ANS NCSD 2013 - Criticality Safety in the Modern Era: Raising the Bar
Wilmington, NC, September 29–October 3, 2013, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2013)

ENHANCEMENTS IN CONTINUOUS-ENERGY MONTE CARLO CAPABILITIES IN SCALE

K. B. Bekar, C. Celik, D. Wiarda, D. E. Peplow, B. T. Rearden, and M. E. Dunn

Oak Ridge National Laboratory
P.O. Box 2008, M.S. 6170, Oak Ridge, TN 37831 USA
bekarkb@ornl.gov

ABSTRACT

Continuous energy Monte Carlo tools in SCALE are commonly used in criticality safety calculations, but new developments in SCALE 6.2 provide enhanced continuous-energy capabilities as well in sensitivity and uncertainty, depletion, and criticality accident alarm system (CAAS) analyses. Recent improvements in the continuous-energy data generated by the AMPX code system and significant advancements in the continuous-energy treatment in the KENO Monte Carlo eigenvalue codes facilitate the use of SCALE Monte Carlo codes to model geometrically complex systems with enhanced solution fidelity. The addition of continuous-energy treatment to the Monaco fixed-source Monte Carlo code, which can be used with automatic variance reduction in the hybrid MAVRIC sequence, provides increased solution fidelity for shielding applications and CAAS modeling. In addition, the Monte Carlo depletion capability in SCALE has been extended from multigroup to continuous energy. This paper describes some of the advancements in continuous-energy Monte Carlo codes within the SCALE code system.

Key Words: SCALE, Monte Carlo, KENO, Continuous Energy

1 INTRODUCTION

SCALE nuclear systems modeling and simulation software is developed and maintained by Oak Ridge National Laboratory (ORNL) to perform analyses of criticality safety, reactor physics, radiation shielding, and spent fuel characterization for nuclear facilities and transportation/storage package designs [1].

The SCALE code system includes the KENO V.a and KENO-VI Monte Carlo (MC) criticality codes; both offer continuous-energy (CE) and multigroup (MG) energy treatments to calculate physical parameters of fissile systems. The major difference between these two codes is their geometry processors; KENO V.a uses very simple geometric components, while KENO-VI uses the SCALE Generalized Geometry Package, which allows more complex geometric modeling. With the introduction of on-the-fly reaction rate tallies and few-group microscopic reaction cross-section calculations, SCALE provides significant capabilities to perform CE depletion and CE sensitivity and uncertainty (S/U) analyses in addition to criticality safety analyses [2, 3]. The addition of CE treatment to the fixed-source MC radiation transport code, Monaco, which can be used with automatic variance reduction in the MAVRIC hybrid deterministic/MC sequence, provides significant enhancements, especially for shielding analyses and criticality accident alarm system (CAAS) modeling [4].

Although these added features enhance the CE capabilities of SCALE for several nuclear applications, their impacts on code performance are significant. One bottleneck of CE MC calculations that can limit the code performance is the memory requirement of the CE data. Reducing this memory footprint enables the SCALE MC codes to be used for a wider range of nuclear analysis applications. In addition to these improvements, some of which directly affect solution accuracy, parallel computation capabilities have been added to KENO to provide reductions in wall clock time, especially for S/U analysis or MC depletion.

This paper summarizes the recent enhancements in CE MC capabilities in SCALE and discusses improvements in code performance that can be accomplished by reducing the memory footprint of CE MC calculations while retaining calculational accuracy in the results. In addition, several new KENO features, which will be available in SCALE 6.2, are presented.

2 ENHANCEMENTS IN CE DATA AND TRANSPORT

The SCALE code system uses nuclear data libraries generated by AMPX, which processes ENDF-formatted nuclear data evaluations to provide CE, MG, and covariance data libraries [5]. In previous SCALE releases, the extent of CE data libraries has been limited by only providing CE neutron data libraries with a specific reaction subset of the ENDF libraries for KENO MC calculations. By improving AMPX capabilities, the contents of the CE data libraries have been extended to support a wide range of reactions for both neutron and gamma interactions and to produce gamma yield data from neutron interactions. This generalized form of the new CE data library uses CE Monte Carlo particle transport in SCALE for various nuclear applications.

2.1 CE Neutron Data

Revisions to CE neutron data have improved both S(α,β) data, resulting in significant bias reduction for thermal systems, and the probability tables that provide CE treatment in the unresolved resonance range, resulting in reduced biases for systems that are sensitive to the intermediate energy range. The updated CE neutron data mitigates the systematic bias trend that was observed for the MC benchmarks with SCALE 6.1 CE data. Additional details are provided in a companion paper on the validation of SCALE 6.2 [6].

2.2 CE Photon Data

In addition to revised neutron libraries, AMPX has been enhanced to generate CE photon data for gamma and coupled neutron-gamma calculations. Photon libraries have the same format as the neutron libraries and are relatively small compared to neutron libraries as data are only available for elements, not nuclides, and only a select set of reactions are required. Coherent scattering, incoherent scattering, pair production, and photoelectric absorption reactions are included in the libraries for each element. Photoelectric absorption is treated as a terminal process with no secondary particles (such as characteristics x-rays) generated.

2.3 Gamma Yield Data

AMPX has also been extended to provide gamma yield data from CE neutron interactions for coupled neutron-photon particle transport. Neutron libraries have been extended to include the associated gamma yield data for available reactions for each nuclide when used in coupled calculations. Eigenvalue calculations skip the gamma yield data because they do not need to

account for the gammas. The representation of yield data provided by ENDF can vary based on the nuclide and reaction type, such that a total gamma yield might be given instead of gamma yield data for a specific reaction (e.g., inelastic gamma yields might be given on the total inelastic reaction, mt = 4, rather than provided separately for each individual discrete level and the continuum). If the gamma yields are requested in the CE physics engine, gammas will be produced after each collision using the associated gamma yield data.

2.4 New CE Physics Engine

A new SCALE package named SCEMPP (SCALE Continuous-Energy Modular Physics Package) has been developed to serve as the CE collision physics engine. SCEMPP models particle collisions in a material and generates the particle(s) resulting from a collision. SCEMPP is essentially an event generator for SCALE. SCEMPP has Fortran and C++ application programming interfaces (APIs) to support both legacy and future developments in SCALE. SCEMPP requires cross sections and kinematics data for each isotope/element and implements the SCALE CE Resource package, which is another new feature in SCALE 6.2 to read and store all the nuclear data from AMPX CE particle libraries and to transfer information via APIs. Communication through APIs, unlike traditional I/O operations in SCALE, and SCEMPP's generic modular structure enable a flexible and powerful integration within modernization and new code developments in SCALE. In addition to creating collision particles, SCEMPP also provides non-transport data, such as reaction responses or point detector data, to Monaco to enable dose calculations and point detector tally estimates.

In SCALE 6.2, SCEMPP is integrated to the Monaco code and MAVRIC sequence to provide CE particle transport for shielding and CAAS analyses.

2.5 CE Transport Updates

An important aspect of CE physics is the energy at which scattering reactions are transitioned from bound S(α,β) data to free-gas treatment. The cutoff energy for the thermal neutron transport treatments in KENO is represented by a single energy for all nuclides. The cutoff value has been revised from 3 eV (default in the previous SCALE releases) to a new default value of 10 eV. Neutron upscatter below the thermal cutoff is modeled by either S(α,β) data or the free-gas treatment. The new cutoff value was selected based on preliminary testing results from models in the Verified, Archived Library of Inputs and Data (VALID) library [7]. Those results indicate that CE neutron transport with a 10 eV thermal cutoff produces more accurate results than the use of the 3 eV cutoff.

In addition, algorithms used to calculate the initialization of the problem-dependent fission spectrum data (the combination of problem-dependent fission spectrum from all fissionable nuclides that is used to generate initial source particles and to sample particles if a fission event appears in particle tracking) for KENO have been improved to provide further enhancements in some cases.

3 ADVANCEMENTS AND NEW FEATURES FOR KENO

Several new features have been introduced in the KENO codes to improve their performance for criticality safety applications and to use them for CE depletion and CE sensitivity and uncertainty analysis.

3.1 On-the-fly Mixture Cross Section Capabilities

Continuous-energy calculations in KENO through SCALE 6.1 have been performed using a "unionized energy grid," where material-dependent cross-section data are generated for each user-defined mixture. The energy grid on which the cross-section data are stored for each mixture is based on a unionization of the individual energy grids for each nuclide in the mixture. Therefore, the storage requirements for the cross-section data in any calculation can increase substantially as new materials are added, particularly if the materials contain a large number of isotopes, as is the case with spent nuclear fuel. The resulting memory requirements can make performing CE calculations intractable for some problems, such as those with multiple spent fuel compositions that could require >200 GB of memory.

A new "on-the-fly mixture cross-section calculation" capability has been introduced as an option in KENO to disable the use of the unionized energy grid approach for all mixtures in the problem. KENO calculations with this new option provide a reduction in memory requirements that can be as much as an order of magnitude, depending on the number and complexity (i.e., number of isotopes required to define a material) of the materials used in the model. This new feature increases the runtime for KENO CE calculations by ~20% or more for most cases but provides the capability to simulate very large problems with multiple mixture configurations that was impossible in SCALE 6.1.

A benchmark suite that includes a collection of problems of varying sizes for different applications was designed to test the performance and accuracy of KENO calculations with various conditions/options. One of the sample sets in this benchmark suite is a simplified KENO core model of ORNL's High Flux Isotope Reactor (HFIR) with different numbers of mixture definitions. This sample set was used to investigate the impact of the number of mixture definitions (each with many isotopes) in a complex model on code performance for an eigenvalue calculation. This sample set was run with MG KENO, CE KENO with the unionized energy grid (UUM=YES), and CE KENO with the new on-the-fly mixture cross-section calculation option (UUM=NO).

Table I presents the timing results and memory requirements of the three KENO calculations (MG KENO, CE KENO with the unionized energy grid option, and CE KENO with the on-the-fly mixture cross section option, respectively). In addition to these results, the eigenvalues for all calculations (k_{eff}), where it was possible for KENO to successfully run to completion, are listed in this table. The computed uncertainty for all eigenvalues is approximately 70 pcm. The results show that CE KENO with the on-the-fly mixture option is capable of simulating all models in the test suite, even the case with 500 mixture definitions. KENO with both the MG treatment and the CE treatment with unionized energy grid approach failed to complete the case with 500 mixture definitions because of the excessive memory requirements. The CPU time for the calculations increases with the number of mixture definitions as expected, but the CE KENO runtimes with the new option are actually less than those for CE KENO cases with the unionized energy grid for 50 and 100 mixtures. Note the dramatic increases in memory requirements for the MG and

CE unionized grid cases as the number of mixtures increases, while the CE on-the-fly cases require almost constant memory regardless of the number of mixtures. These results demonstrate that the on-the-fly mixture option should be used for KENO models that include a large number of mixtures.

Table I. Comparison of KENO code performance. Results were obtained from a simplified KENO HFIR model for different number of mixture definitions in the model

Case	# of mix.	MG KENO-VI			CE KENO-VI (UUM=YES)			CE KENO-VI (UUM=NO)		
		k_{eff}	CPU Time (min)	Memory[a] (GB)	k_{eff}	CPU Time (min)	Memory (GB)	k_{eff}	CPU Time (min)	Memory (GB)
1	5	0.99250	25.30	1.22	0.99066	42.78	3.47	0.99230	60.19	1.03
2	10	0.99283	38.18	2.41	0.99162	53.76	5.69	0.99219	60.11	1.05
3	50	0.99285	159.46	11.90	0.99141	146.95	23.95	0.99175	144.42	1.06
4	100	0.99308	309.65	22.30	0.99300	305.74	46.52	0.99146	275.53	1.06
5	500	-	-	>60.00	-	-	>200.00	0.99161	6373.70	1.10

[a]Memory allocation for MG KENO-VI calculation also includes the memory requirement of cross-section processing tools.

3.2 Reduction in Memory Requirement of CE Internal Storage

Internal storage of CE cross-section data has been converted from double precision to single precision, which results in a further 15–45% reduction in memory footprint, depending on the problem, with no loss of precision in computed results.

Figure 1 illustrates the percent reduction in memory footprint in KENO after this implementation for seven different model problems from the benchmark suite when running KENO with both options: the unionized energy grid (UUM=YES) and the on-the-fly mixture cross section option (UUM=NO). The gain is significant, especially for KENO with UUM=YES because the unionization process needs much more memory space compared to KENO with UUM=NO.

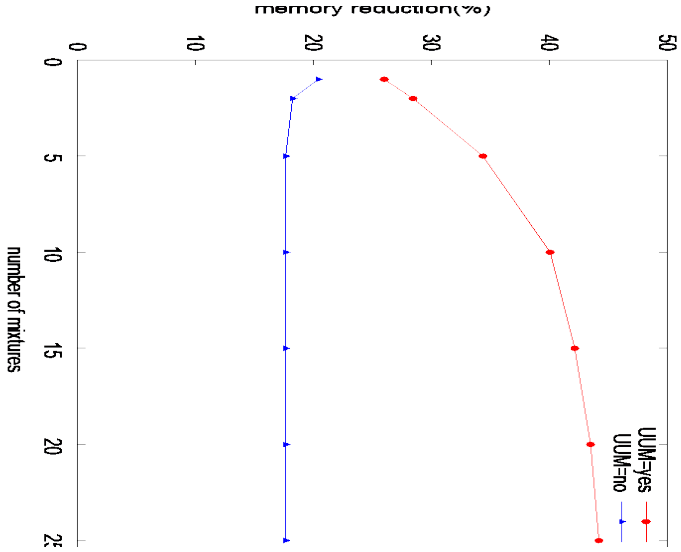

Figure 1. Percent memory reduction after changing internal storage array from double precision to single precision.

3.3 Multigroup Cross Section and Reaction Rate Calculations in CE Mode

To perform depletion calculations with SCALE, accurate few-group cross sections must be calculated for all transmutation reactions at each time step for use in the ORIGEN depletion module. Previous versions of SCALE provide MG Monte Carlo depletion capabilities using the standard region-wise neutron spectra for collapsing multigroup cross sections in a post-processing stage that follows KENO MC calculations. In some cases, the few-group cross sections suffer from limitations of the MG approximation such as inadequate group structure or the inability to properly shield the cross sections for the problem using one-dimensional resonance self-shielding modules.

To address these potential problems, a new few-group microscopic reaction cross-section calculation capability has been added to the KENO codes. This new method produces multigroup cross sections and reaction rates directly in CE mode calculations rather than using a post-processing approach. In each generation, KENO uses track length estimators for the reaction rate tallies for all isotopes in specified regions. At the end of each generation, a subsequent calculation is performed to compute few-group microscopic reaction cross sections for all isotopes in a region as the ratio of the computed reaction rates to the flux averaged over this cell. Finally, KENO computes mean values and statistical uncertainties for all these quantities and saves them in a file, which can be used by ORIGEN.

Although these new tallies in CE mode are more expensive in terms of CPU time, the gain in accuracy makes this method a viable choice. Additionally, CE calculations benefit from the memory management enhancements previously described, where large numbers of depletion regions require substantially less memory in CE mode that in MG. This new method enhances SCALE depletion capabilities by adding a CE depletion option [2].

3.4 Doppler Broadening Rejection Correction

The thermal motion of target nuclides can significantly affect the collision between a neutron and nucleus in the epithermal energy range. As in most MC transport codes, KENO simulates this thermal motion with a free gas scattering model in its CE treatment. However, recent studies have shown that resonance scattering caused by the thermal motion of heavy target nuclides can have a measurable effect upon criticality calculations. The Doppler Broadening Rejection Correction (DBRC) method, which introduces corrections to the Doppler broadening of the scattering kernel with a new sampling equation, has been implemented in the KENO codes [8]. Implementing this method can have a significant impact on criticality calculations due to the increase of neutrons being upscattered into absorption resonances. As temperature increases, more neutrons are upscattered into the resonances, resulting in more absorptions and a lower eigenvalue (k_{eff}).

Results from light water reactor (LWR) pin cell calculations presented in Table II demonstrate the temperature effect on k_{eff} if the new DBRC feature is enabled in KENO. As the temperature increases, the impact of DBRC on k_{eff} becomes more significant in criticality calculations with KENO; the difference approaches 635 pcm at 2400 K. The KENO DBRC results are consistent with those predicted with MCNPX using DBRC.

Although the DBRC implementation in KENO currently is enabled only for ^{238}U, additional nuclides will be tested to evaluate their impact on the results.

Table II. Results of CE KENO with and without DBRC for several temperatures.

Temperature (K)	CE KENO	CE KENO with DBRC	Difference (pcm)
293.6	1.34460	1.34451	-9
600.0	1.33053	1.32932	-121
900.0	1.31759	1.31759	-182
1200.0	1.31029	1.30730	-299
2400.0	1.28113	1.27478	-635

3.5 Multiple Mesh Support

KENO codes in previous SCALE releases only support a single Cartesian mesh definition with a single mesh quantity/tally, which can be either mesh-flux calculations to derive sensitivity coefficients in an S/U analysis or fission source accumulations to use in CAAS analyses. The other limitation is that this single mesh definition must cover the entire geometry.

To enhance the mesh feature in KENO, a new mesh object was defined to support multiple mesh definitions for multiple mesh-based quantities such as mesh flux tally, mesh-based fission source generation, mesh-based fission source convergence diagnostics, and F*(r) mesh in CE sensitivity analysis [3]. The addition of the multiple-mesh capability enables the accumulation of multiple quantities on different grids, which can be defined as covering part of the geometry, rather than the entire geometry.

3.6 Fission Source Convergence Diagnostics

Fission source convergence diagnostic techniques have been implemented in KENO to provide improved confidence in the computed results as well as a reduction in the simulation time for some cases. Confirming convergence of the fission source in addition to k_{eff} is especially useful for flux tallies as needed for reaction rate calculations and S/U analysis.

The KENO source convergence diagnostics rely on Shannon entropy statistics of mesh-based fission source data [9, 10, 11]. By computing the Shannon entropy at each generation, the fission source distribution can be diagnosed in terms of randomness. Random fluctuations in the generation-to-generation Shannon entropy tally make it difficult to determine if and when the Shannon entropy of a system has converged; thus, several tests are typically applied when evaluating the Shannon entropy convergence of a system. Three different tests have been implemented in KENO as part of source convergence diagnostics.

1. **Test 1 – Final Convergence:** The first test assesses the convergence of the fission source at the end of the simulation. Here, the fission source convergence is determined through a comparison between the mean square posterior relative entropy (defined as the statistical distance between the binned fission source and the average fission source over the second half of the active generations) and the centered mean square Shannon entropy. This test states that the fission source is converged if the mean square Shannon entropy is less than the center mean square posterior relative entropy.

2. **Test 2 – First Converged Generation:** The second test verifies that the Shannon entropy of each active generation does not vary significantly from the average Shannon entropy of the system. This test, which should be met over the active generations of the calculation, is especially useful for reporting the generation at which the source converged.

3. **Test 3 – Adequate Active Generations:** The third test verifies that the average Shannon entropy of all the active generations does not differ significantly from the Shannon entropy of the last half of the active generations. This test is useful for detecting fission source convergence in problems where an inadequate number of inactive generations was sampled, as the Shannon entropy would continue to change during the active generations until it eventually converged.

KENO determines the distribution of the fission sites on a spatial mesh, either using a default mesh (5×5×5 mesh within a bounding box surrounding the whole geometry) provided by KENO or a user-specified mesh, before starting particle tracking within a generation. KENO then calculates the Shannon entropy for the current generation and stores it for the convergence test calculations. At the end of the calculations, KENO calculates the posterior entropy and then performs the three tests. The results of these tests and Shannon entropy results are reported in the KENO standard output.

The source convergence diagnostics of KENO were tested with a benchmark problem from the OECD NEA WPNCS Expert Group on Source Convergence. This benchmark problem comprising of three models, Cases 2.1–2.3, represents pin-cell arrays with irradiated LWR fuel

[12], were used to test the source convergence implementation in KENO. The composition of the LWR spent fuel consists of more reactive, low burnup end regions separated by a long, less reactive, high burnup region. The fuel composition differs in certain axial regions for each of the three cases; Case 2.1 has a symmetric isotopic distribution of the pin cell, while Cases 2.2 and 2.3 have higher burnups in one or more of the regions at the bottom of the pin cell. As the Expert Group specifically designed these models to present problematic convergence, a large number of generations are required to reach the converged source [12].

Table III presents the results of the three tests for each model. Figure 2 shows the Shannon entropy variation of the three cases during the simulations. After simulating 5,000 neutron generations (GEN) with 10,000 histories in each generation (NPG), all cases passed Test 1, failed Test 2 (based on 200 skipped generations), and passed Test 3. Case 2.1 was the slowest to converge, requiring 4,581 additional generations after 200 skipped generations. Case 2.2 immediately converges in the beginning of active generations; only 4 additional generations are required after 200 generations skipped. Case 2.3 required 552 additional generations for the source to converge.

Table III. Results of three tests for three models; NPG=10,000, GEN=5,000

Model	Case 2.1	Case 2.2	Case 2.3
Test 1 Final convergence	PASSED	PASSED	PASSED
Test 2 First converged generation (additional skipped generations required shown in parenthesis)	FAILED (4,581)	FAILED (4)	FAILED (552)
Test 3 Adequate active generations	PASSED	PASSED	PASSED

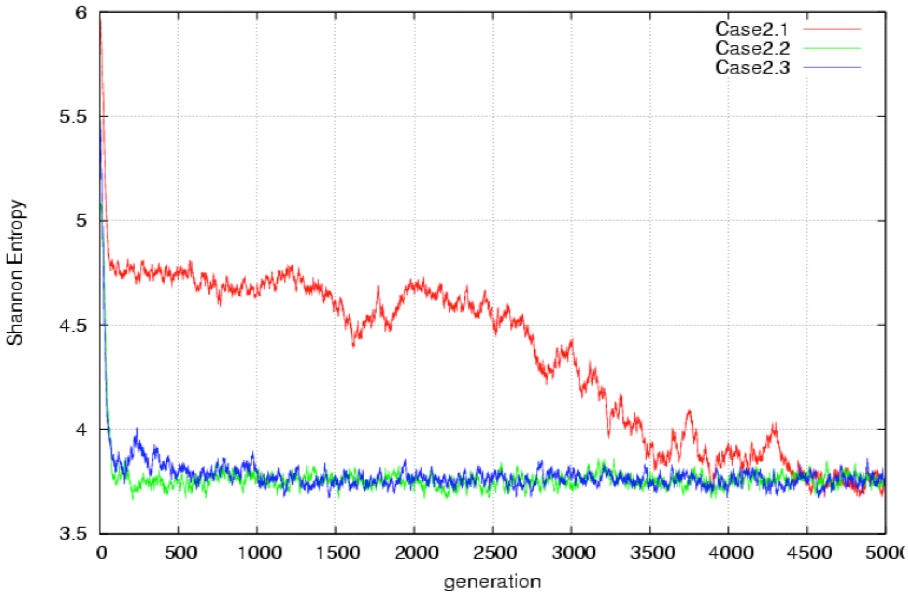

Figure 2. Shannon entropy variation for the three benchmark models.

3.7 Distributed Memory Parallelism via MPI

A standard parallel Monte Carlo algorithm, a simple master-slave approach via Message Passing Interface MPI [13], has been implemented to KENO for using parallel execution, especially for large problems; running problems is limited by the speed of a single processor with the serial version of KENO. In this approach, KENO runs different random walks concurrently on the replicated geometry within the same generation. Fission source and other tallied quantities are gathered at the end of each generation by the master process and are then processed either for final edits or subsequent generations. The KENO random number generator also has been updated to generate history-based random numbers to accomplish identical random walks with different numbers of processors for the same calculation. This feature is helpful especially for SCALE developers to test and verify the code integrity on different architectures with different compilers after a new implementation.

The performance of the MPI implementation in KENO was evaluated through a series of timing comparisons with test problems. In this section, speed-up plots show the "theoretical speed-up," which is proportional to the number of parallel MPI tasks. The theoretical speed-up is calculated as if the code is parallelized 100% (with no serial sections) and has no communication overhead. Therefore, the theoretical speed-up is equal to the number of parallel tasks (M).

Several tests in our benchmark suite were also used to observe the parallel performance of KENO. All these results indicate that the code performance depends on the problem type, its complexity, and requested tally calculations. The complexity of the problem increases the particle tracking time, thus increasing the parallel efficiency. In contrast, multiple tally requests degrade the code performance with increasing numbers of processors because the tallies are updated at the end of each generation, which increases the communication overhead among the processors.

Figure 3 illustrates the parallel performance of KENO-VI with a test in this benchmark suite that is used in a CE depletion calculation for a graphite-moderated reactor model. The parallel performance of KENO-VI is very good (greater than 93% of linear) up to eight processors. The parallel performance begins diminishing as a function of increased number of processors beyond that point but is still above 80% when using up to 20 processors.

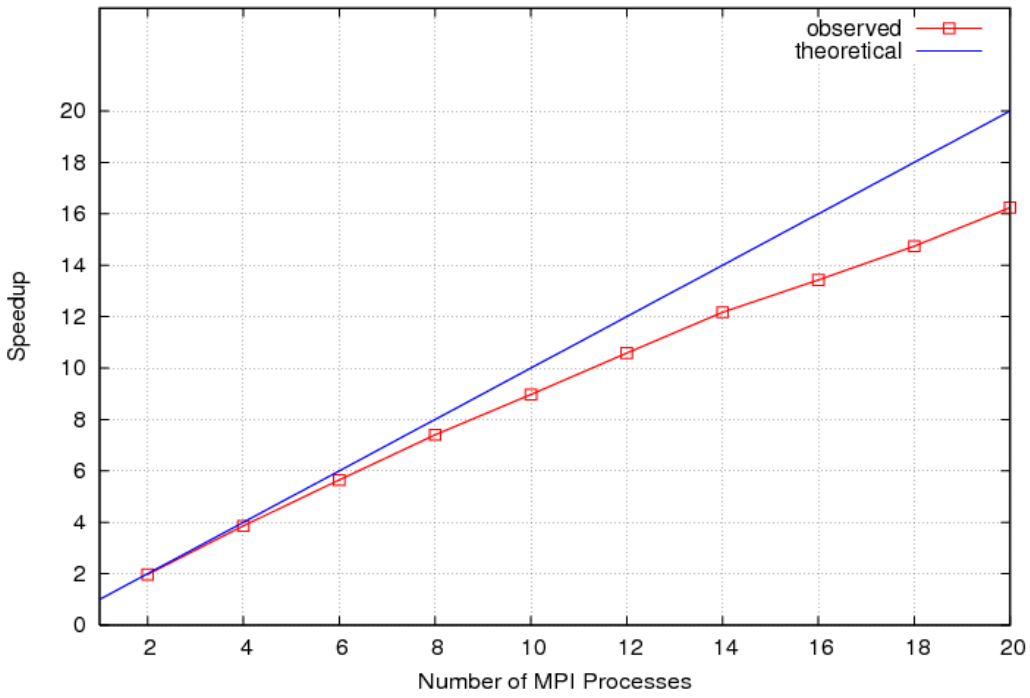

Figure 3. Parallel speed up of KENO-VI code in a depletion calculation.

4 CONCLUSIONS

New features have been implemented in the KENO codes to improve their capabilities for various nuclear applications as well as to improve the performance and accuracy of the codes. Changing the internal storage of CE cross-section data from double precision to single precision and introducing an optional on-the-fly mixture cross section calculation capability significantly reduces the memory requirement of KENO; this reduction is especially useful for systems with a significant number of mixtures. The parallel versions of KENO allow for rapid estimation of k_{eff}, fluxes, and reaction rates, which is especially useful for models that traditionally require very long run times. With new fission source convergence diagnostics based on the Shannon entropy of the fission source implemented in KENO, the user can monitor the convergence of both k_{eff} and the fission source distribution to better ensure accurate results. Recent improvements in the CE cross-section libraries and these significant advancements in the CE treatment in the KENO Monte Carlo eigenvalue codes facilitate the use of SCALE Monte Carlo codes to model geometrically complex systems with enhanced solution fidelity.

5 ACKNOWLEDGMENTS

This work was performed through the sponsorship of the U.S. Department of Energy (DOE) Nuclear Criticality Safety Program, the U.S. Nuclear Regulatory Commission, and the DOE Office of Nuclear Energy.

6 REFERENCES

1. *SCALE: A Comprehensive Modeling and Simulation Suite for Nuclear Safety Analysis and Design*, ORNL/TM-2005/39, Version 6.1, Oak Ridge National Laboratory, Oak Ridge, Tenn. (June 2011). Available from Radiation Safety Information Computational Center at Oak Ridge National Laboratory as CCC-785.

2. S. Goluoglu, K. B. Bekar, and D. Wiarda, "SCALE Continuous-Energy Monte Carlo Depletion with Parallel KENO in TRITON," *Trans. Am. Nucl. Soc.*, **106**, 723–725 (2012).

3. C. M. Perfetti and B. T. Rearden, "Use of SCALE Continuous-Energy Monte Carlo Tools for Eigenvalue Sensitivity Coefficient Calculations," *Proceeding of ANS NCSD 2013—Criticality Safety in the Modern Era: Raising the Bar,* Wilmington, N.C., September 29–October 3, 2013, on CD-ROM, American Nuclear Society, LaGrange Park, Ill. (2013).

4. D. E. Peplow, "Monte Carlo Shielding Analysis Capabilities with MAVRIC," *Nuclear Technology*, **174**, 289–313 (2011).

5. 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**, 118–119 (2002).

6. W. J. Marshall, D. Wiarda, C. Celik, and B. T. Rearden, "Validation of Criticality Safety Calculations with SCALE 6.2," *Proceeding of ANS NCSD 2013—Criticality Safety in the Modern Era: Raising the Bar,* Wilmington, N.C., September 29–October 3, 2013, on CD-ROM, American Nuclear Society, LaGrange Park, Ill. (2013).

7. W. J. Marshall and B. T. Rearden, "The SCALE Verified, Archived Library of Inputs and Data–VALID," *Proceeding of ANS NCSD 2013—Criticality Safety in the Modern Era:Raising the Bar*, Wilmington, N.C., September 29–October 3, 2013, on CD-ROM, American Nuclear Society, LaGrange Park, Ill. (2013)

8. S. Hart et al., "Implementation of the Doppler Broadening Rejection Correction in KENO," *Trans. Am. Nucl. Soc.,* **108**, 423–426 (2013).

9. T. Ueki and Forrest B. Brown**, "**Stationarity and Source Convergence Diagnostics in Monte Carlo Criticality Calculation," Nuclear Mathematical and Computational Sciences Conference (M&C 2003), Gatlinburg, Tenn., April 6–10, 2003.

10. T. Ueki and F. B. Brown, "Stationary Modeling and Informatics-Based Diagnostics in Monte Carlo Criticality Calculations," *Nuclear Science & Engineering* **149**, 38 (2005).

11. M. Wenner and A. Haghighat, “Study of Methods of Stationarity Detection for Monte Carlo Criticality Analysis with KENOV.a,” *Trans. Am. Nucl. Soc*., ***97,*** 647–651 (2007).

12. “Source Convergence in Criticality Safety Analyses,” OECD Report ISBN 92-64-02304-6, Phase I: Results for Four Test Problems, NEA No. 5431 (2006).

13. E. Gabriel et al., “Open MPI: Goals, Concept, and Design of a Next Generation MPI Implementation,” *Proceedings, 11th European PVM/MPI Users’ Group Meeting*, 97–104, Budapest (2004).