CASL-U-2015-0170-000

# SHIFT: A Massively Parallel Monte Carlo Radiation Transport Package

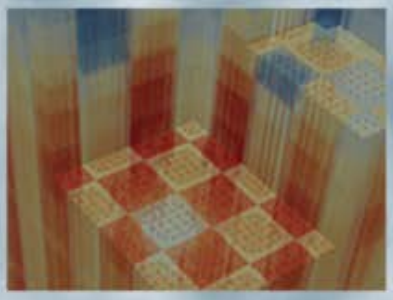

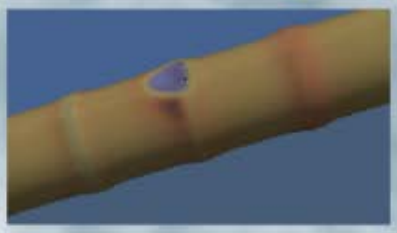

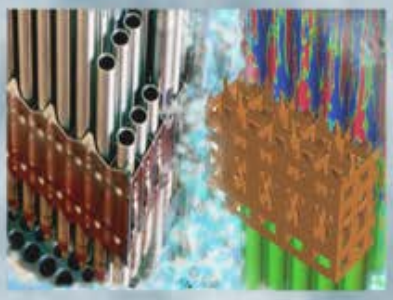

Tara M. Pandya, Seth R. Johnson, Gregory G. Davidson, Thomas M. Evans, and Steven P. Hamilton
**Oak Ridge National Laboratory**

April 19, 2015

U.S. DEPARTMENT OF ENERGY | Nuclear Energy
CASL-U-2015-0170-000

*ANS MC2015 - Joint International Conference on Mathematics and Computation (M&C), Supercomputing in Nuclear Applications (SNA) and the Monte Carlo (MC) Method* · Nashville, Tennessee · April 19–23, 2015, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2015)

# SHIFT: A MASSIVELY PARALLEL MONTE CARLO RADIATION TRANSPORT PACKAGE

**Tara M. Pandya, Seth R. Johnson, Gregory G. Davidson, Thomas M. Evans, and Steven P. Hamilton***
Oak Ridge National Laboratory
1 Bethel Valley Rd., Oak Ridge, TN, 37831 U.S.A.
pandyatm@ornl.gov; johnsonsr@ornl.gov; davidsongg@ornl.gov; evanstm@ornl.gov; hamiltonsp@ornl.gov

## ABSTRACT

This paper discusses the massively parallel Monte Carlo radiation transport package Shift, developed at Oak Ridge National Laboratory and it reviews the capabilities, implementation, and parallel performance of this code package. This code package is designed to scale well on high performance architectures. Scaling results demonstrate very good strong and weak scaling as applied to LWR analysis problems. Also, benchmark results from various reactor problems show that Shift results compare well to other contemporary Monte Carlo codes and experimental results.

*Key Words*: massively parallel, Monte Carlo, radiation transport, Shift

## 1 INTRODUCTION

Shift is a new Monte Carlo (MC) radiation transport code developed for flexible, fast, and accurate transport solutions in a wide variety of application areas. This paper discusses the development, implementation, and verification and validation of Shift for nuclear reactor analysis. Shift serves as the MC package in Exnihilo, a radiation transport code suite that also contains the Denovo deterministic transport package [1, 2]. Shift is designed and optimized for performing MC radiation transport calculations on current and near-future computing architectures, but it can also run most radiation transport problems on laptops or small clusters.

## 2 FEATURES OF SHIFT

### 2.1 Geometry Packages

Transport in Shift is performed using a generic tracking interface, enabling multiple geometry types to be implemented modularly in the code. Shift currently has full support for three geometry types, transporting on the geometry and interpreting the compositions in the file. The first geometry is the Reactor Toolkit (RTK) geometry, a high-speed internal package that models light water reactors (LWRs) for the Consortium for Advanced Simulation of LWRs (CASL) applications. Shift also has support for KENO-VI and KENO-V geometry through the Atlas package in SCALE [3]. Finally, Shift can read materials from and transport on MCNP5 geometry [4] via the Lava library used by ADVANTG [5].

### 2.2 Physics Capabilities and Features

The modular design of Shift allows multiple independent physics implementations, which currently include a separate multigroup (MG) and continuous-energy (CE) physics. Both of these support neutron, photon, and coupled neutron-photon transport.

There are two MG physics packages implemented in Shift. The first is an implementation of the SCALE MG physics algorithm that supports fully anisotropic scattering through $P_N$ moment expansions. Multigroup cross sections can be defined by manually specifying the cross sections in an input file, or they can be automatically built using the compositions read from the physics input. These MG cross sections can be generated directly using SCALE's XSProc module. Shift provides a full interface to the SCALE material specifications so the default infinite homogeneous medium treatment and advanced resonance self-shielding models can be applied when generating MG cross sections [3, 6]. The second is a simple MG physics implementation limited to $P_0$ scattering that is primarily used for code implementation testing.

The SCALE continuous-energy (SCE) physics module in Shift is an implementation of the physics used in CE-KENO. It uses the ENDF [7] data processed through AMPX [8] that is included with SCALE. Besides being a physics engine, this module also provides an efficient C++ and Python interface to underlying SCALE CE data. One difference between the SCE physics module in Shift and the CE physics in KENO is the treatment of fission site sampling. In Shift, the SCE physics module samples the number of fission sites produced at every collision, regardless of the type of collision. More details about fission site sampling in Shift are discussed later.

### 2.3 Source Definitions

As with other components of Shift, the source implementations are modular and easily extensible. There are a few basic separable sources implemented for fixed-source problems, two simple fission sources implemented for the initial source in *k*-eigenvalue problems, and a fixed fission source.

#### 2.3.1 Fixed sources

The fixed source definitions in Shift must be separable in space, angle, and energy. Existing angular distributions include isotropic and monodirectional. The supported energy distributions in Shift include histograms (lower energy bounds and average strength in each bin), line sources (energies and relative intensities), and a Watt fission spectrum. The separable spatial shapes implemented are a point, a rectangular box, a regular cylinder or cylindrical shell, and a sphere or spherical shell. All of these spatial distributions are uniform. Shift can also use a precomputed fission source spatial distribution that is specifically designed to enable fast coupled neutron-photon shielding analyses. This source is based on a fission mesh tally saved from a prior Shift run and particles are sampled according to the fission strength in each mesh cell with rejection according to whether the region contains fissionable material.

#### 2.3.2 Fission sources

As with other *k*-eigenvalue codes, Shift supports two fission sources: an initial distribution of first generation neutrons, and a fission site source for initializing subsequent generations. The initial

fission source in Shift samples uniformly inside a box, rejecting on nonfissionable materials. This fission source has an isotropic angular distribution and the starting energy distribution is determined based on the physics: SCE physics uses a Watt U-235 fission spectrum, and MG physics uses the fission spectrum $\chi$ of the birth material.

The fission source for subsequent generations in Shift uses fission sites sampled from collisions in the prior neutron generation. The fission site contains state data for the particle, e.g. exiting energy and fission site location.

### 2.4 Tallying Capabilities and Features

Shift supports tallying interfaces for several transport events: particles being born, streaming inside a single material, undergoing a collision, inducing a fission, and being killed (by roulette, absorption, leaving the geometry, etc.). Shift uses history-based statistics (as opposed to batch statistics) to determine the accuracy of the tally estimators. The calculated energy-binned and energy-integrated means and estimated variances are reported to the user. The following general tallies and diagnostics are available:

- Cartesian mesh tally – an overlaid Cartesian mesh on which any path length tally can be accumulated;
- cell union tally – geometry cells on which any path length tally can be accumulated;
- collision diagnostic tally – number of collisions per history as function of material, nuclide, and reaction;
- cylindrical mesh tally – an overlaid, translated, and rotated cylinder broken into $(r, z, \theta)$ mesh cells on which any path length tally can be accumulated;
- depletion tally – a fine-group flux tally without variance data;
- path length estimator for $k_{\text{eff}}$ – a path length estimate of the eigenvalue in a cycle;
- history diagnostic tally – detailed particle history event information written to an HDF5 file;
- Shannon Entropy tally – spatially binned source particle distribution that serves as a statistical metric for measuring fission source convergence [9]; and
- source diagnostic tally – source particle density binned in spatial cells.

These tallies can be energy binned and can tally a response definition or a given reaction from a set of supported reactions. Any number of the following multipliers can be tallied in Shift: particle flux, reaction rates, heating rates using multigroup kinetic energy released per unit mass (KERMA) factors, fission heat production rates using fission $Q$ values from SCALE data, and arbitrary energy- and particle-dependent responses input by the user. A serial sensitivity tally capability in Shift is also under development.

### 2.5 Hybrid Capabilities and Variance Reduction

The standard variance MC reduction techniques are implemented in Shift. It uses implicit capture by default with Russian roulette turned on. Shift also implements spatially decomposed, energy-dependent weight windows. These weight windows can be explicitly read from a file or directly assigned from an adjoint Denovo solution as an importance map. A normalization factor is used to make these weight windows "consistent" with the biased source. This consistency condition

between weight windows and biased sources was first shown in the Consistent Adjoint-Driven Importance Sampling methodology [10].

Shift implements a biased source that ensures full consistency given a set of weight of windows and a source definition. Shift creates a space- and energy-dependent cumulative density function (CDF) for directly sampling a spatial mesh cell and energy group. Rejection sampling is used to determine whether the sampled birth point is within the true source region. During the biased source construction, Shift calculates the estimated response,

$$R = \iint q(\boldsymbol{r}, E)\phi^{\dagger}(\boldsymbol{r}, E)\,\mathrm{d}E\,\mathrm{d}V = \iint \frac{q(\boldsymbol{r}, E)}{\mathrm{w}(\boldsymbol{r}, E)}\,\mathrm{d}E\,\mathrm{d}V\,, \tag{1}$$

and renormalizes the weight windows accordingly.

To greatly reduce the variance for certain problems, the Exnihilo framework supports the components for automating the creation of a deterministic problem from a given MC problem definition for hybrid MC–deterministic methods. The compositions from a user's model are converted to multigroup cross sections using SCALE's XSProc module. A ray tracer maps Shift geometries onto a structured Cartesian mesh used by Denovo. Source regions must be manually specified for Denovo, but automatic discretization is under development. Finally, after the approximate deterministic problem is solved, a solution adapter communicates the fluxes between processors and maps them from the Denovo domain decomposition to the Shift boundary mesh decomposition. Because no disk I/O is needed, hybrid calculations can be run on massively parallel architectures.

## 2.6 Parallel Decomposition

### 2.6.1 MSOD algorithm

In order to run Shift efficiently on laptops and leadership-class machines, Shift supports multiple parallel decompositions through the Multiple-Set Overlapping Domain (MSOD) algorithm [11]. This method divides the geometry into structured Cartesian "blocks" based on user-input $x$, $y$, and $z$ grids. An "overlap" fraction, also input by the user, specifies the fraction of each block that is shared with the adjacent neighboring domain. This overlap reduces the amount of communication needed during a transport solve because a particle is not communicated to an adjacent domain until it passes through the overlapping region. Without overlap, a particle could potentially scatter many times between adjacent domains, with a communication each time; therefore, overlapping domains are most effective in highly-scattering media. After decomposing the domain into blocks, the geometry is then replicated across "sets" of particle histories based upon the number of processors and the total number of particle histories the user has requested.

The MSOD setup is determined by a *boundary mesh* in Shift to divide each set into blocks, where the number of cells in the boundary mesh is equal to the number of blocks. The boundary mesh is superimposed on the problem geometry, thereby not requiring any special treatment of the physical geometry to implement MSOD. Furthermore, MSOD naturally reduces to full domain replication or full domain decomposition by defining one block per set. There are two main parallel communication channels needed for MSOD: *intra-set* (block-to-block communication within a set) and *intra-block* (set-to-set communication on the same block). A $k$-eigenvalue calculation in Shift requires three types of communication:

- intra-set communication of particles crossing domain boundaries during transport;
- intra-set and intra-block communication of fission sites between cycles;
- intra-set and intra-block communication of mesh-based tallies at the end of the problem.

When particle histories need to be communicated across blocks, Shift uses a modified domain decomposition algorithm [12, 13] to handle block-to-block communication.

Tallies and fission sites require both intra-set and intra-block communication. Intra-block communication is required to accumulate tally and fission site contributions in overlapping regions to their "parent" blocks depending on overlap and non-aligned meshes. Because $k$-eigenvalue estimators are global, contributions to $k$ must be performed across all sets and blocks using a full global reduction. MSOD scaling results are presented in §3.3. Full details of the MSOD parallel algorithms implemented in Shift will be given in an upcoming journal paper in the *Journal of Computational Physics*.

### 2.6.2 Fission sites

As mentioned previously, for Shift $k$-eigenvalue calculations, the fission sites from a cycle serve as the fission source for the next generation of particles. To account for statistical deviations in the estimate of $k$ for a cycle, particles are born with a modified weight that preserves the global weight of particles in each cycle. This sampling scheme can create load-imbalance across sets in parallel between subsequent cycles. Therefore, the fission bank in Shift must be rebalanced to alleviate this issue.

Shift uses an iterative version of the fission bank rebalance algorithm described in Ref. [14] to redistribute the fission sites across sets. The original algorithm has demonstrated excellent scalability for full domain replication (1 block per set), but it must be extended to handle two specific cases:

1. when sets are out of balance such that nearest neighbors do not have enough fission sites to communicate;
2. when there are multiple blocks per set (MSOD).

To alleviate the first condition, the original algorithm is wrapped in an iteration scheme until full fission bank balance is achieved across all sets. To handle multiple blocks, an additional constraint must be applied that limits set-to-set transfers based on similar block populations of fission sites.

Communication of fission sites is constrained between equivalent blocks. The number of fission sites to pull from each block and communicate is determined by trying to achieve equal numbers of particles on each block in a set for load balance. This condition is adequate for reactor problems but a different condition is needed for problems with large nonfissionable regions. To achieve this, the average number of fission sites per block is calculated in each set to preserve load balance. Then, the number of surplus sites on each block is calculated and sites are communicated across sets from this surplus by sampling uniformly. If not enough surplus sites are present, sites are taken from the regular fission sites. Full details outlining the fission site rebalance algorithm in Shift will be upcoming in a full paper in the *Journal of Computational Physics*.

### 2.7 Depletion Package

Coupling between Shift and ORIGEN [3], a depletion/transmutation analysis code, using ORIGEN's new C++ API is implemented in Exnihilo. Using this new in-memory API avoids the prohibitive cost (on HPC systems) of writing to disk between transport and depletion steps and also enables each depletion step to be performed on each depletable region in parallel. Shift uses the same approach as VESTA [15] to obtain microscopic per-nuclide reaction rates, which tallies ultra-fine-group fluxes in each depletion region and then uses these fluxes to collapse the microscopic CE cross sections into one-group reaction rates. Shift then sends these reaction rates to ORIGEN for a depletion calculation.

Currently, Shift uses constant-power depletion with either forward Euler or a "middlestep" method. The accuracy of the middlestep method has recently been investigated in [16, 17]. These studies show that the accuracy of this method can be adequate for certain problems. In order to optimize the parallel efficiency of the depletion calculation, the depletable regions on each block are distributed amongst the available processors in order to minimize the number of depletion solves performed on each core. After every core has calculated the new concentrations for its depletion regions, the results are broadcast to every other core on the block.

## 3 VALIDATION AND VERIFICATION

A selection of benchmarking results are presented in this section, focused specifically on LWR reactor analysis. Overall, we show that Shift performs very well for these problems, including comparison to experimental and other MC code results. These results also show that Shift can take advantage of leadership-class computing machines and perform efficiently on these machines.

### 3.1 BW1810 Criticality Benchmark

First, the results from the Babcock and Wilcox 1810 (BW1810) criticality experiments [18] are presented. In this paper, we compare eigenvalues, pin powers, and reactivity rod worths calculated using Shift with the published experimental data. The full specification for these small reactor cores is given in the benchmark report [18]. Cores I-XX, not including core XI, were modeled and simulated using Shift. These Shift models did not include any structures past the reactor vessel and did not model the center and top grid plates.

These models are based on critical experiments; therefore, the simulated $k_{\text{eff}}$ values are compared to a reference $k_{\text{eff}}$ of unity [18]. The average, minimum, and maximum reactivity differences between Shift and the measured data over all simulated core configurations are $-47$, $-117$, and 42 pcm, respectively. A difference of less than 100 pcm with the benchmark eigenvalue is considered good for these experiments. Of course, $k_{\text{eff}}$ alone is not an adequate measure for benchmarking these small reactor problems, so we also examine rod worths and pin powers. Figure 1 shows the relative difference in pin powers between Shift and the benchmark for core XIV. These pin powers are at the midplane of the center assembly as reported in the benchmark and cells shown in black are either water holes, control rods, or instrument tubes. All simulated cores have RMS differences at or below 2%.

To approximate the rod worths in these experiments, an approximate calculation of the value of boron parts per million (ppm) per pcm ($1 \times 10^{-5}$) of reactivity was performed using Shift simulations

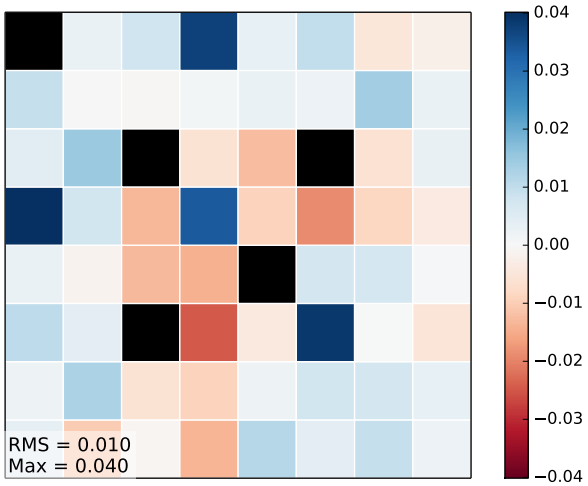


**Figure 1. BW1810 core-XIV relative difference in midplane pin powers for a quarter of the center assembly at the midplane between Shift and the benchmark.**

of three of the BW1810 cores. Figure 2 shows the percent error in $B_4C$ and Gadolinium (Gd) fuel rod worths approximated using Shift for compared to the benchmark values. These values show that Shift approximates the $B_4C$ rod worths to within 1.5% of the benchmark and the Gd fuel rod worth to within 6%. Full details of this benchmark study, including a comparison with OpenMC [19], will be included in a future publication.

### 3.2 Westinghouse AP1000® Simulations

Shift was tested using core physics simulations of the Westinghouse AP1000® PWR startup core. Details about this reactor and its simulation using the Consortium for Advanced Simulation of LWRs (CASL) Virtual Environment for Reactor Analysis (VERA) test stand can be found in [20,21]. The reference solutions for various AP1000® PWR configurations are from simulations using CE KENO-VI. Note that all of these Shift simulations were performed on Titan [22].

In Table I we compare the eigenvalues calculated by Shift and CE KENO-VI. As expected, Shift eigenvalues agree very well with those from CE KENO-VI, yielding an average difference of 18 pcm and maximum difference of 25 pcm over all AP1000® configurations simulated. Next, the power distribution of the AP1000® ARO (AO control rod bank configuration) problem was calculated using Shift and compared to KENO-VI and Denovo results. These results are shown in Table II. These results show that Shift accurately predicts the integrated pin power distributions across this reactor core to less than 1% error.

Rod bank worths were also calculated using Shift. Table III shows a comparison of the results with CE KENO-VI. The largest difference between Shift and CE KENO-VI occurs for a bank closest to the edge of the core with a difference of 2.3% rod worth with an overall RMS of 1.1% over all rod bank comparisons. A strong scaling study of Shift for problem 1 was performed on Titan using full domain replication, multiple sets with one block per set, with a fixed number of

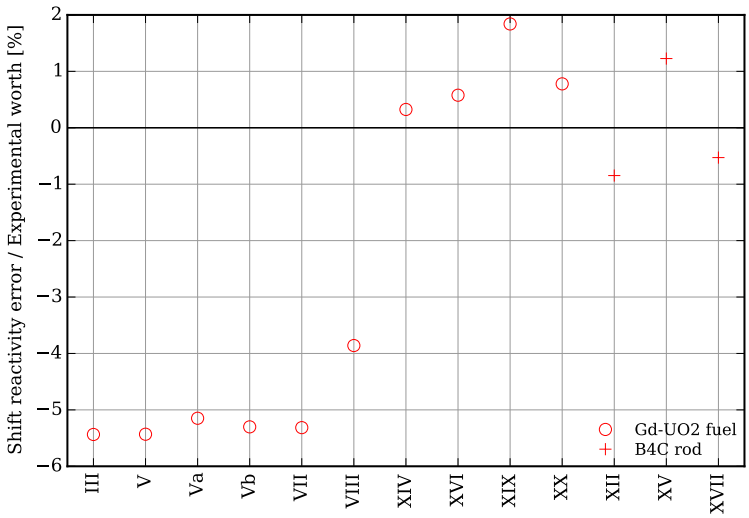


**Figure 2. Absolute error of Shift simulations of Gd-$UO_2$ fuel and $B_4C$ rod worths from the BW1810 experiment.**

**Table I. Shift eigenvalue and benchmark comparison results for assemblies in AP1000®**

| Case ID | KENO-VI $k_{eff}$ | Shift $k_{eff}$ | Diff [pcm] |
|---|---|---|---|
| **ARO** | 1.000870 | 1.001030 | 16 |
| **DBW** | 1.003240 | 1.003450 | 21 |
| **MA** | 0.998258 | 0.998414 | 16 |
| **MB** | 0.998669 | 0.998909 | 24 |
| **MC** | 0.998956 | 0.999148 | 19 |
| **MD** | 0.998496 | 0.998643 | 15 |
| **M1** | 0.994350 | 0.994548 | 20 |
| **M2** | 0.992001 | 0.992185 | 18 |
| **AO** | 0.984609 | 0.984749 | 14 |
| **S1** | 0.990103 | 0.990200 | 10 |
| **S2** | 0.989935 | 0.990183 | 25 |
| **S3** | 0.989650 | 0.989739 | 9 |
| **S4** | 0.995055 | 0.995295 | 24 |
| **Average** | | | **18** |

**Table II. Comparisons of eigenvalues and overall pin powers, ΔP, for AP1000®**

| Code | $k_{eff}$ ± 2pcm | $\Delta k_{eff}$ [pcm] | RMS Asm. [%] | Max Asm. [%] | RMS Pin [%] | Max Pin [%] | Hot Pin [%] |
|---|---|---|---|---|---|---|---|
| **KENO-VI** | 1.00096 | - | - | - | - | - | - |
| **Denovo** | 1.00086 | −10 | 0.2 | 0.4 | 0.4 | 3.0 | 1.0 |
| **Shift** | 1.00131 | 35 | 0.2 | 0.5 | 0.2 | 0.6 | −0.2 |

**Table III. Comparison of rod bank worth calculated using Shift to CE KENO-VI for AP1000®**

| Rod Bank | Material | KENO-VI Worth [pcm] | Shift ΔWorth [pcm] | [%] |
|---|---|---|---|---|
| **MA** | Tungsten | 258 | 4 | 1.4 |
| **MB** | Tungsten | 217 | −5 | −2.3 |
| **MC** | Tungsten | 188 | 0 | 0.1 |
| **MD** | Tungsten | 234 | 5 | 2.0 |
| **M1** | Ag-In-Cd | 651 | −1 | −0.2 |
| **M2** | Ag-In-Cd | 887 | 1 | 0.1 |
| **AO** | Ag-In-Cd | 1635 | 5 | 0.3 |
| **S1** | Ag-In-Cd | 1079 | 9 | 0.8 |
| **S2** | Ag-In-Cd | 1096 | −7 | −0.6 |
| **S3** | Ag-In-Cd | 1124 | 10 | 0.9 |
| **S4** | Ag-In-Cd | 580 | 6 | 0.1 |
| | | **RMS** | **6** | **1.1** |
| | | **Max** | **10** | **2.3** |

particle histories per cycle of $5 \times 10^8$. The excellent scaling Shift attains is shown in Fig. 3 with a parallel efficiency of 97% to 100% on Titan for this problem.

### 3.3 MSOD Scaling Study

We performed a small scaling study of the MSOD and fission site rebalance algorithms discussed in §2.6. This study used a reactor eigenvalue problem consisting of a water-moderated $4 \times 4$ assembly configuration with each assembly made of a $17 \times 17$ array of $UO_2$ fuel pins. This problem ran with SCE physics and RTK geometry with a fixed decomposition of a $2 \times 2$ array of blocks per set. A mesh tally was used to calculate the power in each fuel pin. The strong and weak scaling studies were performed on the OLCF machine, Titan [22].

Strong scaling was demonstrated by keeping the number of particle histories per cycle, $5 \times 10^6$, constant as the number of cores was increased. Results are shown in Fig. 4. This figure shows the speedup achieved per set with various overlap fractions. All of these overlap fractions yield very similar good scaling results on Titan for this problem with an efficiency ranging from 100% down to 84%.

To demonstrate weak scaling, the number of particle histories per set was kept constant at $6 \times 10^4$; therefore, the total number of particle histories in the full problem is increased with increasing number of sets. These results are shown in Fig. 5. This figure shows the Shift parallel efficiency with multiple sets. Note that the efficiencies for various overlap fractions are all within 3% of the ideal scaling curve, and therefore show excellent weak scaling of the MSOD algorithm on Titan for this problem. More rigorous testing of this algorithm as applied to neutron-photon fixed source and eigenvalue problems is underway.

We note that the reactor physics problems solved with Shift to date have all fit in memory with

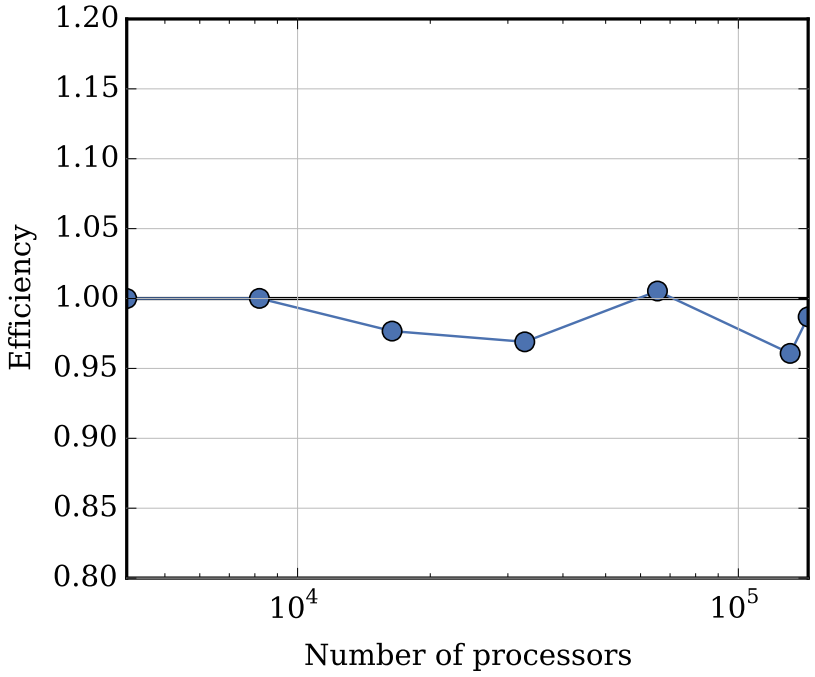


**Figure 3. Shift strong scaling results for AP1000® problem 1 on Titan.**

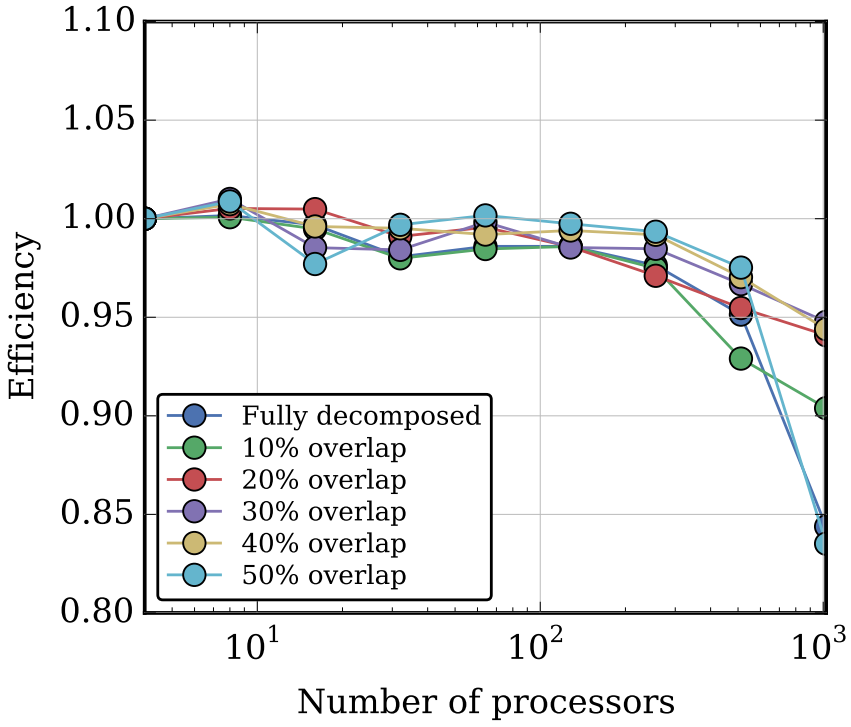


**Figure 4. Strong scaling plot of the $4\times4$ reactor assembly problem on Titan showing efficiency versus the number of sets. Each set is decomposed into $2 \times 2$ blocks.**

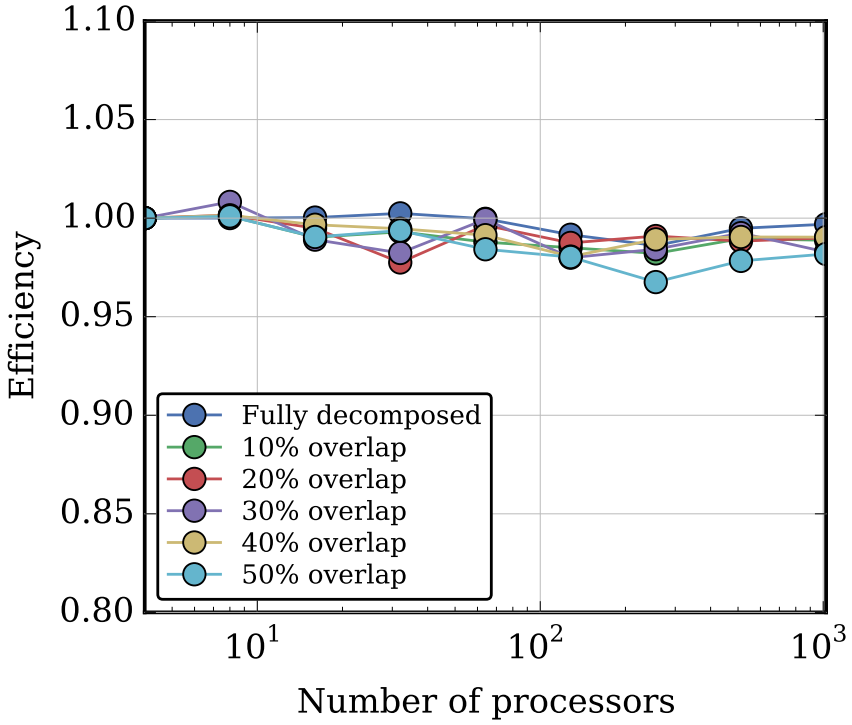


**Figure 5. Weak scaling of the $4 \times 4$ reactor assembly problem on Titan with each set decomposed into $2 \times 2$ blocks. This plot shows efficiency versus number of sets. The multi-set solve time is normalized by the single-set solve time to produce these efficiency results.**

a fully replicated MSOD decomposition. However, the full MSOD decomposition method will be essential for reactor analysis problems planned in the near future, namely middle- and end-of-cycle LWR depletion analyses. Unlike beginning-of-cycle power simulations, these depletion analyses require nuclear data for hundreds of additional nuclides to be stored in memory alongside detailed space- and energy-dependent flux tallies. The MSOD method will allow these detailed flux tallies to be decomposed across processors on large clusters.

## 4 CONCLUSIONS

Overall, Shift was designed to efficiently be able to tackle a wide class of applications. To suit this purpose, the transport components in Shift are designed to be geometry and physics agnostic. This allows the incorporation of multiple geometry and physics implementations with minimal integration effort.

Over the last few years, the development of Shift has focused on solving LWR core physics problems. The benchmarking and scaling results discussed in this paper show some of the successful validation and verification of Shift for LWR core-analysis applications. The implementation of the MSOD parallel algorithm in Shift allows for scaling and a smooth transition between fully domain-replicated topologies and spatially domain-decomposed problems. Results using the MSOD algorithm demonstrate excellent strong and weak scaling for multi-set, single-set and multi-block decompositions.

## 5 ACKNOWLEDGMENTS

Work for this paper was supported by Oak Ridge National Laboratory (ORNL), which is managed and operated by UT-Battelle, LLC, for the U.S. Department of Energy (DOE) under Contract No. DEAC05-00OR22725. This research was also supported by the Consortium for Advanced Simulation of Light Water Reactors (www.casl.gov), an Energy Innovation Hub (http://www.energy.gov/hubs) for Modeling and Simulation of Nuclear Reactors and sponsored by the Laboratory Directed Research and Development Program of ORNL, managed by UT-Battelle, LLC, for the U.S. DOE. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

## 6 REFERENCES

[1] T. M. Evans, A. S. Stafford, R. Slaybaugh, and K. T. Clarno, "Denovo: A New Three-Dimensional Parallel Discrete Ordinates Code in SCALE," *Nuclear Technology*, **171**, pp. 171–200 (2010).

[2] J. J. Jarrell, T. M. Evans, G. G. Davidson, and A. T. Godfrey, "Full Core Reactor Analysis : Running Denovo on Jaguar," *Nuclear Science and Engineering*, **175**, pp. 283–291 (2013).

[3] "SCALE: A Comprehensive Modeling and Simulation Suite for Nuclear Safety Analysis and Design," ORNL/TM-2005/39, Version 6.1, Oak Ridge National Laboratory (2011).

[4] X-5 Monte Carlo Team, "MCNP - A General Monte Carlo N-Particle Transport Code, Version 5," LA-UR-03-1987, Los Alamos National Laboratory (2008).

[5] S. W. Mosher et al., "ADVANTG—An Automated Variance Reduction Parameter Generator," ORNL/TM-2013/416, Oak Ridge National Laboratory (2013).

[6] M. B. Emmett, "MORSE-CGA: A General Purpose Monte Carlo Multigroup Neutron and Gamma-Ray Transport Code System with Array Geometry Capability, Version 2," ORNL-6174, Oak Ridge National Laboratory (1985).

[7] Cross Sections Evaluation Working Group, *ENDF-6 Formats Manual: Data Formats and Procedures for the Evaluated Nuclear Data Files ENDF/B-VI and ENDF/B-VII*, Brookhaven National Laboratory, 2 edition, 2011.

[8] M. E. Dunn and N. M. Greene, "AMPX-2000: A Cross-Section Processing System for Generating Nuclear Data for Criticality Safety Applications," *Trans. Am. Nucl. Soc.*, **86**, pp. 118 (2002).

[9] T. Ueki and F. B. Brown, "Stationarity Modeling and Informatics-Based Diagnostics in Monte Carlo Criticality Calculations," *Nucl. Sci. Eng.*, **149**, *1*, pp. 38–50 (2005).

[10] J. Wagner and A. Haghighat, "Automated Variance Reduction of Monte Carlo Shielding Calculations Using the Discrete Ordinates Adjoint Function," *Nuclear Science and Engineering*, **128**, *2*, pp. 186–208 (1998).

[11] J. C. Wagner, S. W. Mosher, T. M. Evans, D. E. Peplow, and J. A. Turner, "Hybrid and parallel domain-decomposition methods development to enable Monte Carlo for reactor analyses," *Progress in Nuclear Science Technology*, **2**, pp. 815–820 (2011).

[12] T. A. Brunner, T. J. Urbatsch, T. M. Evans, and N. A. Gentile, "Comparison of four parallel algorithms for domain decomposed implicit Monte Carlo," *Journal of Computational Physics*, **212**, pp. 527–539 (2006).

[13] T. A. Brunner and P. S. Brantley, "An efficient, robust, domain-decomposition algorithm for particle Monte Carlo," *Journal of Computational Physics*, **228**, pp. 3882–3890 (2009).

[14] P. Romano and B. Forget, "Parallel Fission Bank Algorithms in Monte Carlo Criticality Calculations," *Nuclear Science and Engineering*, **170**, pp. 125–135 (2012).

[15] "VESTA User's Manual - Version 2.1.0," DSU/SEC/T/2008-331, IRSN (2008).

[16] A. E. Isotalo and V. Sahlberg, "Comparison of Neutronics-Depletion Schemes for Burnup Calculations," *Nuclear Science and Engineering (to appear)*.

[17] A. E. Isotalo, "Comparison of Neutronics-Depletion Coupling Schemes for Burnup Calculations 2," *Nuclear Science and Engineering (to appear)*.

[18] L. W. Newman et al., "Urania-Gadolinia: Nuclear Model Development and Critical Experiment Benchmark," DOE/ET/34212-41 BAW-1810, Babcock and Wilcox Co. (1984).

[19] P. K. Romano and B. Forget, "The OpenMC Monte Carlo Particle Transport Code," *Annals of Nuclear Energy*, **51**, pp. 274 (2013).

[20] F. Franceschini, A. T. Godfrey, J. Kulesza, and R. Oelrich, "Westinghouse VERA Test Stand: Zero Power Physics Test Simulations for the AP1000 PWR," CASL-U-2014-0012-000, CASL (2014).

[21] F. Franceschini, B. Oelrich, and J. C. Gehin, "Simulation of AP1000 first core with VERA," *NEI Magazine* (2014).

[22] Oak Ridge Leadership Computing Facility, "Titan Cray XK7," `https://www.olcf.ornl.gov/computing-resources/titan-cray-xk7/`, December, 2014.