the number of geometric cells can grow exceedingly large. Furthermore, the material compositions, which in the DR case constitute ~ 12% of the memory in the full geometric model, are shown here for a fresh fuel case. In a multi-state case, the number of nuclides in each fuel region will be 40–50× larger due to the contributions from neutron activation, decay, and fission (fuel burnup and depletion). Thus, the full material description in a multi-state case will quickly reach and overwhelm the per-rank, constant, geometry model memory cost.




Figure 8. Peak memory consumption after generation of the SMR model, metadata, and geometry for DR (top) and DD (bottom) simulations. The metadata-builder and generator are class objects that create the metadata, compositions, and geometry objects, respectively.




Figure 9. Memory cost of the material compositions in the non-depleted fuel SMR problem for DR (top) and DD (bottom).

Nonetheless, the geometric model memory consumption is still too large when considering unique pins. A new geometric implementation is being developed in Shift that supports replicated geometric constructions with overlaid material id mappings (under a different research endeavour). This implementation should reduce the per-rank geometry model memory cost substantially in the future.

3.3.2 5-mini Model

This first test problem run with domain decomposition is the same 5-mini model used for the threading testing in Sect. 3.1.1. The purpose of showing all of the results below with this test problem is to do a deeper dive into the effects of the various relevant options when running a smaller and simple domain-decomposed ex-core problem. Different from the previously detailed 5-mini problem, this version has a second state point that includes depletion to investigate the effect of transferring pin-wise depleted isotopic compositions on memory and computation time.

First, this problem was run in forward mode with feedback and with unique pins; therefore, the pin materials and the MC particle transport were the two things decomposed in Shift. This also means the following were transferred from MPACT to Shift: pin-wise average fission source, boron concentration, and isotopic compositions. Although feedback was on, this first case purposely did not transfer the pin temperatures so a comparison could be made in the next set of results when transferring everything. This unique pin problem had 532,554 unique material compositions in Shift. Finally, the Shift decomposition used for this problem was intentionally not load balanced.

This problem was run on the ORNL CADES cluster with 1×10^6 particle histories and a default vessel fluence tally. Table 5 shows the memory usage, total computational time, and Shift transport solve times for both state points. Note that all of these configurations were run on 4 computational nodes with MPACT using 96 processors and CTF running on 69 of those processors.

There are several aspects of the results shown in Table 5 to note. Since the second state point includes depleted materials instead of the fresh materials used for the first state point, these computation times are longer and the memory usage is larger across all cases. Second, the DD calculations always take longer than the replicated calculations, as expected due to communication. Finally, there is only a small decrease in memory usage comparing the DD cases to the replicated cases for this particular problem. This result is not quite expected and could be due to several factors. First, the memory diagnostic being reported is deceiving, since it is a system call on each node. Second, the size of the geometry compared to the size of the compositions is much larger, and since the geometry is not decomposed, it is dominating the memory usage in all cases.

Next, this test problem was run with pin temperature transfer enabled. The *transfer* parameter in the SHIFT block was set to **all**, which added the transfer of fuel and/or depletable pin region temperatures and moderator densities. The *temp_transfer* parameter was set to **all** to also include the transfer of the pin clad temperatures. See the VERA ex-core manual for full details on these parameters and their options [5]. To reduce memory further, pole data was used for on-the-fly temperature broadening [5] [16] [17].

Table 6 shows the memory usage and computational performance when transferring all quantities that can currently be transferred from MPACT to Shift. For all of these results, MPACT ran on 96 processors with CTF running on 69 of those processors; therefore, a total of 4 nodes was used. First, when comparing the results in Table 6 to Table 5, there is not a significant increase in memory usage even though temperatures

Table 5. Domain-decomposed 5-mini forward case memory usage and computation times at each state point with isotopic transfer. MPACT used 96 processors.

Shift Procs	Decomp [X×Y×Z]	Shift Total time [s] St1/St2	Shift Transport time [s] St1/St2	Memory [GB] St1/St2
1	Replicated	391/456	381/447	5.9/8.2
2	Replicated	218/250	208/241	11.3/15.0
4	Replicated	127/142	117/133	22.5/30.0
8	Replicated	78/91	66/80	45.3/60.0
2	2×1×1	348/396	339/389	11.1/15.0
4	2×2×1	312/351	302/342	21.2/28.4
8	2×2×2	235/279	225/270	41.8/55.0
16	2×2×4	201/195	169/185	83.1/109.1

are now being transferred. This is one of the main advantages of using pole data versus the typical up-front temperature broadening of CE cross section data. Second, there is an overall increase in computation time between these two tables, which can mostly be attributed to doing on-the-fly doppler broadening.

Table 6. Domain-decomposed 5-mini forward case memory usage and computation times at each state point transferring all. MPACT used 96 processors.

Shift Procs	Decomp [X×Y×Z]	Shift Total time [s] St1/St2	Shift Transport time [s] St1/St2	Memory [GB] St1/St2
1	Replicated	685/773	636/748	6.0/8.5
2	Replicated	343/414	329/391	11.4/15.3
4	Replicated	194/235	179/213	22.7/30.6
8	Replicated	117/146	101/123	45.3/61.0
2	2×1×1	545/662	533/643	11.4/15.2
4	2×2×1	516/603	504/585	21.6/28.6
8	2×2×2	380/489	368/471	41.9/55.4
16	2×2×4	277/310	263/292	83.3/108.8

Finally, this 5-mini problem was run in four different scenarios with CADIS enabled and the multigroup cross-section reduction technique described in Sec. 2.2. As with the previous cases, this problem was run with transferring the pin isotopic compositions only and also with transferring everything. All of the DD cases were run with the same non-load balancing as the previous cases and with proper load balancing (denoted by LB in the following tables).

First, memory usage and computation times for each state point are shown in Table 7 when only transferring the fission source and pin isotopic compositions (analogous to Table 5). A reminder when looking at the Denovo decomposition column, that Denovo, which performs the adjoint calculation for CADIS, uses Koch-Baker-Alcouffe (KBA) decomposition so it is not decomposed axially [18]. The most

obvious and expected thing to point out is that the total Shift computation times for the first state point are all longer than the previous cases because this includes the CADIS part of the calculation. The Shift transport times reported are only for the MC transport solve. Second, the memory usage for the $2 \times 2 \times 2$ non-load balanced and load-balanced cases is less than the replicated counterpart, which is starting to show the advantage of using DD over replicated domains. Finally, the computation times for the load-balanced DD cases are all smaller than the non-load balanced cases which shows that load balancing the Shift domains is important for the MC transport part of the calculation.

Next, Table 8 shows the memory usage and computation time for each state point when transferring everything for these CADIS simulations. There are two main things to note: first, the use of the pole data and the reduced multigroup cross sections help keep the memory usage on par with the values in Table 7; second, when comparing the load-balanced DD runs to the previous results, we see that adding the transfer of pin-wise temperature has not significantly increased the Shift computation times.

Table 7. Load-balanced versus not load-balanced domain-decomposed 5-mini CADIS case memory usage and computation times for each state with isotopic transfer. MPACT used 96 processors and LB refers to load-balanced results.

Shift Procs	Shift Decomp [X×Y×Z]	Denovo Decomp [X×Y]	Shift Total time [s] St1/St2	Shift Transport time [s] St1/St2	Memory [GB] St1/St2
1	Replicated	–	2938/324	298/308	8.4/9.7
2	Replicated	–	1544/187	174/174	13.0/16.3
4	Replicated	–	858/120	97/108	23.8/30.6
8	Replicated	–	612/78	53/66	46.7/60.3
2	$2 \times 1 \times 1$	2×1	1791/466	401/453	13.0/16.0
4	$2 \times 2 \times 1$	2×2	1088/390	331/378	22.7/28.8
8	$2 \times 2 \times 2$	4×2	806/303	244/293	43.0/55.1
16	$2 \times 2 \times 4$	4×4	500/207	176/196	83.9/109.5
2	$2 \times 1 \times 1$ LB	2×1	1617/302	264/295	12.7/15.9
4	$2 \times 2 \times 1$ LB	2×2	905/175	160/167	22.6/29.0
8	$2 \times 2 \times 2$ LB	4×2	666/177	115/169	43.0/55.0
16	$2 \times 2 \times 4$ LB	4×4	403/113	81/105	84.2/109.2

Overall, the main takeaways from this DD analysis of problem 5-mini are: first, the user needs to be aware of the pincell load balancing when defining their boundary mesh domains as it will affect memory usage and computation times; and, second, running properly decomposed and fully coupled CADIS ex-core calculations is now achievable. For the analysis performed in this section, the number of particles simulated was kept the same for the forward and CADIS calculations. However, in practice, fewer particles can be transported with CADIS while achieving MC uncertainty for ex-core quantities of interest on the same order as forward calculations. This would also further reduce the computation time of DD CADIS calculations.

Table 8. Load-balanced domain-decomposed 5-mini CADIS case memory usage and computation times at each state transferring all. MPACT used 96 processors and LB refers to load-balanced results.

Shift Procs	Shift Decomp [X×Y×Z]	Denovo Decomp [X×Y]	Shift Total time [s] St1/St2	Shift Transport time [s] St1/St2	Memory [GB] St1/St2
1	Replicated	–	3091/546	477/518	8.5/10.2
2	Replicated	–	1652/312	267/283	12.9/16.6
4	Replicated	–	899/190	141/165	23.9/31.2
8	Replicated	–	645/123	82/98	46.5/61.8
2	2×1×1	2×1	1778/514	388/485	12.7/16.1
4	2×2×1	2×2	1098/425	335/405	22.6/29.3
8	2×2×2	4×2	822/339	255/321	43.2/56.2
16	2×2×4	4×4	510/229	181/209	84.1/110.0
2	2×1×1 LB	2×1	1625/328	277/317	12.8/16.0
4	2×2×1 LB	2×2	908/192	164/181	22.6/29.4
8	2×2×2 LB	4×2	676/183	119/170	42.9/54.9
16	2×2×4 LB	4×4	406/123	83/112	83.8/110.0

3.3.3 SMR Model

This second test problem is the same SMR model used for the multithreading testing in Sect. 3.1.2. As with the previous test problem, a second state point to include depletion was added to this model. These simulations were run on the same compute cluster using 16 processors for the MPACT and CTF part of the calculation and varying numbers of processors for the Shift part of the calculation. All of these calculations were performed with using only 8 processors per node in order to be sure the problem fit in memory.

First, Table 9 and Table 10 show memory usage and computation time when running this SMR problem in forward mode with and without transferring temperatures. The main takeaway from these results is that the DD calculations run and produce comparable results (not shown here). These results certainly do not show the advantage of running with DD since the replicated problem can still be fit in memory and run in less time. Further investigation into the optimal decomposition for this problem is needed.

Table 11 shows the computation time and total memory usage for this SMR problem run with CADIS. The only conclusions that can be drawn from these results are that more sets should be run for the DD case to reduce the computational time. The amount of memory needed for this geometry, which is replicated, is large compared to the size of the unique compositions therefore memory comparisons are not useful for this problem.

Finally, the fission source and the tallied neutron flux are shown from the DD CADIS calculation. Figure 10 shows the ^{235}U fission source for each state point and Figure 11 shows the total neutron flux and the associated relative error for each state point. These can be compared to the SMR results in the multithreading section to see that CADIS produces much lower relative error and the same fission source as expected.

Table 9. Domain-decomposed SMR forward case memory usage and computation times at each state with isotopic composition transfer. MPACT used 16 processors.

Shift Procs	Decomp [X×Y×Z]	Shift Total time [s] St1/St2	Shift Transport time [s] St1/St2	Memory [GB] St1/St2
64	Replicated	1921/2112	1849/2084	375/474
2	2×1×1	2307/2554	2287/2535	368/472
4	2×2×1	2534/2602	2518/2586	365/469
8	2×2×2	3331/3327	3283/3312	370/467
16	2×2×4	4755/4790	4742/4779	389/492
32	2×2×8	6361/6445	6343/6433	411/521


Table 10. Domain-decomposed SMR forward case memory usage and computation times at each state transferring all. MPACT used 16 processors.

Shift Procs	Decomp [X×Y×Z]	Shift Total time [s] St1/St2	Shift Transport time [s] St1/St2	Memory [GB] St1/St2
64	Replicated	3434/3226	3363/3199	390/488
2	2×1×1	4071/3801	4047/3781	384/492
4	2×2×1	4407/4185	4389/4169	382/484
8	2×2×2	5643/5857	5593/5800	392/482
16	2×2×4	8150/8452	8130/8435	418/513
32	2×2×8	10821/11444	10801/11425	450/554


As mentioned at the beginning of this section, only the short list of depleted nuclides was tracked and transferred between MPACT and Shift. If however, these computations used the full nuclide depletion list, the size of the compositions compared to the geometry should be much larger and therefore domain decomposition would be needed to fit these calculations in memory. Most ex-core calculations do not need this full nuclide list, but further investigation into when and what problems need domain decomposed Shift is warranted.

Table 11. Domain-decomposed SMR CADIS case memory usage and computation times at each state. MPACT used 16 processors.

Shift Procs	Shift Decomp [X×Y×Z]	Denovo Decomp [X×Y]	Shift Total time [s] St1/St2	Shift Transport time [s] St1/St2	Memory [GB] St1/St2
64	Replicated (transfer iso)	–	5742/4184	3719/4157	411/503
64	Replicated (transfer all)	–	2923/3189	2753/3158	392/492
64	2×2×4 (4 sets, transfer iso)	8×8	7380/9791	7197/9772	416/512
64	2×2×4 (4 sets, transfer all)	8×8	7370/9915	7181/9892	410/507



(a) State 1 fission source



(b) State 2 fission source

Figure 10. Quarter symmetry SMR ^{235}U fission source distribution for each state point of DD calculation. The NE quadrant of the geometry is shown here because Shift internally models this quadrant as a reflection of the SE quadrant.

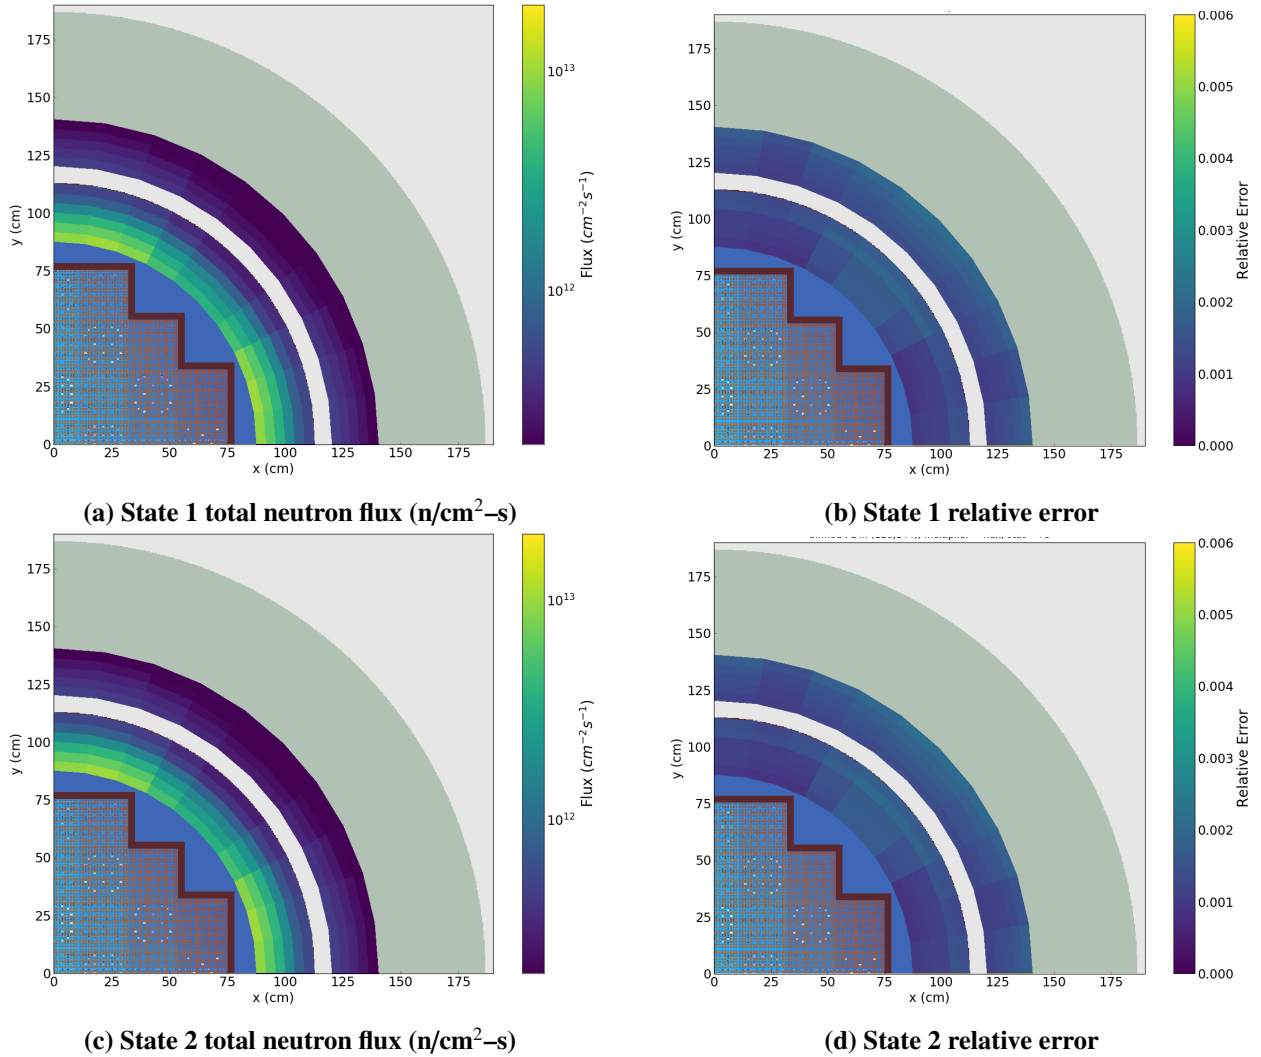


Figure 11. Total neutron flux and relative error in vessel and 20 cm of bioshield of quarter symmetry SMR at an elevation of 120 cm run with DD and CADIS. The NE quadrant of the geometry is shown here because Shift internally models this quadrant as a reflection of the SE quadrant.

3.4 FW-CADIS

FW-CADIS is enabled by setting *problem_mode* to **fwcadis** in the SHIFT block of the VERA XML file. Since FW-CADIS mode is a hybrid calculation, it requires the same parameters to be specified in the SHIFT block that are needed for a CADIS run. These hybrid parameters are described in the ex-core VERA user manual [5]. As with CADIS mode, the user must specify the tallies that they would like to optimize.

The FW-CADIS implementation through VERA was tested using a modified version of the SMR ex-core model described in Section 3.1.2. VERA has the ability to automatically generate the supplemental ex-core file via parameters defined in the standard VERA input, as demonstrated in [5, 6]. This feature was used to generate an ex-core file with a bioshield and 3 cylindrical ex-core detectors located in wedge-shaped wells within the bioshield. Detectors 1 and 2 are centered on the core midplane at 30° and 60° as shown in Fig. 12a. These detectors each have two axial tally regions, an upper and a lower, with a height of 52.4 cm . Fig. 12b shows detector 3, which is located near the top of the core, has a deeper well, and has an axial height of 26.2 cm . All three detectors are modeled as a void cylinder with a stainless steel shell. The hybrid transport calculations shown in this section optimized for the flux in all three of these detectors.

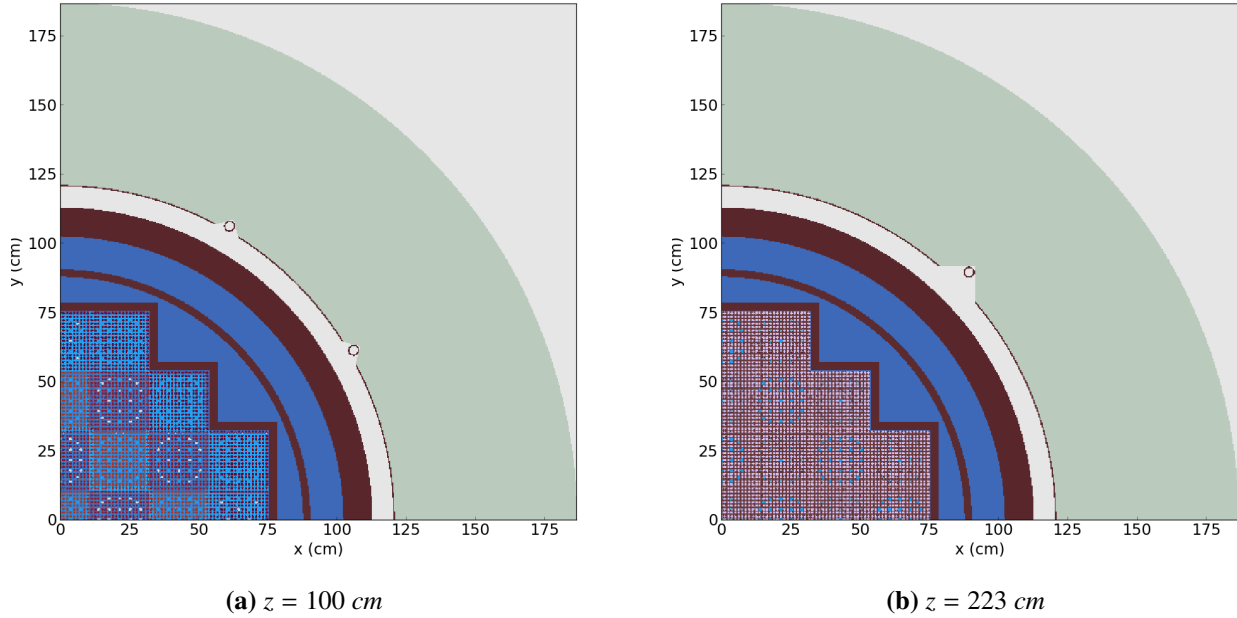


Figure 12. Axial slice of the quarter symmetry SMR model with ex-core detectors and nonunique pins at elevations of (a) 120 cm and (b) 123 cm. The colors represent the materials in each region.

All calculations were run on the CADES computing cluster described in Section 3.1.2. The geometry was run with nonunique pins, so only the fission source was transferred from MPACT to Shift. All simulations ran across 4 nodes with 64 MPACT processors and 64 Shift processors. Shift transported 10^8 particles and Denovo ran with an 8×8 decomposition. Simulations with CADIS mode and FW-CADIS mode were run for comparison.

A CADIS mode simulation with this geometry treats the 3 detectors as a single tally to optimize. In this case, the two tallies at the core midplane converge more quickly than the third elevated tally. A simulation

using FW-CADIS uses an initial forward deterministic solution to estimate the response at all tallies that are being optimized, as shown in Fig. 13. This response is then used to scale the adjoint source used in the adjoint deterministic calculation. In this example, the response in the third detector should be significantly lower than the response in the two midplane detectors.




Figure 13. Quarter symmetry SMR Denovo total flux at elevations of (a) $z = 120$ cm and (b) $z = 230$ cm.

The adjoint source spatial distribution used for the Denovo adjoint calculation when running in CADIS mode is shown in Fig. 14 overlaid with the geometry at the elevations of the detectors. The corresponding distribution when running in FW-CADIS mode is given in Fig. 15. Comparing the adjoint source strength between the two methods, the FW-CADIS simulation increases the source strength in the smaller detector near the top of the core relative to the two detectors at the midplane, because this detector has a lower response. The adjoint function spatial distribution resulting from the FW-CADIS adjoint sources is shown in Fig. 16. The calculated variance reduction parameters for the FW-CADIS calculation are more effective than the CADIS method at obtaining low variance in all requested detectors.

The detector tally results for CADIS and FW-CADIS modes are given in Table 12. As expected, the FW-CADIS calculation results in a significant improvement in the detector 3 relative error for the same number of particle histories; in this case, the relative error was cut in half compared to the CADIS solution. A small increase in the lower tallies relative errors for detectors 1 and 2 is observed due to the FW-CADIS variance reduction parameters driving more particles towards detector 3.

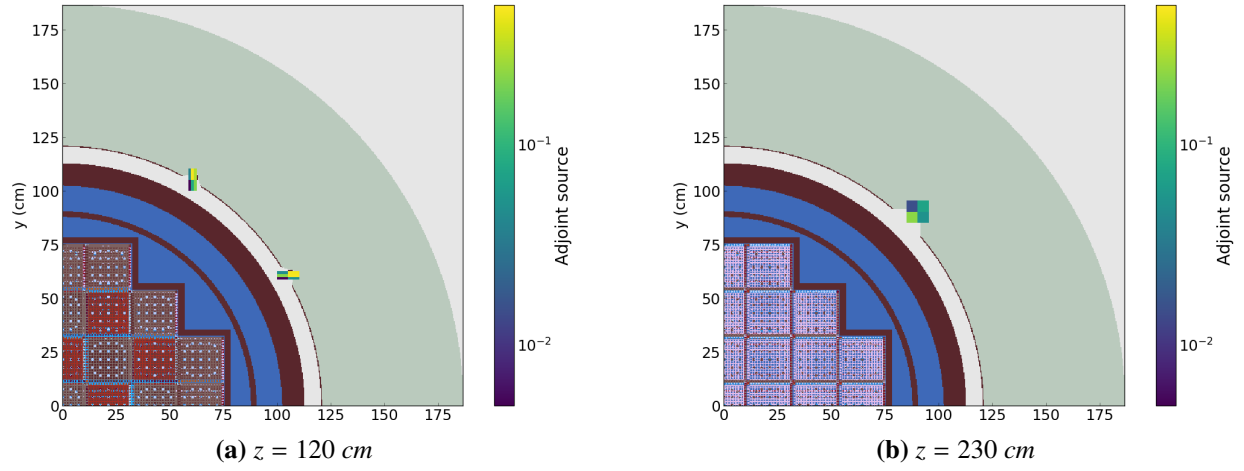


Figure 14. Adjoint source distribution for CADIS mode with quarter symmetry SMR problem at elevations of (a) 120 cm and (b) 230 cm.

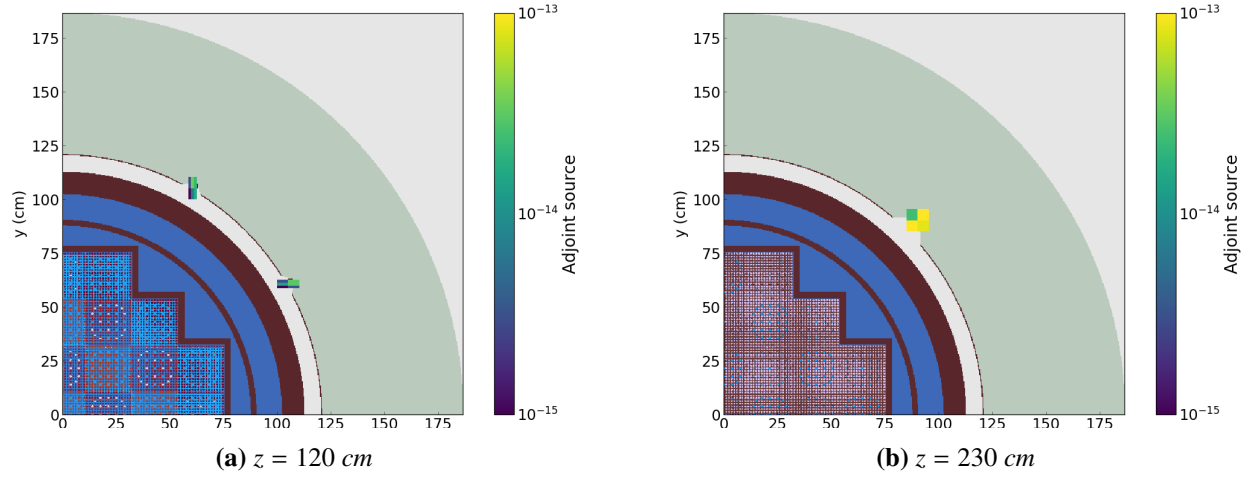




Figure 15. Adjoint source distribution for FW-CADIS mode with quarter symmetry SMR problem at elevations of (a) 120 cm and (b) 230 cm.

Table 12. Detector flux tally results for SMR ex-core geometry.

Detector	CADIS		FW-CADIS	
	Flux [$\text{cm}^{-2}\text{s}^{-1}$]	RE [%]	Flux [$\text{cm}^{-2}\text{s}^{-1}$]	RE [%]
1 (lower)	3.78E11	0.25	3.76E11	0.29
1 (upper)	3.12E11	0.26	3.12E11	0.26
2 (lower)	3.78E11	0.28	3.78E11	0.32
2 (upper)	3.12E11	0.29	3.11E11	0.27
3	4.01E10	0.99	4.00E10	0.49



(a) $z = 120 \text{ cm}$



(b) $z = 230 \text{ cm}$

Figure 16. Quarter symmetry SMR adjoint function spatial distributions at elevations of (a) $z = 120 \text{ cm}$ and (b) $z = 230 \text{ cm}$.

4. SUMMARY

In summary, the optimization of Shift and VERA for ex-core analysis was demonstrated through enabling several new features, including multithreading, domain decomposition, and FW-CADIS, to reduce the memory footprint and enhance the computational performance of VERA. With these optimization methods, fully coupled ex-core calculations can now be performed on moderately sized high performance computing (HPC) machines. Fully coupled in this context means the transfer of moderator boron concentration, pin-wise fission source, depleted compositions, temperatures, and moderator densities from MPACT (with CTF) to Shift. Without these optimization methods, many processors on a compute node would be left unused for the entire ex-core calculation. Therefore, these methods enable the user to better use the resources available and reduce computation time.

4.1 FUTURE WORK

Future work will involve additional testing of these features with larger problems such as AMA Problem 9. Additional options for the new FW-CADIS capability will also be added, including optimization for mesh tallies and exposing the binned energy treatment to the user. The DD capability could also benefit from the optimal decomposition algorithm implemented in Shift as part of previous work under CASL.

References

- [1] John A. Turner, Kevin Clarno, Matt Sieger, Roscoe Bartlett, Benjamin Collins, Roger Pawlowski, Rodney Schmidt, and Randall Summers. The Virtual Environment for Reactor Applications (VERA): Design and architecture. *Journal of Computational Physics*, 326:544–568, December 2016.
- [2] T. M. Pandya, S. R. Johnson, T. M. Evans, G. G. Davidson, S. P. Hamilton, and A. T. Godfrey. Implementation, Capabilities, and Benchmarking of Shift, a Massively Parallel Monte Carlo Radiation Transport Code. *J. Comp. Phys.*, 308:239–272, 2016.
- [3] Daniel Jabaay. CASL MPACT User’s Manual. Technical Report CASL-U-2019-1873-000, CASL, July 2019.
- [4] R. Salko, M. Avramova, A. Wysocki, J. Hu, A. Toptan, N. Porter, T. Blyth, C. Dances, A. Gomez, C. Jernigan, and J. Kelly. CASL CTF User’s Manual. Technical Report CASL-U-2019-1885-000, CASL, March 2019.
- [5] Tara Pandya, Thomas Evans, Katherine Royston, Kevin Clarno, Benjamin Collins, Shane Stimpson, and Shane Henderson. Excore Radiation Transport Modeling with VERA User Manual. Technical Report CASL-U-2018-1556-002, CASL, September 2019.
- [6] Tara Pandya, Katherine Royston, and Thomas Evans. Utilizing Shift in VERA for Ex-Core Calculations. In *CASL Symposium*. American Nuclear Society, 2020.
- [7] John C. Wagner, Scott W. Mosher, Thomas M. Evans, Douglas E. Peplow, and John A. Turner. Hybrid and Parallel Domain-Decomposition Methods Development to Enable Monte Carlo for Reactor Analyses. *Prog. Nucl. Sci. Technol.*, 2(1):815–820, 2011.
- [8] J. Austin Ellis, Thomas M. Evans, Steven P. Hamilton, C.T. Kelley, and Tara M. Pandya. Optimization of Processor Allocation for Domain Decomposed Monte Carlo Calculations. *Parallel Comput.*, 87:77–86, September 2019.
- [9] John C. Wagner and Alireza Haghighat. Automated variance reduction of Monte Carlo shielding calculations using the discrete ordinates adjoint function. *Nuclear Science and Engineering*, 128(2):186–208, 1998.
- [10] John Wagner, Douglas Peplow, and Scott Mosher. FW-CADIS Method For Global And Regional Variance Reduction Of Monte Carlo Radiation Transport Calculations. *Nuclear Science and Engineering*, 176(1), January 2014.
- [11] T. Evans, A. Stafford, R. Slaybaugh, and K. Clarno. Denovo—A new three-dimensional parallel discrete ordinates code in SCALE. *Nucl. Technol.*, 171:171–200, 2010.
- [12] S. R. Johnson, T. M. Evans, G. G. Davidson, S. P. Hamilton, and T. M. Pandya. Exnihilo transport code manual. Technical Report Release 6.2.0 (Dev), Oak Ridge National Laboratory, 2017.
- [13] A. T. Godfrey. VERA Core Physics Benchmark Progression Problem Specifications. Technical Report CASL-U-2012-0131-004, CASL, August 2014.
- [14] A. T. Godfrey, E. E. Davidson, G. Wolfram, B. Collins, C. Gentry, G. Ilas, S. Palmtag, T. Pandya, and K. Royston. Watts bar unit 1 source range detector response validation during refueling. Technical Report CASL-U-2018-1561-000, CASL, 2018.

- [15] Katherine Royston and Tara Pandya. Full Core Subcritical Multiplication Ex-Core Calculations with Shift. Technical Report CASL-U-2020-1953-000, CASL, May 2020.
- [16] Shichang Liu, Xingjie Peng, Colin Josey, Jingang Liang, Benoit Forget, Kord Smith, and Kan Wang. Generation of the windowed multipole resonance data using Vector Fitting technique. *Annals of Nuclear Energy*, 112:30–41, 2018.
- [17] C. Josey, P. Ducru, B. Forget, and K. Smith. Windowed multipole for cross section Doppler broadening. *Journal of Computational Physics*, 307:715–727, 2016.
- [18] R. S. Baker and K. R. Koch. An Sn algorithm for the massively parallel CM-200 computer. *Nuclear Science and Engineering*, 128:312–320, 1998.