Performance of the Initial Implementation of the Shift Monte Carlo Code in SCALE 6.3[*]

W.J. Marshall and T.M. Greene
Oak Ridge National Laboratory
1 Bethel Valley Road
Oak Ridge, TN 37831-6170
marshallwj@ornl.gov, greenetm@ornl.gov

[Digital Object Identifier (DOI) placeholder – to be added by ANS during production]

ABSTRACT

The Shift Monte Carlo code will be introduced in SCALE 6.3 as an alternative to the KENO V.a, KENO-VI, and Monaco codes. Calculations were performed to establish the performance of Shift for criticality safety analyses within the criticality safety analyses sequence (CSAS) based on models in the Verified, Archived Library of Inputs and Data (VALID). This test suite contains over 600 critical experiment models covering a broad range of fissile materials and neutron energy spectra. The comparisons presented include calculated k_{eff} values and runtime performance for serial calculations and a selection of parallel calculations. Comparisons are presented for multigroup (MG) and continuous-energy (CE) calculations for KENO V.a and KENO-VI models.

Results generated with a beta version of SCALE 6.3 indicate excellent agreement in k_{eff} values between KENO and Shift. The largest differences in the average k_{eff} value calculated for the 15 categories of KENO V.a models are 0.00020 ± 0.00011 Δk for MG calculations and 0.00011 ± 0.00005 Δk for CE calculations. Similar comparisons in three categories using KENO-VI result in the largest differences for MG calculations: as 0.00004 ± 0.00003 Δk, and -0.00002 ± 0.00003 Δk for CE.

The preliminary results also indicate that Shift is faster than KENO on a per particle basis, especially for fast spectrum systems. The uncertainty per history is also higher, however, so the Monte Carlo figure of merit is higher for KENO for thermal and intermediate spectrum systems. As expected, Shift generally has better speedup than KENO for parallel calculations, regardless of neutron energy spectrum.

Key Words: KENO, Shift, validation, verification, eigenvalue

1 INTRODUCTION

Versions of the KENO Monte Carlo code have been used in criticality safety analyses since the 1960s; improving its usability for infrequent users was one of the reasons the SCALE code system was created in the late 1970s. Shift is a new Monte Carlo transport code developed recently at Oak Ridge National Laboratory (ORNL) which is being included alongside the KENO V.a and KENO-VI transport codes in SCALE 6.3. Plans call for retirement of the KENO codes in SCALE 7. Inclusion of all three transport codes

---

[*] Notice: This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (http://energy.gov/downloads/doe-public-access-plan).

in the SCALE 6.3 release provides users with a period to test Shift, provide feedback for improvement, and acquaint themselves with the new code. This paper presents an initial set of calculations using the benchmark models included in the Verified, Archived Library of Inputs and Data (VALID) [1] maintained at ORNL. The library contains 618 critical configurations from the International Handbook of Evaluated Criticality Safety Benchmark Experiments (ICSBEP Handbook) [2] and has been used in recent validation reports for SCALE criticality safety calculations [3].

The results indicate generally excellent agreement between the KENO and Shift results using both KENO V.a and KENO-VI in both 252-group multigroup (MG) and continuous-energy (CE) calculation modes. The k_{eff} agreement appears to be slightly better for CE calculations than for MG calculations. Across 15 categories of experiments modeled in KENO V.a, the largest difference in average MG-calculated k_{eff} values is $0.00020 \pm 0.00015\ \Delta k_{\mathrm{eff}}$ for the LEU-SOL-THERM experiments. The CE results show the largest difference for U233-MET-FAST experiments, with the average k_{eff} differing by $0.00011 \pm 0.00005\ \Delta k$. For the three categories of systems modeled in KENO-VI, the largest differences are for the HEU-MET-FAST experiments, and the differences are 0.00004 ± 0.00003 and $-0.00002 \pm 0.00003\ \Delta k$ for MG and CE calculations, respectively. More detailed examination of these results is presented in Section 3, following a brief discussion of the changes necessary to convert criticality safety analyses sequence (CSAS) input to run with Shift instead of KENO in Section 2.

This testing also revealed some minor differences between the KENO and Shift implementations that cause a small number of the models in VALID to fail in Shift without modification. Some of these are related to geometry package difference, some are caused by differences in the starting source approach, and one is caused by a difference in the implementation of the uncertainty-controlled calculation termination. These issues are discussed in Section 4. More than 93% of current VALID inputs run with the Shift transport code.

The last comparison discussed here is the two codes' run-times. In general, Shift performs particle transport faster. Unfortunately, the uncertainty resulting from the same number of particles is higher in Shift than in KENO, especially for thermal systems. The resulting figure of merit (FoM) is generally higher for KENO than Shift, except for fast systems. Shift has been designed for efficient parallel calculations, and it outperforms KENO in parallel performance. These results are presented in Section 5.

2 USING SHIFT WITHIN THE SCALE CSAS SEQUENCES

Shift has been implemented into SCALE so that existing CSAS5 and CSAS6 inputs will run successfully. The only change required in an existing input is to append "-shift" to the existing CSAS sequence declarator. An example for the simplified PU-MET-FAST-001 benchmark from the ICSBEP Handbook is shown in Figure 1.

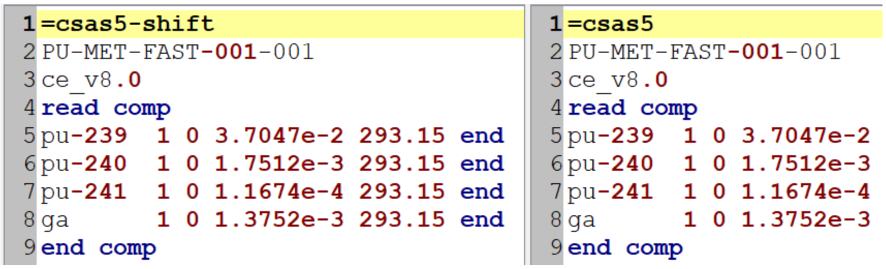

```
=csas5-shift
PU-MET-FAST-001-001
ce_v8.0
read comp
pu-239  1 0 3.7047e-2 293.15 end
pu-240  1 0 1.7512e-3 293.15 end
pu-241  1 0 1.1674e-4 293.15 end
ga      1 0 1.3752e-3 293.15 end
end comp
```

```
=csas5
PU-MET-FAST-001-001
ce_v8.0
read comp
pu-239  1 0 3.7047e-2
pu-240  1 0 1.7512e-3
pu-241  1 0 1.1674e-4
ga      1 0 1.3752e-3
end comp
```

Figure 1. Input change to select Shift transport (left) instead of KENO (right).

The input and cross section processing and output generation are largely unchanged, so the user experiences in these regards are quite similar. All data libraries are supported by both KENO and Shift transport. Many KENO edit options are not yet implemented in Shift, so the text output typically contains less information. In general, reported parameters are formatted in the same way to minimize impacts on user-generated post-processing scripts. Shift generates several plots to assist users with source convergence assessment. In addition, Shift generates an HDF5 archive with a significant amount of detailed information from the calculation. This file can be viewed or processed with any available HDF5 viewers or utilities.

3 *K*EFF RESULTS

VALID contains models of 618 evaluated critical experiments from the ICSBEP Handbook [3]. These experiments cover a combination of different fissile nuclides, physical forms, and neutron energy spectra, with 15 different ICSBEP benchmark categories represented: three categories are modeled with experiments in KENO-VI, and benchmarks from all 15 categories are modeled in KENO V.a. A complete list of experiments is provided in Table 3 of Saylor et al. [3]. The results presented in this paper consider both MG and CE particle transport with both KENO and Shift.

Of the 618 models included in VALID, 577 can be reliably executed with the CSAS-Shift sequences without modification. The remaining 41 cases require minor modifications to account for issues associated with geometry, initial source sampling, or uncertainty-controlled termination. These issues are discussed in more detail in Section 4. The results presented here come from the 577 models that did not require modifications, although the changes required could have been made and results incorporated. These cases are excluded so that results from only identical inputs are considered in this assessment.

The results presented here consider only the stochastic uncertainty associated with the Monte Carlo particle transport solution, and not uncertainties resulting from the experiments or their evaluation. This is judged to be the most appropriate uncertainty because the primary comparison here is of two codes with fixed models, and not a validation of either code to measurements. The results presented are also average k_{eff} values, and not calculated-to-expected (C/E) ratios, for the same reason: the calculations are used to quantify performance of the Shift code relative to the relevant KENO code. The results presented here could be recast as validation if the k_{eff} values were normalized with the expected benchmark values and the associated uncertainties were incorporated. It is expected that the SCALE 6.3 validation report will include C/E values and will provide the initial validation of Shift within the SCALE CSAS sequences.

3.1 KENO V.a Models

The 15 different ICSBEP categories are represented here with 521 KENO V.a benchmark models. The number of cases in each category ranges from 2 (MIX-COMP-FAST) to 123 (LEU-COMP-THERM); the average k_{eff} values for each code are reported for each category of experiments, along with the difference between the two average values. The difference in the average k_{eff} values and its uncertainty is shown in bold for cases in which the difference is more than 2 standard deviations, and it is shown in italics if the difference is between 1 and 2 standard deviations. The results for the MG calculations are provided in Table I and for CE calculations in Table II.

The results in Table I show generally good agreement between Shift and KENO. The difference between the average k_{eff} values is less than one standard deviation for 5 categories, between one and two standard deviations for 5 categories, and greater than two standard deviations for 5 categories. The largest magnitude difference is 0.00020 ± 0.00011 Δk_{eff} (LEU-SOL-THERM), and the largest relative difference is 2.57 standard deviations (U233-SOL-INTER). Across all categories, the calculated k_{eff} values are within one standard deviation for 348 cases, between one and two standard deviations for 140 cases, between two and three standard deviations for 30 cases, and more than three standard deviations apart for 3 cases. The cumulative percentages for this distribution are 66.8% within one standard deviation, 93.7% within two standard deviations, and 99.4% within three standard deviations. These cumulative distributions are very close to the theoretical distribution expected from a normal distribution, indicating only random differences between the sets of results. Overall, this assessment indicates very good agreement between Shift and KENO V.a for MG calculations of k_{eff}.

Table I. Average k_{eff} values for KENO V.a models, MG calculations

Evaluation Category	No. of Cases	Shift		KENO		Difference (S − K)	
		Avg k_{eff}	σ	Avg k_{eff}	σ	Avg k_{eff}	σ
HMF	23	1.00266	0.00002	1.00258	0.00002	**0.00008**	**0.00003**
HST	52	0.99907	0.00002	0.99905	0.00002	0.00002	0.00002
IMF	11	1.00316	0.00003	1.00312	0.00003	0.00004	0.00005
LCT	123	0.99921	0.00002	0.99916	0.00001	**0.00005**	**0.00002**
LST	19	0.99870	0.00009	0.99850	0.00006	*0.00020*	*0.00011*
MCF	2	0.99137	0.00008	0.99122	0.00007	*0.00015*	*0.00011*
MCT	21	0.99867	0.00003	0.99867	0.00002	0.00000	0.00003
MST	10	0.99326	0.00004	0.99329	0.00003	-0.00002	0.00005
PMF	12	0.99965	0.00003	0.99957	0.00003	*0.00008*	*0.00004*
PST	81	0.99756	0.00001	0.99754	0.00001	*0.00002*	*0.00002*
UCT	3	0.99924	0.00007	0.99922	0.00006	0.00001	0.00009
UMF	10	0.99883	0.00004	0.99871	0.00003	**0.00011**	**0.00005**
USI	29	0.97938	0.00002	0.97930	0.00002	**0.00008**	**0.00003**
USM	8	0.97540	0.00004	0.97533	0.00004	*0.00007*	*0.00006*
UST	117	0.99672	0.00001	0.99668	0.00001	**0.00004**	**0.00001**

Table II. Average k_{eff} values for KENO V.a models, CE calculations

Evaluation Category	No. of Cases	Shift		KENO		Difference (S − K)	
		Avg k_{eff}	σ	Avg k_{eff}	σ	Avg k_{eff}	σ
HMF	23	1.00143	0.00002	1.00141	0.00002	0.00002	0.00003
HST	52	0.99831	0.00002	0.99830	0.00002	0.00000	0.00003
IMF	11	1.00095	0.00003	1.00101	0.00003	*-0.00007*	*0.00005*
LCT	123	0.99873	0.00002	0.99875	0.00001	-0.00002	0.00002
LST	19	0.99867	0.00009	0.99858	0.00007	0.00009	0.00011
MCF	2	0.98813	0.00009	0.98809	0.00007	0.00003	0.00011
MCT	21	0.99891	0.00003	0.99890	0.00002	0.00001	0.00003
MST	10	0.99357	0.00004	0.99354	0.00003	0.00002	0.00005
PMF	12	0.99939	0.00003	0.99942	0.00003	-0.00003	0.00004
PST	81	0.99769	0.00001	0.99772	0.00001	**-0.00004**	**0.00002**
UCT	3	0.99915	0.00008	0.99914	0.00006	0.00000	0.00010
UMF	10	0.99871	0.00004	0.99860	0.00003	**0.00011**	**0.00005**
USI	29	0.97946	0.00003	0.97945	0.00002	0.00001	0.00004
USM	8	0.97552	0.00006	0.97546	0.00004	0.00006	0.00007
UST	117	0.99735	0.00001	0.99738	0.00001	**-0.00004**	**0.00002**

The results in Table II show excellent agreement between Shift and KENO for CE calculations. The difference between the average k_{eff} values is less than one standard deviation for 11 categories, between one and two standard deviations for one category, and greater than two standard deviations for three categories. The largest difference in average k_{eff} values is 0.00011 ± 0.00005, 2.31 standard deviations, for the U233-MET-FAST category. When pooled together, the differences between individual k_{eff} values are less than one standard deviation for 340 cases, between one and two standard deviations for 137 cases, between two and three standard deviations for 35 cases, and greater than 3 standard deviations for 9 cases. The cumulative distribution is 65.3% within one standard deviation, 91.6 % within two standard deviations, and 98.3% within three standard deviations. The result is a mismatch in which the average k_{eff} values are in better agreement for CE calculations than MG calculations, but the distribution of individual k_{eff} values is better for MG calculations than for the CE results. It is also worth noting that 7 of the 9 CE calculations that differ by more than 3 standard deviations are in the U233-SOL-THERM category, and an additional case is in the U233-SOL-INTER category. The overall summary is one of generally excellent agreement in calculated k_{eff} between Shift and KENO V.a for CE calculations.

3.2 KENO-VI Models

KENO V.a and KENO-VI differ primarily in their geometry description, with KENO V.a including a restricted geometry package for faster execution, and KENO-VI containing a generalized geometry description for representing more complicated models. Shift has only one native geometry description and one set of transport algorithms, so no difference is expected in Shift performance. Previous validation studies for KENO V.a and KENO-VI have shown that both codes generate equivalent results for the same models [3]. The comparison of Shift and KENO-VI is at this point more limited than for KENO V.a, because only 57 KENO-VI models are included in VALID; 56 of these models are included in this comparison. The average k_{eff} values for both codes using MG and CE transport are provided in Table III. The formatting is consistent with Tables I and II: the one category with a difference of more than one standard deviation is shown in italics. None of the three categories represented with KENO-VI differ by more than 2 standard deviations with either MG or CE transport.

Table III. Average k_{eff} values for KENO-VI models

Evaluation Category	No. of Cases	Shift		KENO		Difference (S − K)	
		Avg k_{eff}	σ	Avg k_{eff}	σ	Avg k_{eff}	σ
MG Results							
HMF	27	0.99751	0.00002	0.99746	0.00002	*0.00004*	*0.00003*
IMF	2	1.00377	0.00009	1.00380	0.00007	-0.00003	0.00011
MCT	27	0.99544	0.00002	0.99542	0.00002	0.00002	0.00003
CE Results							
HMF	27	0.99765	0.00002	0.99767	0.00002	-0.00002	0.00003
IMF	2	1.00484	0.00009	1.00482	0.00007	0.00002	0.00011
MCT	27	0.99581	0.00002	0.99579	0.00002	0.00001	0.00003

The results show excellent agreement between Shift and KENO, and again slightly better agreement in average k_{eff} values with CE transport than with MG transport. The HEU-MET-FAST category has the largest difference, both absolute and relative, in average k_{eff} values for CE calculations and the largest relative difference for the MG calculations. Based on only two models, the IMF results show a larger absolute difference in average k_{eff} values for the MG results. The individual case results are also in good agreement. Considering MG calculations, 37 cases have differences of less than one standard deviation, 16 are between one and two standard deviations, and the remaining three have differences of more than two standard deviations. For the CE calculations, 40 cases are within one standard deviation of each other, and the remaining 16 are between one and two standard deviations apart. These distributions also agree well with theoretical results, although with only 56 points, they are less reliable than the KENO V.a distributions. Overall, the results of the KENO-VI models in VALID support the conclusion that Shift-calculated k_{eff} values are very similar to KENO results. This is not a surprising result considering the same nuclear data are used in both transport codes, but it lends further credence to the belief that most code bias is a result of nuclear data errors and is not related to the geometry representation or physics models in the codes.

4 SHIFT DISCREPANCIES

Three categories of issues have been identified for direct use of KENO models with Shift transport in SCALE 6.3. Each of these is discussed briefly in this section, along with the simple modifications needed to resolve the problems and allow reliable results to be generated with Shift. The first problem is related to the uncertainty-based termination of a calculation and was only encountered in one KENO-VI calculation. The second difference is related to the implementation of controls on the sampling of initial source sites; this issue impacted 23 U233-SOL-THERM models. The final discrepancy is the most general and relates to handling of small offsets in the geometry; here, it affected 17 LEU-COMP-THERM cases.

4.1 Uncertainty-Based Execution Termination

Shift supports termination of execution when the cumulative k_{eff} reaches a user-specified uncertainty in much the same manner as KENO V.a and KENO-VI in recent SCALE releases. The desired uncertainty is specified using the *sig*= value in the PARAMETER block. KENO requires at least twice as many active generations as skipped generations to be run before terminating the calculation with this parameter, regardless of the cumulative k_{eff} uncertainty. Shift has not implemented this same control, and in one case, the first four active generations were very similar. This resulted in an estimated uncertainty that was less than the specified uncertainty level. This is an unlikely scenario, but it cannot necessarily be avoided without changing the method of termination. A very short runtime is the first indication that this has occurred, and an examination of the output file will indicate that only a very small number of generations were executed. Changing the random number seed has been shown to resolve the issue. A series of 9 calculations with different random number seeds ran between 1,188 and 1,305 active generations to achieve the same desired uncertainty that the base calculation had achieved after only 4 generations. An efficient solution to this problem will be implemented in the future, although it is not yet clear if it will be the same as the rule used by KENO.

4.2 Initial Source Site Selection

The source sites for the first generation of an iterated-source Monte Carlo transport simulation must be selected according to some process, because no previous generation exists from which to select them. A range of source descriptions exist for KENO in the CSAS sequences, although only a limited number of these have been implemented in Shift for the SCALE 6.3 release. Both KENO and Shift default to uniform sampling over the model volume for the initial source sites. Transport codes typically include a check to ensure that source sites can be located for the first generation. KENO has historically used a generation time limit for this function, with a user-specified time per generation which can be set using the *tba*= value in the PARAMETER block. This parameter is defaulted to 10 minutes. However, the Shift implementation selects up to a fixed number of random sites. Execution is terminated if no sites have been selected in fissionable material before the set number of points has been sampled. The number of source sites to sample is not a user-modifiable parameter in SCALE 6.3. This difference can be encountered in models with a small fraction of the overall volume containing fissionable material and in prematurely terminated calculations for 23 U233-SOL-THERM systems. These models have small solution tanks for which room return is an important model feature, resulting in a large model volume with a small fraction of fissionable material. The solution to this difficulty is to provide a better starting guess in the START block. Shift implementation in SCALE 6.3 supports both start type 0 and start type 6 in CSAS. Start type 0 selects points uniformly within a volume; by default, this volume is the entire problem volume. A smaller volume can be specified as a cuboid in the START block, thus ensuring that fissile material is sampled quickly. Start type 6 allows for the specification of individual points in the global coordinate system for initial source sites. Either option will allow execution of the entire simulation. The relevant VALID cases were modified with start type 0 with a cuboid around the solution tank to allow Shift execution of the models.

4.3 Small Geometry Discrepancies

The identified discrepancy that seems most likely to be a generic issue relates to small discrepancies in the geometry specification. When KENO encounters such small gaps, typically on the order of 0.00001 cm, the particle is advanced forward a small distance and again assesses its location. Historically, the distance of this forward bump has been tailored to accommodate most discrepancies that are not the result of a geometry error. Small gaps of this type have typically been introduced in some KENO models to eliminate coincident surfaces, especially between holes and other geometry components. This is often a direct result of the "hole error" identified in SCALE 5 in 2005 [4]. Shift, on the other hand, identifies that such close surfaces exist and issues a warning message. This leads to an error message almost immediately in transport, as shown in Figure 2. The solution is straightforward, if not always quick and easy: the nearly coincident surfaces must be identified and made coincident. This fix allows Shift to track particles through the geometry without difficulty and to generate a reasonable k_{eff} value. The difficulty of this modification is obviously model dependent, and in some cases, it may be quite laborious. It is hoped that future developments in the geometry package will remove this limitation. All 17 cases in the LEU-COMP-THERM-008 evaluation within VALID have this discrepancy, as do ORNL-generated models of the GBC-32 [5] and GBC-68 [6] computational benchmark casks.

```
!!! Geometry error in particle 0:2: In universe 'unit 41'->'unit 13'-
>'u13a3'->'unit 1': Failed to exit volume EXTERIOR (local volume #3) at local
point {-0.81789,-0.571888,147.038} along {0.0390323,-0.318898,-0.946985}
 ^^^ at {-49.89069,-13.6579675066892,147.038101127412} along
{0.0390323181251486,-0.318898325993042,-0.946984865676645}
```

Figure 2. Error message generated by Shift caused by near-coincident surfaces.

5 TIMING RESULTS

A final aspect of performance to be documented here is the runtime of Shift compared to that of KENO. Shift has been designed for efficient operation in parallel execution, even for massively parallel calculation on leadership-class computing platforms. Criticality safety calculations do not require this level of computing, so an assessment of performance for serial calculations and more modest parallel calculations are presented here, again using the results from the calculations of the benchmarks from VALID. The primary focus in this paper is to analyze these serial results for a large number of calculations. A limited study of parallel performance for one fast and one thermal system—both based on KENO V.a models—is also documented here.

5.1 Serial Calculations

The runtime and k_{eff} uncertainty information has been collected from the CSAS calculations discussed in Section 3 and analyzed in this effort. As with the agreement in k_{eff} values, there are differences in the results for MG and CE calculations. The CE results are provided in Table IV, and the MG results are provided in Table V. In both tables, the runtime ratios are on the left side of the table, and the FoM results are on the right. The runtime ratio is calculated as KENO runtime divided by Shift runtime, so the numbers greater than unity indicate longer runtimes for KENO and better performance for Shift. The FoM ratios are calculated as Shift FoM divided by KENO FoM. This is the opposite of the runtime ratio, but it simplifies the comparisons by ensuring that numbers greater than 1 indicate better performance for Shift. For each category of systems, the average ratio is provided, along with the maximum and minimum individual case ratios, to provide some indication of the range of results. These results lump the KENO V.a and KENO-VI results together for the HEU-MET-FAST, IEU-MET-FAST, and MIX-COMP-THERM categories; further studies for potential differences between the two will likely be performed and published in the future.

Table IV. Runtime comparisons for Shift and KENO, CE calculations

Evaluation Category	Runtime Ratio (K/S)			FoM Ratio (S/K)		
	Avg.	Min.	Max.	Avg.	Min.	Max.
HMF	1.97	0.86	4.08	1.56	0.76	2.92
HST	0.91	0.37	1.54	0.44	0.14	0.72
IMF	2.34	0.77	3.27	1.66	0.70	2.28
LCT	0.78	0.40	1.47	0.53	0.29	0.89
LST	0.77	0.50	1.17	0.52	0.34	0.86
MCF	2.44	2.26	2.63	1.45	1.32	1.57
MCT	0.52	0.35	0.79	0.40	0.23	0.66
MST	0.87	0.72	1.01	0.52	0.38	0.63
PMF	3.22	2.36	4.40	2.67	2.08	4.06
PST	0.79	0.46	1.19	0.46	0.31	0.70
UCT	0.75	0.64	0.92	0.43	0.33	0.54
UMF	2.97	2.17	3.78	2.41	1.82	3.42
USI	0.93	0.57	1.20	0.41	0.31	0.63
USM	1.06	0.93	1.20	0.43	0.35	0.52
UST	1.12	0.56	2.16	0.52	0.30	0.97

Table V. Runtime comparisons for Shift and KENO, MG calculations

Evaluation Category	Runtime Ratio (K/S)			FoM Ratio (S/K)		
	Avg.	Min.	Max.	Avg.	Min.	Max.
HMF	1.01	0.45	2.57	0.81	0.32	2.52
HST	0.49	0.13	0.92	0.38	0.10	0.67
IMF	1.22	0.52	2.09	0.87	0.52	1.47
LCT	0.43	0.17	0.94	0.46	0.15	0.90
LST	0.91	0.43	1.56	0.58	0.24	0.89
MCF	1.19	1.01	1.37	0.86	0.85	0.87
MCT	0.54	0.31	1.04	0.44	0.21	0.84
MST	0.67	0.51	1.09	0.45	0.38	0.60
PMF	2.31	1.43	4.16	1.87	1.23	3.08
PST	0.85	0.46	1.46	0.67	0.36	1.21
UCT	0.40	0.32	0.54	0.25	0.21	0.31
UMF	2.49	1.28	3.93	1.95	0.98	2.59
USI	0.71	0.54	0.98	0.55	0.39	0.82
USM	0.87	0.76	0.92	0.66	0.57	0.79
UST	0.78	0.45	1.10	0.62	0.25	0.95

For CE calculations, Shift runtimes are less than KENO runtimes for fast systems and are generally slightly slower for thermal systems. The larger uncertainty also represents a significant penalty in the FoM calculation, and Shift performance is better only for fast spectrum systems. Thermal spectrum systems are much more efficient in KENO than in Shift.

For MG calculations, Shift is again generally faster for fast spectrum systems, although the difference is less than that for CE calculations. Thermal systems are faster in KENO, and again, there is a clear difference from the CE calculations: KENO has a greater speed-up in MG mode than Shift. Future optimization may improve MG performance in Shift. As with the CE results, consideration of uncertainty gives a stronger edge to KENO. KENO is clearly more efficient for thermal systems, and it is even more efficient for some of the fast spectrum systems.

The serial runtime results summarized in this section indicate that Shift is generally faster than KENO for fast spectrum systems, whereas KENO is as fast or faster for the thermal and intermediate spectrum systems in the VALID library. KENO also results in lower uncertainty than Shift for the same number of particles, increasing the relative efficiency of KENO. KENO is clearly more efficient than Shift for thermal systems, especially for MG calculations; Shift generally has higher performance than KENO for fast spectrum systems.

5.2 Parallel Calculations

Two systems were selected to investigate parallel performance of KENO V.a and Shift: PU-MET-FAST-006-001 and LEU-COMP-THERM-080-001. These choices were largely arbitrary, but one fast spectrum system and one thermal spectrum system were selected given the differences in code performance as a function of spectrum. The number of neutrons per generation was not increased, because the number of cores was increased, although in general it is recommended that the number of neutrons per core be maintained and the number of neutrons per generation be maximized. This approach minimizes message passing time and aids parallel efficiency. In this case, both codes were handicapped in the same way. The thermal system used 40,000 neutrons per generation, and the fast system used 10,000. The results are shown in Figure 3, indicating that Shift parallel performance is generally better than KENO V.a. While Shift speedup is better, the runtime is still lower for KENO than for Shift for LEU-COMP-THERM-080-001 on 96 cores. The Shift runtime advantage for PU-MET-FAST-006-001 grows larger as the number of cores increases. A more extensive study is needed to fully quantify relative parallel performance, and KENO-VI should also be considered in future studies.

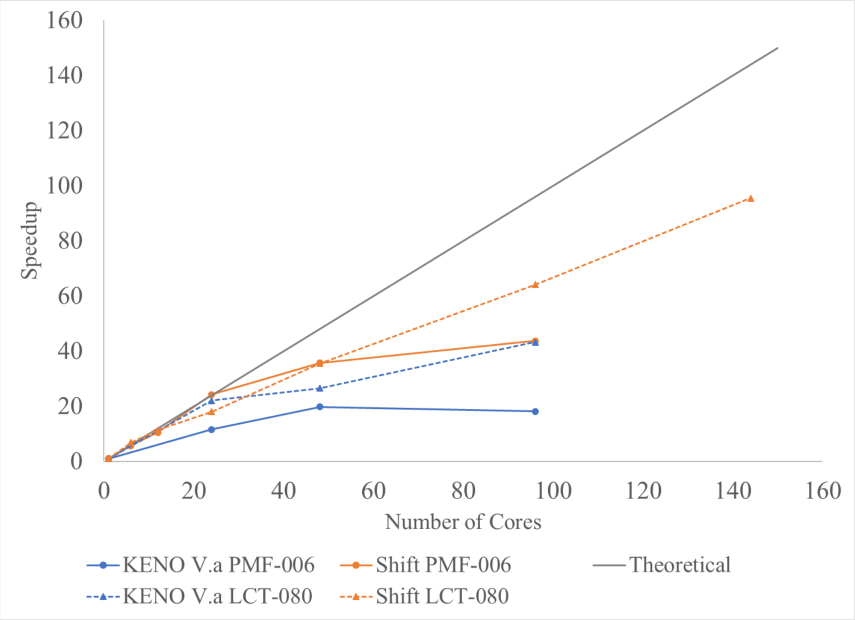

Figure 3. Parallel performance of KENO V.a and Shift.

6 CONCLUSIONS

This paper presents an assessment of the initial deployment of Shift within the CSAS sequences of SCALE 6.3. The calculated k_{eff} values are very similar between Shift and the KENO V.a and KENO-VI Monte Carlo codes deployed in SCALE 6.3 and prior releases. The differences in the average k_{eff} values according to the ICSBEP experiment category are in good agreement, and the distribution of individual differences closely follows the normal distribution expected for only random differences in the results. Shift can be viewed as generating largely equivalent results to KENO in both MG and CE modes.

A few discrepancies have been identified that may require modification of existing CSAS inputs before they can be executed with Shift transport. The issue most likely to cause significant work to solve relates to Shift's inability to handle surfaces that are nearly coincident. Unfortunately, surfaces with this type of small spacing have been built into some models specifically to protect against a code error in prior versions of KENO. Other identified issues relate to the initial guess for fission sites and the calculation termination when controlled with uncertainty. These issues can be readily addressed.

Shift runtime performance is generally favorable compared to that of KENO for fast spectrum systems, especially for CE calculations. KENO runtime performance is superior for thermal systems. The resulting uncertainty is also generally lower for KENO than for Shift, so the FoM performance changes the balance for both CE and MG calculations more in favor of KENO. Parallel calculations indicate that Shift scales better for the number of processors typically used in criticality safety calculations. This improved parallel performance is generally not sufficient to overcome the lower performance for thermal systems. Optimization of MG calculations and neutron scattering physics may improve Shift performance in future releases.

7 ACKNOWLEDGMENTS

This work was supported by the Nuclear Criticality Safety Program, funded and managed by the National Nuclear Security Administration for the US Department of Energy. Additional support has been provided by the US Nuclear Regulatory Commission office of Nuclear Regulatory Research.

8 REFERENCES

1. W. J. Marshall and B. T. Rearden, "The SCALE Verified, Archived Library of Inputs and Data – VALID," *Proceedings of the ANS Nuclear Criticality Safety Division Topical Meeting (NCSD2013),* Wilmington, NC, September 29 – October 4, 2013 (2013).
2. *International Handbook of Evaluated Criticality Safety Benchmark Experiments*, NEA/NSC/DOC(95)03, Organisation for Economic Co-operation and Development, Nuclear Energy Agency, Paris, France (2021).
3. E. M. Saylor, W. J. Marshall, J. B. Clarity, Z. J. Clifton, and B. T. Rearden, *Criticality Safety Validation of SCALE 6.2.2*, ORNL/TM-2018/884, Oak Ridge National Laboratory, Oak Ridge, TN (2018).
4. "Potentially Non-Conservative Error in Modeling Geometric Regions in the KENO V.a Criticality Code," NRC Information Notice 2005-13, US Nuclear Regulatory Commission, Washington, DC (2005).
5. J. C. Wagner, *Computational Benchmark for Estimation of Reactivity Margin from Fission Products and Minor Actinides in PWR BUC*, NUREG/CR-6747 (ORNL/TM-2000/306) Washington, DC (2001).
6. D. E. Mueller, J. M. Scaglione, J. C. Wagner, and S. M. Bowman, *Computational Benchmark for Estimation of Reactivity Margin from Fission Products and Minor Actinides in BWR Burnup Credit*, NUREG/CR-7157 (ORNL/TM-2012/96) Washington, DC (2013).