ORNL/TM-2021/2280

Methods and Usability Enhancements in Shift for Non-LWR Applications

Tara Pandya
Matthew Jessee
Tarek Ghaddar
Friederike Bostelmann

November 7, 2021

OAK RIDGE National Laboratory

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

DOCUMENT AVAILABILITY

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

Website osti.gov

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

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

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

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

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.

ORNL/TM-2021/2280

Nuclear Energy and Fuel Cycle Division

METHODS AND USABILITY ENHANCEMENTS IN SHIFT FOR NON-LWR APPLICATIONS

Tara Pandya
Matthew Jessee
Tarek Ghaddar
Friederike Bostelmann

Date Published: November 7, 2021

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

LIST OF TABLES

ACRONYMS

API	application programming interface
CRAB	Comprehensive Reactor Analysis Bundle
CSAS	Criticality Safety Analysis Sequence
GG	General Geometry
HDF5	Hierarchical Data Format 5
HTGR	high-temperature gas-cooled reactor
INL	Idaho National Laboratory
LWR	light-water reactor
M&S	modeling and simulation
MC	Monte Carlo
NEAMS	Nuclear Energy Advanced Modeling and Simulation
NRC	Nuclear Regulatory Commission
ORANGE	Oak Ridge Advanced Nuclear Geometry Engine
ORNL	Oak Ridge National Laboratory
TRISO	tristructural isotropic

ABSTRACT

Several development and analysis tasks were undertaken in FY21 under the Nuclear Energy Advanced Modeling and Simulation program to enhance modeling of non light–water reactors (LWRs) with Shift. Specifically, these efforts targeted enhancements for tristructural isotropic (TRISO) fuel modeling. A new Shift user interface was developed that allows for much better usability and ease of modeling for non-LWR problems and TRISO fuel. Performance studies were conducted using an HTR-10 fuel pebble model by comparing different geometry packages in Shift, KENO-VI, and Serpent. These studies showed that the new geometry package in Shift performs well compared to Serpent for TRISO fuel modeling with consistent tracking options between both packages. The studies also identified the most critical areas of improvement for more efficiently performing Monte Carlo transport on TRISO fuel models with Shift. Tally calculations in Shift were optimized for non-LWR cross section generation and depletion calculations, and areas for further optimization and accuracy improvements were identified. Finally, initial collaboration efforts were formed between Idaho National Laboratory, Argonne National Laboratory, and the Nuclear Regulatory Commission to use Shift for Comprehensive Reactor Analysis Bundle support.

1. INTRODUCTION

The US Department of Energy Office of Nuclear Energy's Nuclear Energy Advanced Modeling and Simulation (NEAMS) program develops modeling and simulation (M&S) tools to accelerate the deployment of advanced nuclear energy technologies. The reactor physics technical area focuses on the development of M&S tools for modeling reactor physics phenomena (i.e., neutral particle transport and isotopic depletion/decay). The three primary codes supported in the reactor physics technical area are (1) Griffin for non light-water reactor (LWR) physics, (2) MPACT for LWR physics, and (3) the Shift Monte Carlo (MC) code used for reference solutions, ex-core dose assessment, cross section generation, and sensitivity and uncertainty assessment.

The purpose of this report is to present the Shift modifications and other work performed to support tristructural isotropic (TRISO) fuel modeling, code usability, performance enhancements, and Comprehensive Reactor Analysis Bundle (CRAB) support. Section 2 presents the Shift tally enhancements, both implemented and planned. Section 3 presents the layout of the user interface enhancements for Shift. Section 4 presents analysis of past and current TRISO modeling performance efforts. Section 5 presents the collaboration efforts with the Nuclear Regulatory Commission (NRC) for CRAB integration. Finally, Section 6 presents a summary of this milestone work.

2. SHIFT TALLY ENHANCEMENTS

As part of the Shift tally enhancement work performed under this milestone, the team drafted a journal article that describes work built upon the Shift L2 milestone completed earlier in FY21 [1]. The draft of this journal article will be submitted to the *Annals of Nuclear Energy* after it completes the Oak Ridge National Laboratory (ORNL) internal review process [2].

The previous L2 milestone report listed both needed and planned tally improvements for the two-step neutronics workflow for Shift and Griffin to improve ease of use and accuracy. This list is reproduced below along with the current status and estimated completion (in FY22) of each tally improvement. An additional tally improvement was added for the calculation of mu-bar and the transfer matrix.

1. The scatter matrix cross section is computed from the scattering event path–length tally, which includes multiplicity of nonfission production reactions. The secondary energy distribution is computed from the scatter matrix probability collision tally, which does not include multiplicity. This approach conserves neutron balance, but the secondary energy distribution would be more accurate if multiplicity were included by using the postcollision particle weight instead of the precollision weight. Estimated Completion: In Progress, 11/2021.

2. The scattering matrix collision tally can be extended to include higher-order moments for use with Griffin transport solvers. Estimated Completion: In Progress, 11/2021.

3. The fission spectrum is currently computed using the birth tally method. Although this method provides accurate macroscopic values that include prompt and delayed contributions, nuclide-dependent and prompt/delayed χ distributions can be computed by including additional nuclide-dependent ν-fission path-length tallies per cross section set. These nuclide-dependent tallies can also be utilized to compute forward-weighted homogenized kinetic parameters for each cross section region. Estimated Completion: 03/2022.

4. For the previous L2 milestone, the κ-fission cross sections were computed outside of the Shift calculation by scaling the fission cross section tally by a prescribed κ value. In future work, Shift can be modified to compute energy-released and energy-deposited cross sections based on the library, κ, and heating values. Estimated Completion: 03/2022.

5. The Empire calculations from the L2 milestone were performed using cell union tallies. This approach requires the user to provide unique cell and universe identifiers for each pin cell to compute the required homogenized cross section and reference flux values. This approach will be significantly improved by enhancing the Shift hexagonal mesh tally capability. Estimated Completion: In Progress, 02/2022.

6. Treatment for diffusion-coefficient calculation for thermal systems dominated by hydrogen scattering can be improved. Although recent R&D has focused on the cumulative migration method, for which an initial implementation exists in Shift, an easier approach would be to implement hydrogen-based transport correction ratios, as done in VERA, SCALE, and Serpent (with the *set trc* card). Estimated Completion: 03/2022.

7. A single tally class within the Shift tally system should be established to compute the requisite reaction rates for Griffin and convert them to the final format before archiving them in Hierarchical Data Format 5, commonly known as HDF5. Estimated Completion: In Progress, 11/2021.

8. Additional investigation of the Empire calculations revealed large diffusion coefficient and scattering matrix tally uncertainties for the various cross section sets. The large uncertainties are owing to the small number of collision events in both the spatial domain of the tallies and the energy bins of the cross section structure. A hybrid approach revealed improved diffusion coefficient and scattering matrix tallies. This hybrid approach used a combination of material-based collision tallies for mean scattering angle and transfer matrix energy distribution along with local path–length tallies for the total and scattering cross sections. This approach should be automated as part of the improvements to the diffusion coefficient calculation. Estimated Completion: 11/2021.

3. SHIFT USER INTERFACE ENHANCEMENTS

To more succinctly and easily define reactor models, an enhanced user interface for Shift is under development. The new interface enhancements are listed below.

1. A two-layer input design includes the following features:
   - a user-facing layer constructed with a standardized input format (e.g., JSON, XML, YAML, SON from NEAMS Workbench),
   - an in-code layer built on a lightweight serializable data structure for serialization and validation, and
   - an input design that supports low-level and high-level specification of materials, geometries, and state definitions.
2. Geometry enhancements listed below are also included:
   - automated volume calculation,
   - automated unique materials for repeated model components,
   - automated cell and surface tagging for queries, and
   - hierarchical geometry definition built from engineering-style universes (assembly/lattice/pin) and/or primitive shape universes (polygons/circles/spheres/etc).

A significant portion of the reactor geometry infrastructure was developed in FY21, and the focus of FY22 will be on TRISO-based microreactor support and workflow automation based on multiple state definitions.

Figure 1 shows the different SCALE packages utilized by this interface. The internal Shift application programming interface (API) allows external code packages to run Shift with Oak Ridge Advanced Nuclear Geometry Engine (ORANGE), which is an improved constructive solid geometry package within SCALE. This improved geometry package reduces the memory footprint of the geometry object stored as well as the tracking performance for TRISO fuel problems.

Figures 2 and 3 show the Empire benchmark and an HTR-10 fuel pebble modeled with an initial prototype of the new user interface. Figure 4 shows the portion of the input for the HTR-10 fuel pebble used for modeling the geometry. Table 1 shows the different input sizes for describing a reactor assembly in the Empire microreactor assembly, the Empire microreactor core, and an HTR-10 fuel pebble.

Table 1. The number of input lines for three inputs using the improved interface and SCALE/Shift

Input	Improved Interface Input Lines	SCALE/Shift Input Lines
Empire Assembly	299	12,421
Empire Core	517	223,290
HTR-10 Fuel Pebble	117	412,601

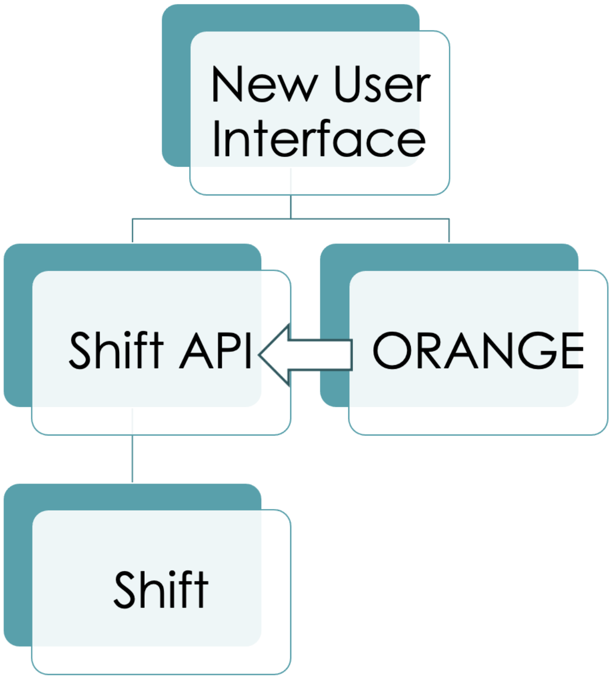

Figure 1. The package structure for `Shift`'s new user interface.

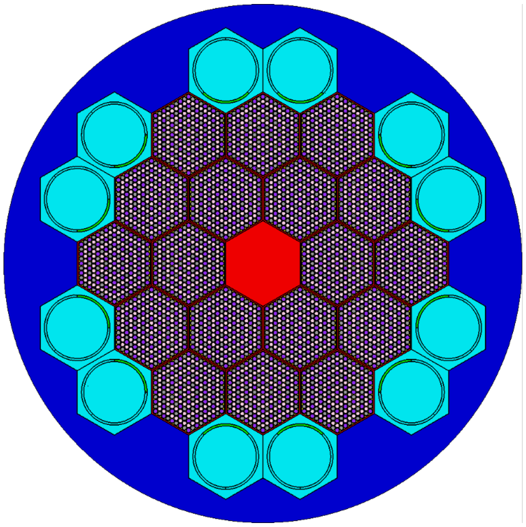

Figure 2. The Empire benchmark modeled with ORANGE using `Shift`'s new user interface.

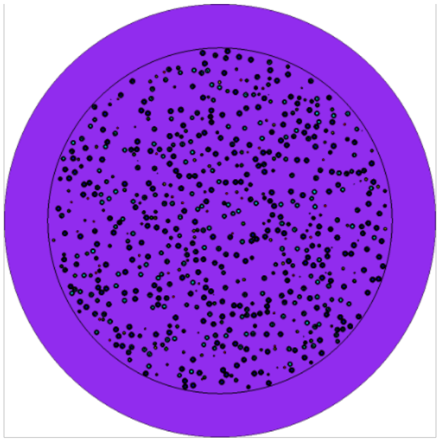

Figure 3. A 2D slice of the HTR-10 fuel pebble modeled with ORANGE using `Shift`'s new user interface.

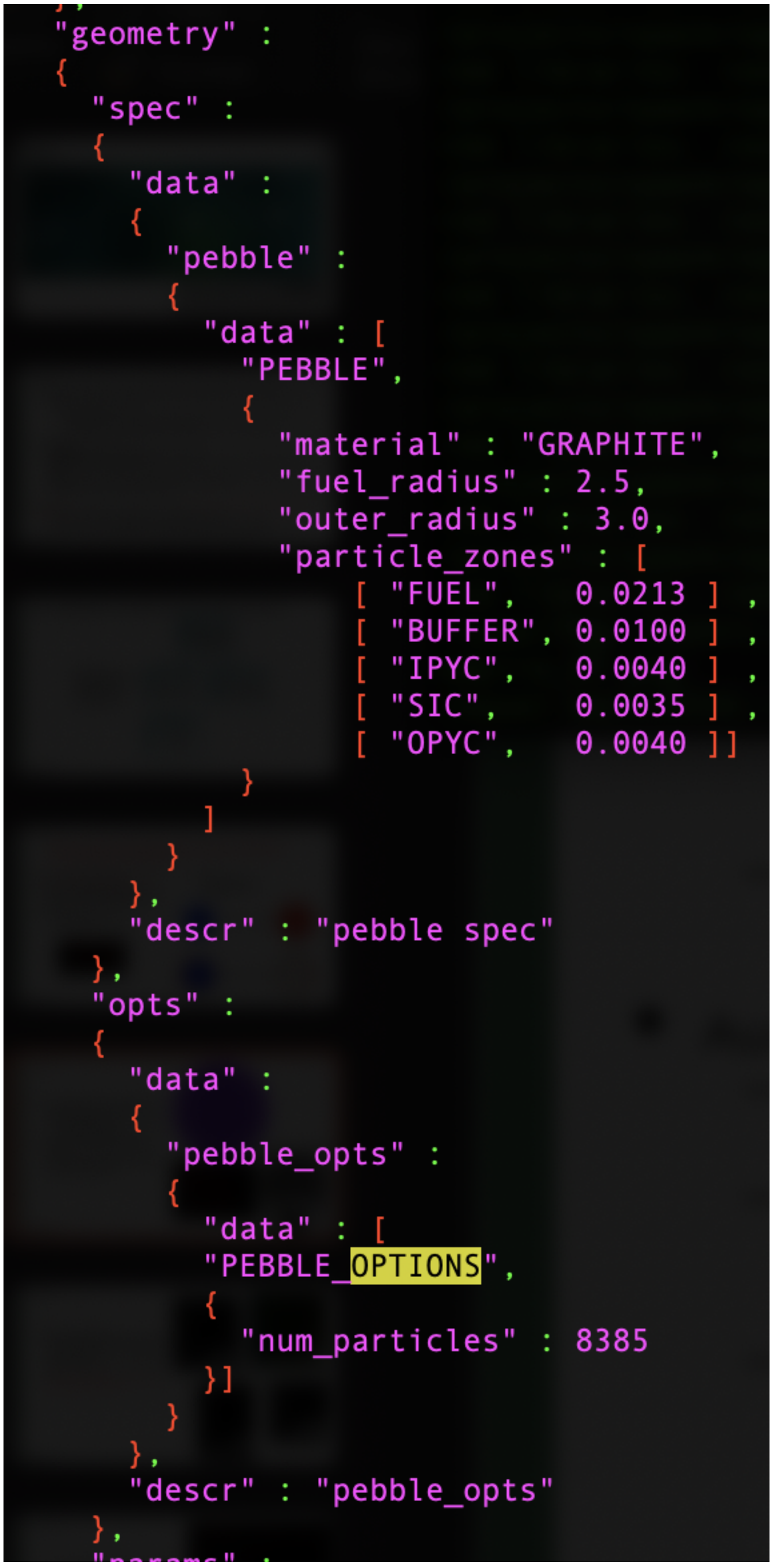

```
 },
 "geometry" :
 {
   "spec" :
   {
     "data" :
     {
       "pebble" :
       {
         "data" : [
           "PEBBLE",
           {
             "material" : "GRAPHITE",
             "fuel_radius" : 2.5,
             "outer_radius" : 3.0,
             "particle_zones" : [
                 [ "FUEL",    0.0213 ] ,
                 [ "BUFFER", 0.0100 ] ,
                 [ "IPYC",    0.0040 ] ,
                 [ "SIC",     0.0035 ] ,
                 [ "OPYC",    0.0040 ]]
           }
         ]
       }
     },
     "descr" : "pebble spec"
   },
   "opts" :
   {
     "data" :
     {
       "pebble_opts" :
       {
         "data" : [
         "PEBBLE_OPTIONS",
         {
           "num_particles" : 8385
         }]
       }
     },
     "descr" : "pebble_opts"
   },
   "params" :
```

Figure 4. The geometry input for the HTR-10 fuel pebble with the enhanced Shift user interface.

4. TRISO PERFORMANCE

Previously, Shift was used to model high-temperature gas-cooled reactor (HTGR) fuel systems, as documented by Bostelmann et al. [3]. In this previous work, Shift was applied to calculate the initial criticality of the HTR-10 benchmark and was also used in performance tests for a single fuel–pebble model. Consistent results were obtained between Shift, SCALE's KENO-VI [4] MC code, and the Serpent [5] MC code.

The study herein provides eigenvalue comparisons between Shift, SCALE/KENO-VI, and Serpent for a single fuel–pebble model based on the HTR-10 benchmark [6], as already used by Bostelmann et al. [3]. Reflective boundary conditions were specified on the pebble model's bounding cuboid to model an infinite, simple-cubic lattice of fuel pebbles surrounded by He. Such an HTGR model provides challenges for the MC transport calculation owing to the large number of surfaces and cells required to model the large number of TRISO fuel particles distributed in the fuel region.

4.1 HTR-10 FUEL PEBBLE DESCRIPTION

The HTR-10 fuel pebble consists of a fuel zone that is 5 cm in diameter and contains 8,385 TRISO fuel particles distributed randomly in a graphite matrix; this matrix is surrounded by a 5 mm graphite layer, which results in a pebble with a 6 cm outer diameter (Figure 5). The TRISO fuel particle is 0.91 mm in diameter and includes a microfuel kernel composed of UO_2 that is 0.5 mm in diameter. The fuel kernel is enclosed by four concentric coatings: a porous graphite buffer, an inner pyrolytic carbon layer, a ceramic silicon carbide layer, and an outer pyrolytic carbon layer. The ^{235}U enrichment of the fuel is 17 wt % in this HTR-10 configuration.

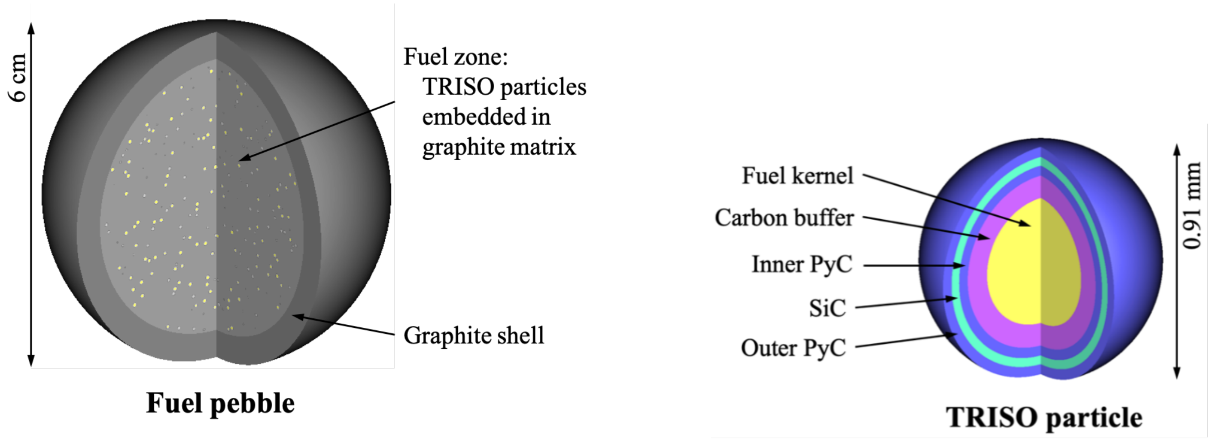

Figure 5. HTR-10 fuel pebble and TRISO model.

4.2 FUEL PEBBLE MODELS

A single-heterogeneous model with a smeared-fuel region and various double-heterogeneous models with explicit consideration of the fuel particles in the fuel region were developed to assess the impact on computation time, memory, and reactivity prediction.

1. **Volume-homogenized:** The fuel kernel, all layers of the TRISO particles, and the graphite matrix material were volume homogenized and placed into the fuel region.

2. **Square particle lattice:** The TRISO particles were placed in a square lattice. The locations of the particles in the lattice were chosen so that no particle intersected with the outer boundary of the fuel region. The lattice pitch was chosen so that the correct number of particles was included in the model.
3. **Random particle distribution:** The TRISO particles were placed at random locations in the fuel region. The locations were chosen so that no particle intersected with any other particle or the outer boundary of the fuel region. The random particle distributions were modeled in five different ways:
   (a) using Shift's internal *randommix* capability, which automatically determines random locations of the particles based on a simple user input within the main particle and pebble details,

   (b) explicitly placing the particles in the geometry using individual definitions of surfaces for each particle within the fuel region,

   (c) using the same explicitly placed particles but dividing the fuel region with an overlaid mesh to reduce surface search time,

   (d) defining a TRISO particle as its own unit and placing the particles as holes into the fuel pebble unit at the same coordinates as the previous models, and

   (e) using the above approach with TRISO particles as holes but modeling the fuel region model with an overlaid mesh.

Examples of each of these modeling approaches are shown in Figure 6.

Shift was run three different ways to understand the optimal execution of the code:

1. Shift ran through the Omnibus general front end using an input format that closely resembles the implementation of Shift's General Geometry (GG) engine.
2. Shift ran through the SCALE Criticality Safety Analysis Sequence (CSAS) using the input format of SCALE/KENO-VI's geometry. This is the traditional way that many users would interact with Shift. The geometry is processed through Hyas for the GG engine. Using slightly different assumptions when modeling the geometry can lead to different performance characteristics when compared with the first approach.
3. Lastly, a new geometry processing engine, ORANGE, is currently being implemented for Shift. ORANGE can be called through CSAS, as described above.

The key difference between GG and ORANGE is that GG tries to make a fully consistent and essentially unfolded global definition of the geometry, in which unfolded means that every geometry unit embedded in another unit or array has a separate instance within each higher-level universe. Therefore, particles are transported in a single universe level at a time. In contrast, ORANGE uses a layered approach in which particles are transported through its universe hierarchy simultaneously because higher levels automatically mask lower universe levels.

All calculations were performed using the same version of code, except for the last approach, which used a branch of the current SCALE source code to enable the geometry processing through ORANGE instead of Hyas/GG. For Serpent, calculations were performed with delta tracking and with surface tracking [5].

All MC calculations were run using $10,000$ neutrons per cycle with 50 active cycles and 25 inactive cycles. The ENDF/B-VII.1 nuclear data library was used for all calculations. The calculations were run at ORNL on a Linux cluster of AMD Opteron 6378 CPUs running at 2.4 GHz. Each node consists of 32 CPUs with 128 GB of memory available per node. The full node was reserved when running these calculations, but only 1 out of 32 CPUs was utilized to avoid the impact of parallel implementations.

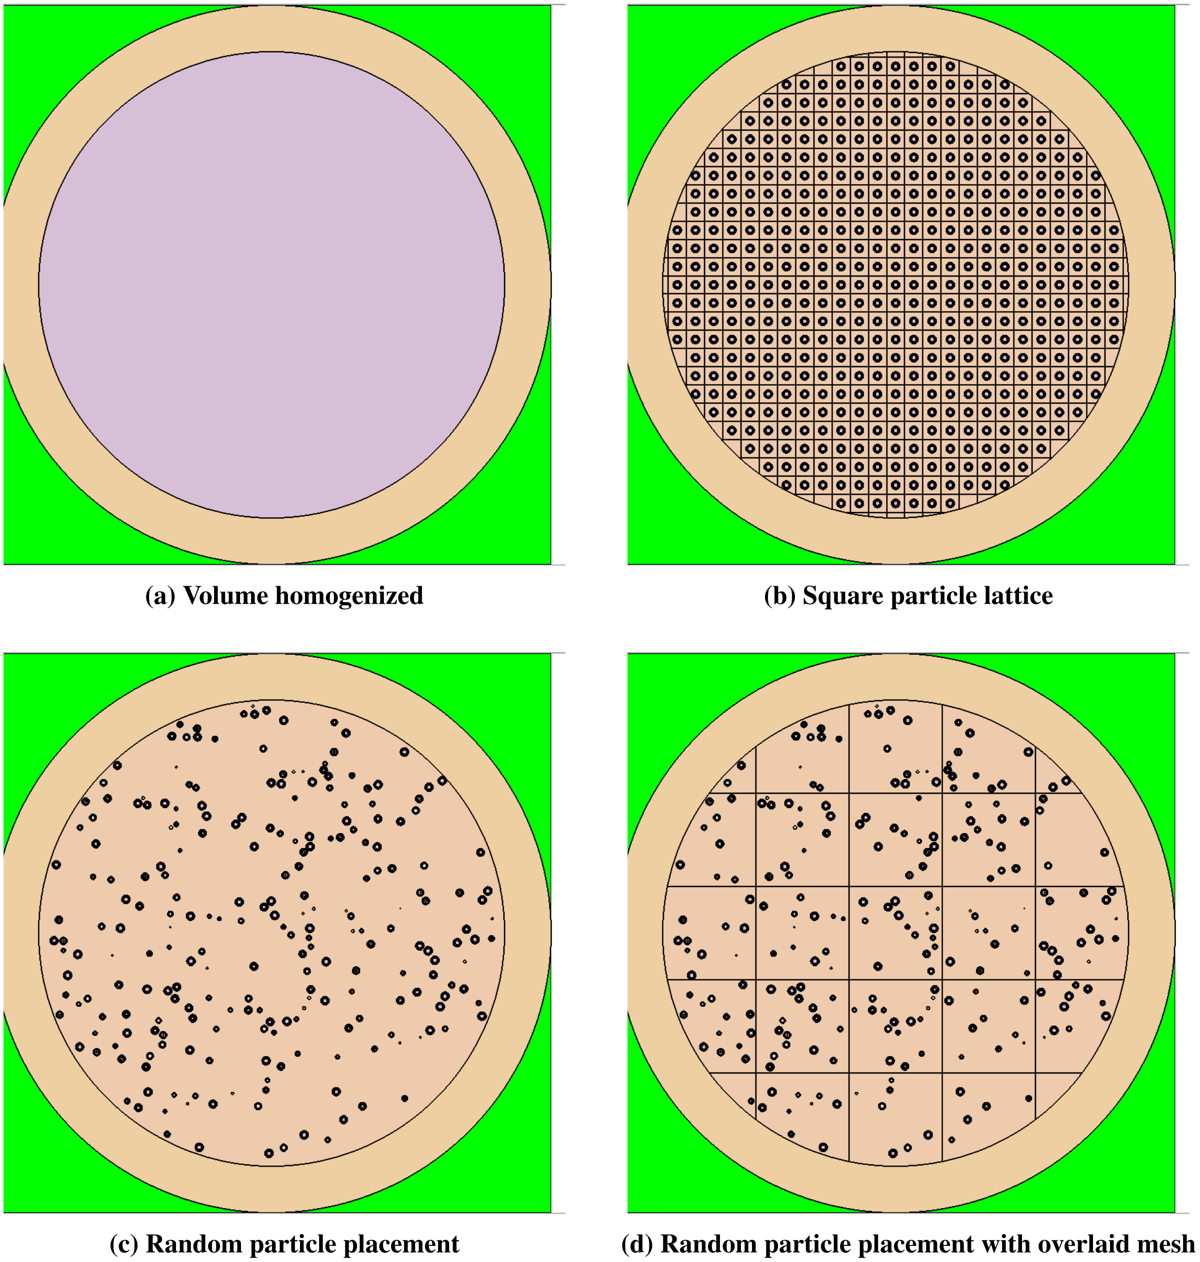

(a) Volume homogenized

(b) Square particle lattice

(c) Random particle placement

(d) Random particle placement with overlaid mesh

Figure 6. Various methods used to model TRISO particles.

4.3 PERFORMANCE ASSESSMENT

Table 2 details the computation time, memory footprint, and eigenvalue k_{inf} for each of the cases discussed above. The first column in the table indicates the MC transport code, how this code is called (e.g., through SCALE or through Omnibus), and which geometry engine was used.

The performance metrics show that the single-heterogeneous, volume-homogenized models are faster and use less memory than the double-heterogeneous models. All codes show similar computation times and memory footprints for the calculation of the volume-homogenized model. The additional surfaces in the double-heterogeneous models cause significant increases in computation time and memory use depending on the choice of code and the modeling. The eigenvalue for the volume-homogenized model is consistently different from any of the double-heterogeneous models. This is a well known impact of TRISO fuel owing to the spatial shielding effects of the fuel that cannot be captured by volume homogenizing directly. This case is only used to provide a baseline computation time for traditional single-heterogeneous fuel components.

For the double-heterogeneous models, the Shift calculations that use the SCALE front end are much slower than those that use the native Omnibus front end. This is caused by additional geometry calculations (i.e., surface intersection checks) that are activated through the SCALE geometry model. Furthermore, when running through Omnibus, Shift uses slightly less memory than SCALE/Shift. For example, SCALE/Shift takes 2× longer for the lattice model and 3.8× longer for the random model using holes.

When comparing SCALE/Shift calculations with the different geometry engines, different trends were observed for the computation time. For models developed using holes, the calculations running through ORANGE resulted in a significant run time decrease. The most significant decrease (a factor of 3.5) was observed for the model that inserted all TRISO particles as holes in the fuel region. In contrast, the computation time was greatly increased (by a factor of 1.6) when using the random model with TRISO particles modeled explicitly in the fuel region using individual surfaces. For the lattice model, ORANGE improved both the computation time and the memory footprint: the run time was decreased by almost half compared to Hyas/GG, and the memory footprint was reduced by almost a quarter compared to Hyas/GG.

When using either geometry engine, the *randommix* block resulted in a run time improvement compared to the explicit TRISO particle models (i.e., no mesh) because the applied geometry conditions in this block allowed for fewer geometry checks compared to an explicit model.

Comparing SCALE/Shift running through ORANGE with SCALE/KENO-VI, similar run times were obtained for the homogenized case and the lattice case. However, for the random mesh calculation, KENO-VI required a substantially longer run time compared to Shift. The KENO-VI calculations for the other random particle models took significantly more time—owing to the significant processing time required for the geometry—and were not completed in time for this report.

Looking at the performance of Shift using ORANGE vs. Serpent using the traditional surface tracking, which is also used in Shift, similar performance was obtained for the volume-homogenized models and the lattice models. However, Serpent clearly outperformed Shift when using models with random particle distributions. The Shift models with one fuel region (i.e., no subdivision through a mesh) resulted in computation times 14×, 26×, and 112× longer than the Serpent random calculation. However, the subdivision of the fuel region into multiple zones using an overlaid mesh caused a significant run time decrease because the distance to the boundary check is only performed for TRISO particles within each mesh cell, as opposed to checking the entire set of TRISO particles within the fuel region. Using holes to place TRISO particles into such a coarse mesh, Shift using ORANGE resulted in slightly shorter run times and a slightly smaller memory footprint compared with Serpent.

Table 2. A comparison of performance metrics between Shift, SCALE/KENO-VI, and Serpent

Code	Model	Time (h)	Memory (GB)	k_{inf}	σ
Omnibus/Shift	homogenized	0.23	0.17	1.57051	0.00194
GG	lattice	**1.11**	1.03	1.67295	0.00259
	random (mesh overlay, holes)	1.53	0.22	1.67884	0.00222
	random (explicit, holes in global unit)	**45.83**	0.20	1.67585	0.00196
	random (explicit, surfaces in global unit)	92.48	0.67	1.67588	0.00246
SCALE/Shift	homogenized	0.36	0.30	1.56759	0.00270
GG through Hyas	lattice	**2.27**	**1.18**	1.67640	0.00280
	random (mesh overlay, holes)	2.11	0.38	1.67982	0.00270
	random (mesh overlay, surfaces)	3.33	1.74	1.67566	0.00238
	random (explicit, holes in global unit)	**175.32**	0.36	1.67338	0.00267
	random (explicit, surfaces in global unit)	**131.19**	1.70	1.67676	0.00217
	random (randommix block)	32.82	0.32	1.67560	0.00249
SCALE/Shift	homogenized	0.41	0.30	1.56916	0.00243
ORANGE	lattice	**1.26**	**0.32**	1.67347	0.00352
	random (mesh overlay, holes)	**1.59**	0.37	1.67632	0.00299
	random (mesh overlay, surfaces)	7.77	1.71	1.67205	0.00257
	random (explicit, holes in global unit)	**49.31**	0.36	1.67529	0.00366
	random (explicit, surfaces in global unit)	**214.65**	1.25		
	random (randommix block)	25.86	0.31	1.67317	0.00280
SCALE/KENO-VI	homogenized	0.30	0.36	1.56625	0.00098
KENO	lattice	1.10	0.37	1.67463	0.00099
	random (mesh overlay, surfaces)	183.0	0.69		
Serpent	homogenized	0.21	0.38	1.56716	0.00074
Delta tracking	lattice	**0.29**	0.48	1.67442	0.00065
	random (explicit)	**0.30**	0.49	1.67661	0.00067
Serpent	homogenized	0.28	0.38	1.56728	0.00077
Surface tracking	lattice	**1.38**	0.48	1.67468	0.00067
	random (explicit)	**1.90**	0.48	1.67574	0.00079

By default, Serpent uses Woodcock delta tracking. Delta tracking provides significant advantages, especially for double-heterogeneous systems with thousands of fuel particles, because particles travel over material boundaries [7]. For the TRISO particle lattice or the random TRISO particle distribution, the calculations using delta tracking took only 15%–20% of the time required by those using surface tracking.

In addition to delta tracking, Serpent uses a unionized energy grid constructed by merging the energy grids of all nuclides in the problem [8]. This approach can significantly reduce computation time because it minimizes the number of grid-search iterations. However, because the model investigated in this study has only a small number of materials and nuclides, the benefits of using a unionized grid are not significant in this particular case.

Despite using the same MC parameters, Shift shows standard deviations a factor of 3 larger than Serpent. The improvement in the standard deviation estimator used by Serpent is believed to be a difference in k_{eff} estimators, and improvements to Shift are planned in the near future. A more in-depth study of the true standard deviation should be performed using different random seeds to understand the true variance between these cases instead of the estimated variance.

4.4 PERFORMANCE CONCLUSIONS

In general, Serpent outperforms Shift for models with TRISO particles randomly distributed in a fuel component. The primary reasons for shorter computation times in Serpent are (1) the use of Woodcock delta tracking, which is especially efficient for systems with a large number of cells, and (2) the use of the same unionized energy grid for all point-wise cross sections. Shift achieves shorter computation times when using the SCALE/Shift randommix block, which applies certain restrictions and reduces the number of geometry checks performed. Performance competitive with Serpent was obtained using Shift when overlaying a coarse mesh over the fuel-component region. Different options for improving the Shift computation time for these double-heterogeneous models will be explored. An extension of the randommix block with an additional overlaid mesh could contribute to shortening the computation time with a small modeling effort.

5. CRAB SUPPORT

As part of the coordination with Idaho National Laboratory (INL), Argonne National Laboratory, and the NRC, initial steps were taken to collaborate on using `Shift` and `Griffin` for the CRAB workflow. A NEAMS collaboration workshop was held on May 5, 2021, at which the `Shift` capabilities and the two-step neutronics work currently under way were presented to relevant parties. Interest was expressed in future development efforts for integrating `Shift` and `Griffin` into the CRAB workflow to target specific advanced reactor analysis.

6. CONCLUSIONS AND FUTURE WORK

The work performed in FY21 under NEAMS for Shift non-LWR development has yielded (1) a new Shift user interface for easier modeling that takes advantage of a new geometry engine in Shift; (2) a performance study of an HTR-10 fuel pebble comparing Shift, KENO-VI, and Serpent; (3) tally optimizations and enhancements for more accurate non-LWR cross section generation; and (4) initial collaboration on using Shift for CRAB support. These efforts have identified a variety of areas for future research that would benefit TRISO fuel modeling efforts and non-LWR applications with Shift. Below is a summary of proposed future efforts.

The enhanced Shift user interface will continue development with the following priorities: test against criticality benchmarks, implement parallel support via input and the Shift API to take advantage of Shift's existing parallelism, and further simplify and optimize the input.

Additionally, the unionized-energy grid feature in Shift will be enabled to decrease computation time for TRISO models. The randommix block will be extended with an added overlaid mesh option to further decrease the computation time for TRISO models. Hybrid or full Woodcock delta tracking will be implemented in Shift. The detailed list of Shift tally optimizations needed for usability and accuracy of cross section generation for two-step neutronics workflow will be added (see Section 2 for full list).

7. ACKNOWLEDGMENTS

The authors would like to thank the Griffin developers and users for their invaluable assistance. SCALE models for the HTR-10 fuel pebble used here were developed as part of previous studies with support from Cihangir Celik (ORNL) and Eva Davidson (ORNL).

This research made use of INL computing resources, which are supported by the US Department of Energy's Office of Nuclear Energy and the Nuclear Science User Facilities under Contract No. DE-AC07-05ID14517.

REFERENCES

[1] Tara Pandya, Matthew Jessee, Rike Bostelmann, and Javier Ortensi. Demonstration of Two-Step Neutronics Capability with Shift and Griffin. Technical Report ORNL/TM-2021/2048, Oak Ridge National Laboratory, 2021.

[2] Two-Step Neutronics Calculations with Shift and Griffin for Advanced Reactor Systems. *Annals of Nuclear Energy*, Draft.

[3] Friederike Bostelmann, Cihangir Celik, Mark L. Williams, Ronald J. Ellis, Germina Ilas, and William A. Wieselquist. SCALE Capabilities for High Temperature Gas-Cooled Reactor Analysis. *Annals of Nuclear Energy*, 147:107673, 2020.

[4] William A. Wieselquist, R. A. Lefebvre, and Matthew A. Jessee. SCALE Code System, Version 6.2.4. Technical Report ORNL/TM-2005/39, Oak Ridge National Laboratory, Oak Ridge, TN, 2020.

[5] Jaakko Leppanen, Maria Pusa, Tuomas Viitanen, Ville Valtavirta, and Toni Kaltiaisenaho. The Serpent Monte Carlo Code: Status, Development and Applications in 2013. *Annals of Nuclear Energy*, 82:142–150, 2015.

[6] William K. Terry, Leland M. Montierth, Soon Sam Kim, Joshua J. Cogliati, and Abderrafi M. Ougouag. Evaluation of the Initial Critical Configuration of the HTR-10 Pebble-Bed Reactor. Technical Report HTR10-GCR-RESR-001, NEA/NSC/DOC(2006)1, Rev. 0, OECD/NEA, 2007.

[7] Jaakko Leppänen. Performance of Woodcock Delta-Tracking in Lattice Physics Applications Using the Serpent Monte Carlo Reactor Physics Burnup Calculation Code. *Annals of Nuclear Energy*, 37(5):715–722, 2010.

[8] Jaakko Leppänen. Two Practical Methods for Unionized Energy Grid Construction in Continuous-Energy Monte Carlo Neutron Transport Calculation. *Annals of Nuclear Energy*, 36(7):878–885, 2009.