

## LA-UR-18-27656

Approved for public release; distribution is unlimited.

Title:                   Diagnostics for Undersampling and Clustering in Monte Carlo  
                                  Criticality Calculations

Author(s):            Brown, Forrest B.  
                                  Josey, Colin James

Intended for:        MCNP documentation  
                                  Report

Issued:                2018-08-10

---

**Disclaimer:**

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

# Diagnosics for Undersampling and Clustering in Monte Carlo Criticality Calculations

Forrest B. Brown, Colin J. Josey  
Monte Carlo Methods, Codes, & Applications Group, LANL

## 1.0 Introduction

For the past 40 years, it has been known from the theoretical work of Gelbard [1] and Brissenden [2] that the results from Monte Carlo (MC) criticality calculations are biased if the number of neutrons per generation is not sufficiently large. The computed k-effective eigenvalue will be lower than the actual eigenvalue, and the computed eigenfunction will be too low in high-importance regions and too high in low-importance regions. This difficulty is commonly called the undersampling problem for criticality calculations. The biases for undersampling were confirmed for practical, realistic problems with modern MC codes by Brown [3]. The solution to undersampling is simple – use a larger number of neutrons per generation to reduce the biases to levels that are negligible.

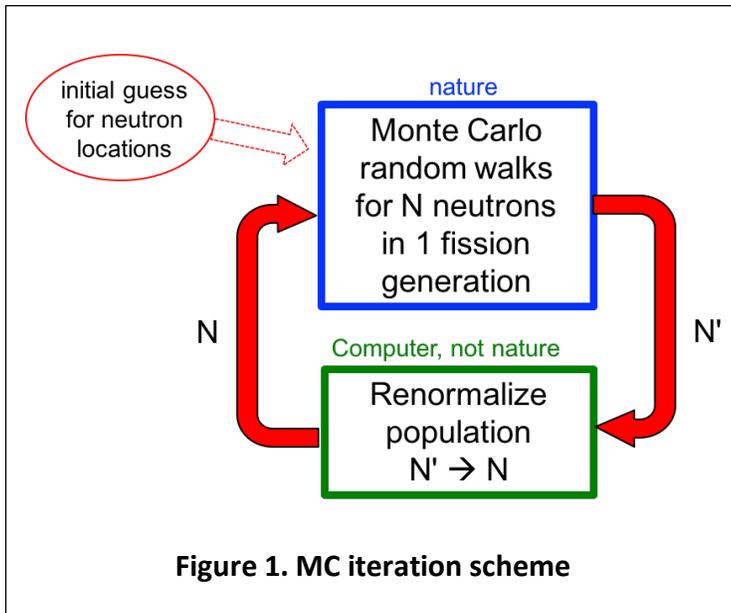
Over the past few years, the OECD-NEA-WPNCs Expert Group on Advanced Monte Carlo Techniques (EGAMCT) [4,5,6] investigated an extreme form of undersampling called clustering. When the number of neutrons per generation is small, neutrons in successive generations tend to group together (cluster) due to correlations, and the clusters tend to migrate through the problem space.

While the problems of undersampling and clustering are well known and well understood theoretically, there were no existing diagnostics for MC codes to determine whether a given calculation is using sufficient neutrons per generation to eliminate these problems. That is, no tools were available to diagnose biases due to undersampling or clustering. The tools described herein resolve this longstanding problem by providing 2 robust methods for detecting undersampling and clustering in MC criticality calculations.

The diagnosis of undersampling and clustering is part of a larger effort to thoroughly overhaul the calculational methods for criticality problems, providing for automated sampling of the initial fission distribution, acceleration of the iteration convergence process, automated detection of convergence in the iterations and starting of the active tally cycles, and the diagnosis of undersampling. All of these new methods have been prototyped in a local modified version of mcnp6.2 [7] and tested on a variety of critical systems.

## 1.1 Background on MC Criticality Calculations

The biases due to undersampling and the manifestation of clustering in extreme cases are strictly artifacts of the calculational algorithm used to solve MC criticality calculations. A generation-based iteration scheme is used, and the neutron population is renormalized between successive iterations, as shown in Figure 1.



Each generation is started with the same total number of neutrons  $N$  (or equivalently for MCNP, total weight) and a number of neutrons  $N'$  is produced.  $N'$  is a stochastic variable, hence the renormalization to  $N$  neutrons for the next generation is a biased process. Further, the renormalization reduces the number of independent neutron fission chains, introducing correlation [8]. The correlation manifests as clustering for very low  $N$ .

The bias in  $k$ -effective introduced by the renormalization process is negative and proportional to  $1/N$ . For small to moderate sized problems, it has been

shown that using  $N > 10,000$  neutrons per generation effectively removes the bias in  $k$ -effective [3]. For larger physical systems, 100,000 or 1 M or more neutrons per generation may be needed. The eigenfunction (i.e., fission distribution) is also biased proportional to  $1/N$ , with additional complications of being too low in high-importance regions and too high in low-importance regions.

The MC iteration scheme begins with an initial guess for  $k$ -effective and the fission distribution. Iterations (called inactive cycles) are performed without tallies until  $k$ -effective and the fission distribution have converged to their stationary state. After convergence, tallies are turned on, and iterations (called active cycles) are continued until sufficiently small uncertainties are obtained for desired results. Concerns for the different phases of the MC iteration process are illustrated in Figure 2 [3]. The diagnostic tests for undersampling cannot be performed until after  $k$ -effective and the fission distribution have converged to their stationary states.

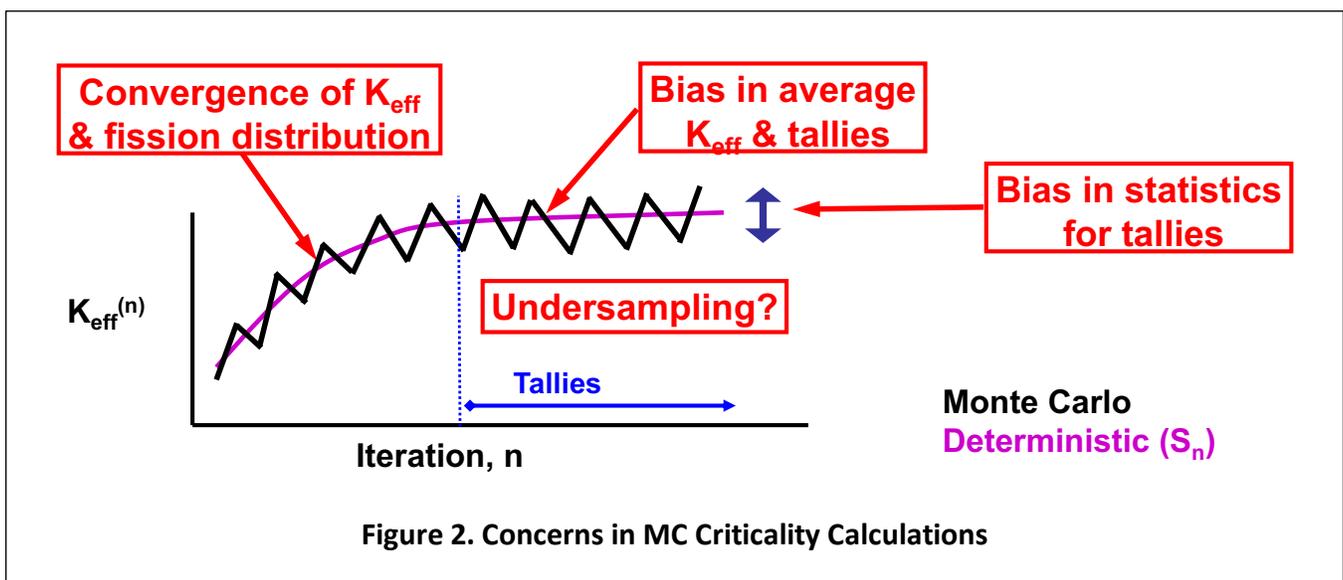
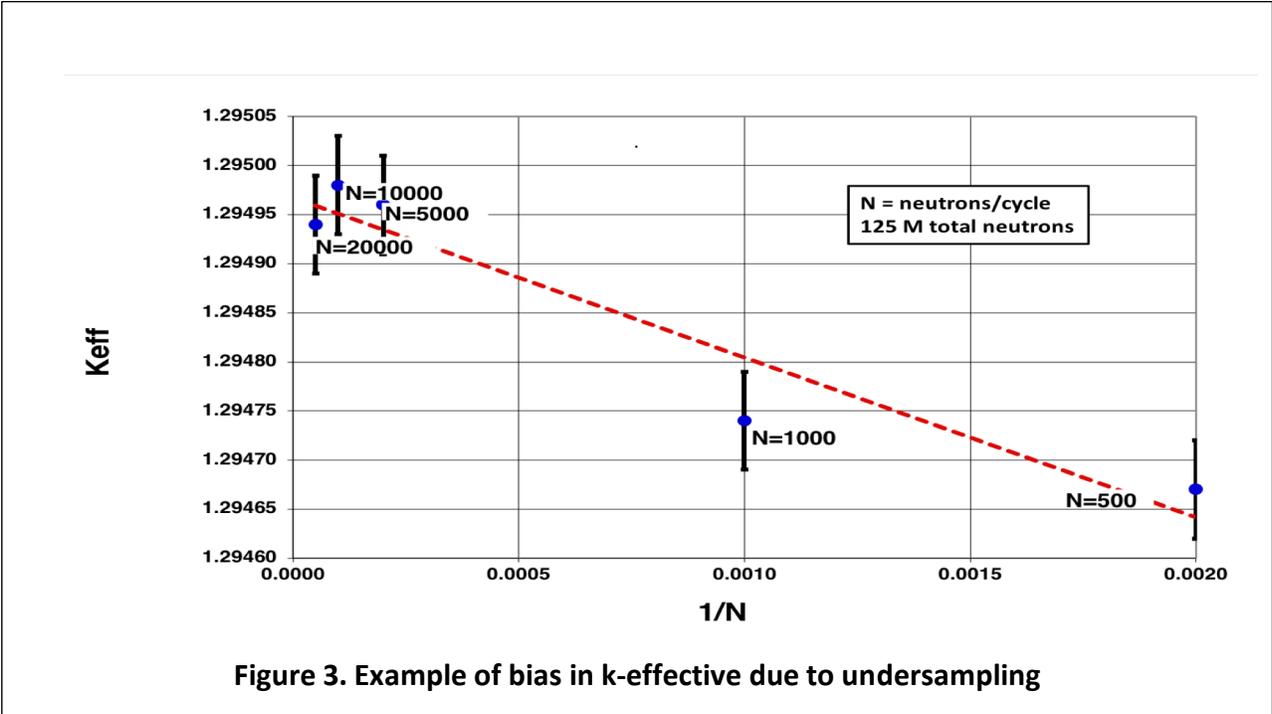
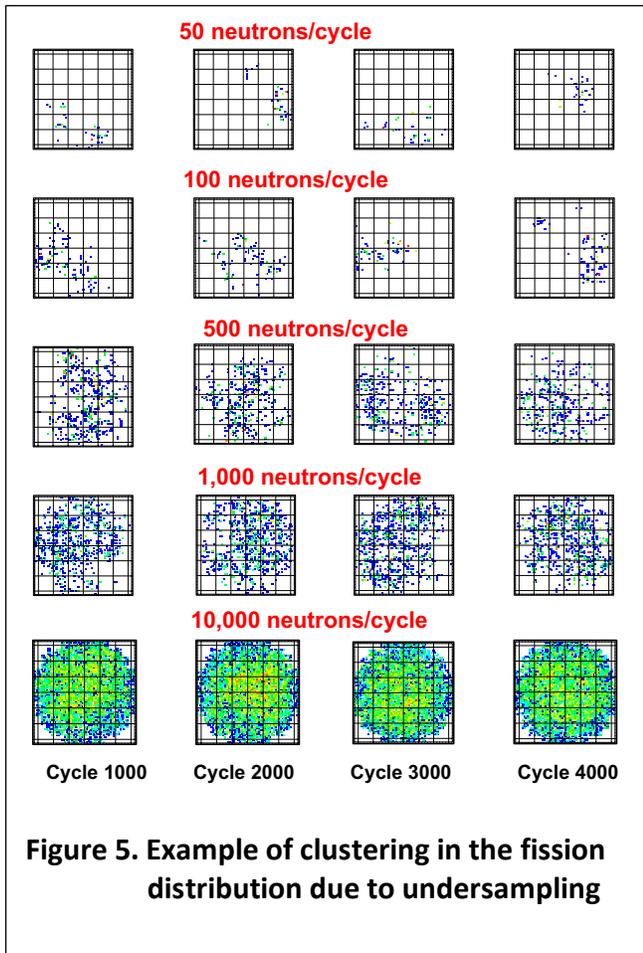
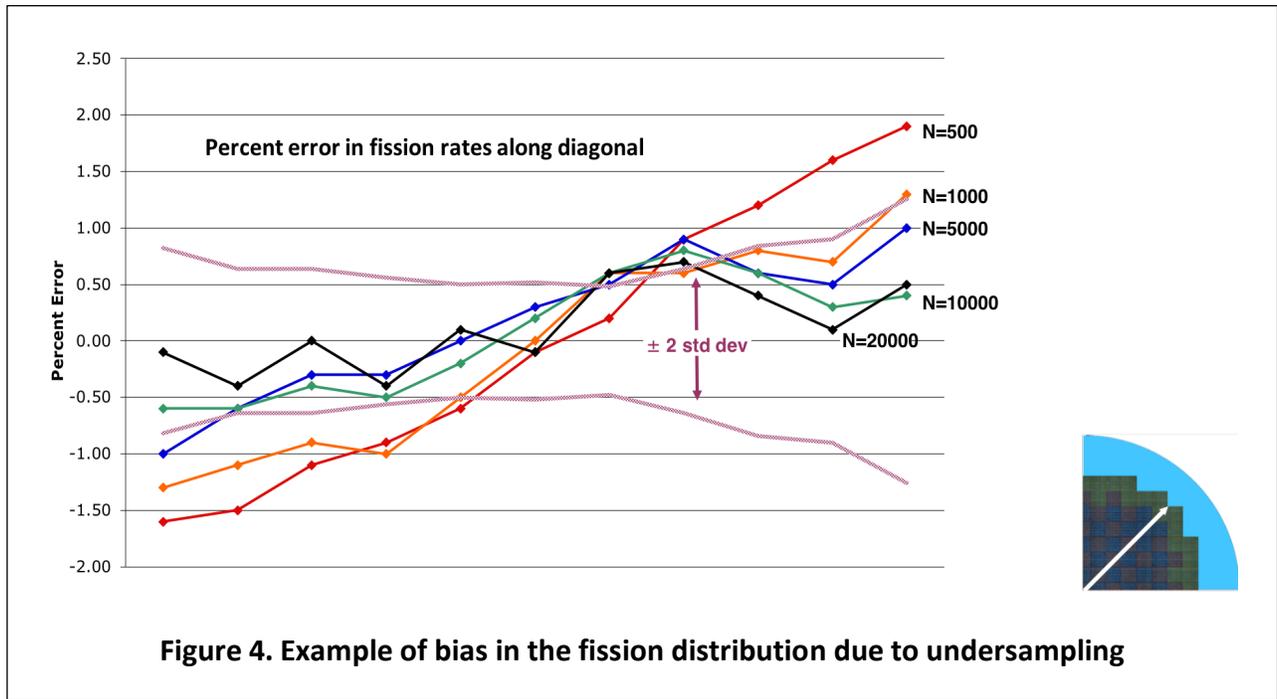


Figure 2. Concerns in MC Criticality Calculations

In the present work, we are concerned with diagnosing biases in k-effective and the fission distribution due to undersampling. Figures 3-5 are examples of the biases and clustering in a standard PWR model for different numbers of neutrons per generation [3]. Figure 3 clearly shows the nonconservative bias in k-effective and the 1/N dependence (indicated by the red dashed line). For this problem, at least 5,000 neutrons/cycle are needed to reduce the bias in k-effective to a negligible level compared to problem statistics. Figure 4 shows the corresponding bias in the fission distribution as a function of neutrons/cycle. For small N, there are significant biases in shape, with too few fissions in important regions and too many fissions in unimportant regions. With 10,000-20,000 neutrons/cycle, the shape bias is reduced and within problem statistics. A good analyst would use 50,000, 100,000, or more neutrons/cycle to avoid issues with biased fission distribution shapes. It is important to note that the biases illustrated in Figures 3 and 4 will not go away even if an infinite number of cycles are run. The biases are strictly due to the renormalization that is performed every cycle, and depend on N, the number of neutrons/cycle. The biases do not depend on the number of active cycles that are run in a problem. Thus, running an infinite number of neutrons with small N will still result in significant biases in k-effective and the shape of the fission distribution. The Figure 5 illustrates the extreme cases of undersampling, where clearly visible clusters of neutrons form and migrate through the problem space. Even with 10,000 neutrons/cycle, close examination of the fission distribution shows high and low spots migrating around, even though there are no large gaps, indicating that there is some undersampling.

The discussion above applies equally well to every criticality problem run with Monte Carlo using power iteration. Figures 3-5 illustrate the effects of the renormalization biases. Until now, there have been no reliable tests or metrics for diagnosing undersampling. The analysis and computer resources required for generating Figures 3-5 for a particular problem are prohibitive for routine, daily use of Monte Carlo for criticality calculations.





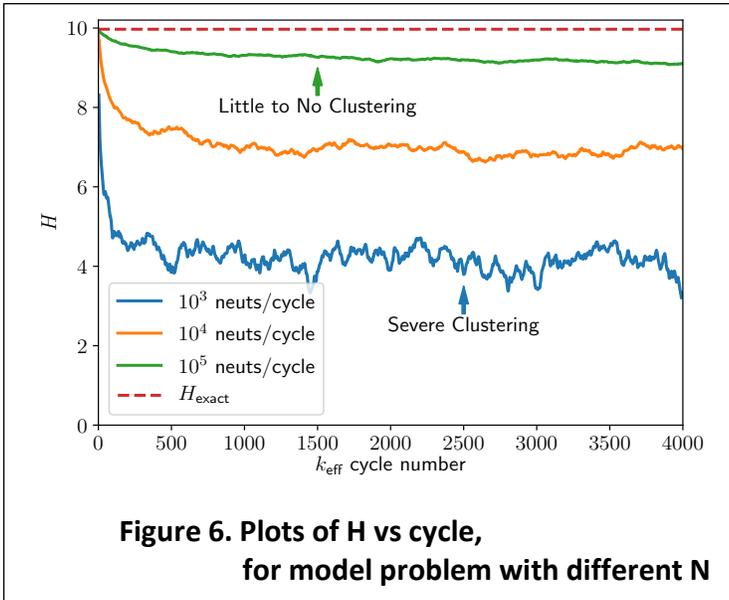
## 1.2 Background on Shannon Entropy

The Shannon entropy,  $H$ , of the fission neutron distribution was developed and introduced into MCNP to provide a metric for assessing convergence of the iteration process in criticality calculations [9,10].  $H$  is computed at the end of each iteration and plotted vs iteration; when  $H$  has converged, the fission distribution has converged and active cycles can begin. ( $k$ -effective converges in fewer iterations, as discussed in [10].)

To compute  $H$ , a grid covering all fissile regions is superimposed on the problem, and the number of source neutrons from fission are counted in each of the grid regions. After conversion to probabilities,  $p_i$ ,  $H$  is computed as

$$H = - \sum_{i=1}^B p_i \log_2 p_i$$

where  $B$  is the number of bins in the mesh ( $p_i \log_2 p_i = 0$  when  $p_i = 0$ ).  $H$  is a convex function, with a minimum value of 0 (for all neutrons in the same bin) and a maximum value of  $\log_2 B$  (for neutrons uniformly distributed among bins).



When the use of Shannon entropy for assessing convergence was first introduced, there was the assumption that the number of neutrons in a generation was much larger than the number of bins, so that the resultant  $p_i$  values accurately represented the fission neutron distribution and were not “too noisy.” A simple thought experiment illustrates the difficulties that arise when the number of neutrons is small: Consider a homogeneous fissionable material in a cube with reflecting boundaries, a Shannon entropy mesh with  $10 \times 10 \times 10$  bins, and 1,000 neutrons/generation. Since the exact theoretical neutron

distribution is known for this system (i.e., a uniform distribution), the true value of the Shannon entropy is known,  $H = \log_2 1000 = 9.97$ . If all neutrons are distributed uniformly, 1 per mesh bin, then  $H = 9.97$ , the correct maximum value. If there are 2 neutrons in every other bin, 500 clusters of 2, then  $H = 8.97$ . For 250 clusters of 4,  $H = 7.97$ ; for 125 clusters of 8,  $H = 6.97$ ; etc. For all of the neutrons in 1 bin, 1 cluster of 1000,  $H = 0$ . Clustering reduces the computed value of  $H$ . Further, when there are clusters of neutrons migrating through a problem, the computed values of  $H$  for successive cycles will vary in a much noisier manner as the clusters partially cross into different sets of bins in the mesh. This is illustrated in Figure 6, based on MCNP calculations with different numbers of neutrons per generation. As the number of neutrons per generation increases, the computed  $H$  approaches the correct value and exhibits less noise.

It should be clear from the thought experiment and the results shown in Figure 6 that, if an exact (or at least better) solution for the fission distribution is available, it can be used to compute a reference  $H$  that can be compared with the actual  $H$  for the neutron distribution in the calculation. Significant difference between  $H$  for the reference solution and  $H$  for the actual neutron distribution is an indicator of clustering.

## 2.0 Diagnostic Tests for Undersampling/clustering in Criticality Calculations

The discussion and examples in previous sections provide motivation for 2 types of statistical testing to diagnose the effects of undersampling. One type of test would examine the shape of the fission distribution compared to some reference solution. If the shapes of the 2 distributions do not match within some tolerance, then there are probably biases due to undersampling. The second type of test would compute some metrics for the neutron distribution and a reference distribution, such as Shannon entropy or relative entropy (i.e., the Kullback-Leibler discrepancy), and use the metrics for assessing undersampling effects. These tests could be used both for assessing convergence of  $k$ -effective and the fission distribution and for detecting undersampling. For convergence issues, the statistical tests would be performed during inactive cycles. For undersampling issues, the statistical tests would be performed only during active cycles (i.e., after convergence).

In References [11-13], a number of statistical tests were proposed based on information theory, involving the use of Shannon entropy, Kullback-Leibler discrepancy, and other schemes. At the time there was unfortunately no means for obtaining a reference solution prior to determining convergence. The proposed methods involved saving all tally results from all cycles, running a large number of cycles, and then doing posterior analysis to determine convergence and diagnose undersampling. The complexity of the methods raised issues (e.g., how many cycles must be run to collect information for the posterior analysis), and the computer resource requirements were prohibitive (e.g., for problems with large tally sets of 100s of GB or more, saving all the tallies for 100s or 1000s the cycles is clearly prohibitive for routine calculations). The reference solution was obtained by assuming that the last half of the cycles were converged (a questionable assumption for some problems), and that the average for the last half of the problem provided an accurate reference solution for use in statistical testing. In addition, such a reference solution would still exhibit undersampling biases.

Several years after the work described in [11-13], a new method for obtaining a reference solution was developed – the sparse-storage fission matrix method. The fission matrix method using sparse-storage is described in detail in references [14, 15] and an update on the current status is in preparation [16]. References [14, 15] demonstrated that the fission matrix method provides an accurate solution to the k-eigenvalue criticality problem if the mesh spacing is fine enough that the “flat-source” approximation is valid. The use of a sparse-storage scheme solved the longstanding difficulty of prohibitive memory requirements for the matrix, such that there are now no memory-size limits. Reference [16] provides details on the current status, including the physics-based method for assuring that the mesh is chosen appropriately, fine enough to provide accuracy, but not so fine that statistical tallies for the matrix elements introduce too much noise.

The present work is based largely on the ideas developed in [11-13], with a reference solution provided by the sparse-storage fission matrix method [14, 15, 16], and additional testing based on traditional statistical methods for comparing distributions. The tests proposed for diagnosing convergence of the iteration process will be discussed in a separate report. In the present work, the discussion below targets the diagnosis of undersampling in Monte Carlo criticality calculations.

In the sections below, 2 different tests are described for diagnosing the presence of undersampling/clustering problems. It is assumed in all that follows that convergence to the stationary state was reliably detected and that all cycles are active cycles.

- The first test involves comparing the average Shannon entropy for a block of cycles to the Shannon entropy for the cumulative fission source for the block. If these 2 estimates of Shannon entropy differ, then it is likely that undersampling/clustering is present. This test involves only information that is already present in MCNP or any other MC code (e.g., Tripoli, Scale, Serpent, MC21, etc.).
- The second test presumes that a reference solution for the fission source distribution is present, one that is not subject to source renormalization bias. For MCNP, the reference solution for the fission source distribution is provided by the sparse-storage fission matrix method. When an accurate reference solution for the fission source distribution is available,

biases in the source shape due to the renormalization process can be readily diagnosed. (While this test can be used with MCNP6, other MC codes may not have capabilities for producing the fission matrix, or may be memory-limited due to not having a sparse-storage scheme capable of providing a suitably accurate fission matrix.)

## 2.1 Undersampling Diagnostic Based on Jensen's Inequality and Shannon Entropy

Jensen's inequality [17] states that for any continuous convex function  $f(x)$ , the following is true:

$$f\left(\frac{1}{L} \cdot \sum_{k=1}^L a_k\right) \leq \frac{1}{L} \cdot \sum_{k=1}^L f(a_k)$$

For a block of  $L$  successive cycles in a criticality calculation, setting  $a_k = p_i^{(k)}$ , where  $(k)$  denotes the  $k^{\text{th}}$  cycle, and  $f(x) = x \log_2(x)$  (a convex function), this becomes

$$\left[ \frac{1}{L} \cdot \sum_{k=1}^L p_i^{(k)} \right] \log_2 \left[ \frac{1}{L} \cdot \sum_{k=1}^L p_i^{(k)} \right] \leq \frac{1}{L} \cdot \sum_{k=1}^L p_i^{(k)} \log_2 p_i^{(k)}$$

Summing  $i$  over all of the bins  $B$  in the entropy mesh and multiplying both sides by -1 gives

$$-\sum_{i=1}^B \frac{1}{L} \cdot \sum_{k=1}^L p_i^{(k)} \log_2 p_i^{(k)} \leq -\sum_{i=1}^B \left[ \frac{1}{L} \cdot \sum_{k=1}^L p_i^{(k)} \right] \log_2 \left[ \frac{1}{L} \cdot \sum_{k=1}^L p_i^{(k)} \right]$$

Then, interchanging the order of the summations on the left side,

$$\frac{1}{L} \sum_{k=1}^L \left[ -\sum_{i=1}^B p_i^{(k)} \log_2 p_i^{(k)} \right] \leq -\sum_{i=1}^B \left[ \frac{1}{L} \cdot \sum_{k=1}^L p_i^{(k)} \right] \log_2 \left[ \frac{1}{L} \cdot \sum_{k=1}^L p_i^{(k)} \right]$$

The quantity in brackets on the left is simply the Shannon entropy for the single cycle  $(k)$ , which will be denoted as  $H^{(k)}$ . The righthand side is the Shannon entropy of neutron distribution accumulated over the entire block of  $L$  cycles, denoted as  $H_L$ . That is, on the left side, the Shannon entropy is computed for each single cycle and then averaged over the block; on the right side, the neutron distribution is first averaged over the block of cycles, and then the Shannon entropy is computed once for the block. Thus,

$$\frac{1}{L} \sum_{k=1}^L H^{(k)} \leq H_L$$

This inequality reflects the fact that there is less information about the source distribution in any single cycle than in a block of  $L$  cycles. If undersampling is present, and especially if there is significant clustering, the  $H^{(k)}$  values for cycles in the block will vary. Clusters or just high/low regions will shift with iterations in the block, resulting in noisy and low values for the  $H^{(k)}$ .  $H_L$ , however, is computed at the end of the block, after accumulating the fission neutrons for all  $L$  cycles in the block. It will be less sensitive to clustering (i.e., clustering effects will average out as they move around). Note that if there are sufficient neutrons per cycle such that undersampling/clustering are not significant, all  $H^{(k)}$  will (approximately) equal  $H_L$ , and the equality will (approximately) hold.

Based on the above discussion, the following is proposed as a test for undersampling/clustering in MCNP criticality problems:

- After active cycles have begun, a block of  $L$  consecutive cycles is examined. Currently,  $L=10$  is recommended. Shannon entropy for the block is computed in 2 ways: First,  $H^{(k)}$  is determined for each cycle in the block and then averaged to determine  $\langle H^{(k)} \rangle$  for the block. Second, the fission neutron source distribution for each of the  $L$  cycles in the block is accumulated, normalized, and then used to determine  $H_L$  for the block.
- If  $\langle H^{(k)} \rangle$  differs significantly from  $H_L$ , then a warning is issued to alert the end-user that undersampling/clustering may be present and that more neutrons/cycle should be used for the problem.
- The tolerance band  $\varepsilon_1$  for  $(\langle H^{(k)} \rangle - H_L) / H_L$ , the relative difference in the 2 entropy estimates, is arbitrary for now. A value of  $\varepsilon_1=1\%$  for the tolerance limit on the relative difference in entropy has worked well in practical testing on a variety of problems.

This test for undersampling/clustering has been implemented and tested on a variety of problems in a local version of mcnp6 and has worked very well in testing. For problems where the warning is issued, problems were rerun with more neutrons/cycle. If the tests were still not passed, the neutrons/cycle was again increased and the problem was rerun, repeating this process until the tests were passed. For most problems tested, 50k – 100k neutrons/cycle was sufficient to pass the test. Rarely, more than 100k neutrons/cycle are required, typically for very large problems or loosely coupled problems.

Comment is in order concerning the values chosen for  $L$  and  $\varepsilon_1$ . In nearly all testing to date,  $L=10$  and  $\varepsilon_1=1\%$  have been used. If a larger value is chosen for  $L$ , then a smaller value for  $\varepsilon$  may be appropriate. Choosing too large a value for  $L$  may delay the testing, since it is performed only at the end of a block of  $L$  cycles, and some computer time may be wasted. Choosing too small a value for  $L$  risks unreliable testing results due to statistical noise in computing  $\langle H^{(k)} \rangle$ . Choosing  $\varepsilon_1$  too small risks unnecessary and unwarranted warnings, while too large results in possibly not detecting the presence of undersampling/clustering.

## 2.2 Undersampling Diagnostic Based on Shannon Entropy and the Fission Matrix Eigenfunction

For the present work, we assume that the fission matrix eigenfunction is a nearly “exact” representation of the fission neutron distribution. (Even if not exact, it provides a more accurate solution than the MC fission neutron distribution, since it is cumulative over all cycles, including inactive cycles.) What matters for diagnosing undersampling/clustering is the fact that the fission matrix eigenfunction does not involve the renormalization process that occurs at the end of each neutron generation. The fission matrix is accumulated over all cycles without renormalization (except the first cycle, which is used to obtain physics parameters for establishing sufficient meshing), while the fission neutron distribution is renormalized after each cycle. The renormalization process is the cause of the undersampling biases in k-effective and the fission neutron distribution.

Based on the above discussion, the following is proposed as a test for undersampling/clustering in MCNP criticality problems:

- After active cycles have begun, a block of  $L$  consecutive cycles is examined. Currently,  $L=10$  is recommended. During the block, the fission neutron source starting each cycle is accumulated, and at the end of the block used to compute Shannon entropy for the block,  $H_L$ . (This is the same quantity as  $H_L$  in the previous section.)
- The fission matrix equations are solved to determine the eigenvalue and fundamental-mode eigenfunction. Then the Shannon entropy of the fission matrix eigenfunction,  $H_{FM}$ , is computed.
- If the relative difference between  $H_L$  and  $H_{FM}$  is greater than some tolerance  $\varepsilon_2$ , then the shapes of the neutron distribution and fission matrix eigenfunction are different, indicating that renormalization biases are present for the neutron distribution due to undersampling. A warning is issued to alert the end-user that undersampling/clustering may be present and that more neutrons/cycle should be used for the problem.
- The tolerance band  $\varepsilon_2$  for  $(H_L - H_{FM}) / H_{FM}$ , the relative difference in the 2 entropy estimates, is arbitrary for now. A value of  $\varepsilon_2=1\%$  for the tolerance limit on the relative difference in entropy has worked well in practical testing on a variety of problems.

As for the diagnostic test described in Section 2.1, the values of  $L$  and  $\varepsilon_2$  are somewhat arbitrary, but values of  $L=10$  and  $\varepsilon_2=1\%$  have worked effectively in testing to date. The remarks in Section 2.1 apply as well to the tradeoffs in selecting  $L$  and  $\varepsilon_2$  – extreme values may result in misdiagnosis.

### 3.0 Conclusions and Further Work

Section 2 has described the statistical tests performed during active cycles of an MCNP criticality problem, as implemented in a local version of MCNP6.

These diagnostic tests have been applied to an assortment of criticality problems, including: a 2D PWR model, the ATR (advanced test reactor), the 3D Kord Smith Challenge problem (OECD-NEA 3D reactor computer-performance benchmark), the August Winkelman research reactor, the ICSBEP benchmark case LEU-COMP-THERM-078 (a Sandia experiment), a large 3D storage pool with checkerboard arrangement (OECD-NEA EG on source convergence benchmark), a 400 cm tall single reactor fuel-pin unit cell with reflecting boundary conditions, and the Whitesides problem (k-effective of the world). In all cases tested the diagnostic tests for undersampling/clustering were effective. Following the warning advice, increasing the number of neutrons/cycle to larger values resulted in passing the tests (no warning issued).

Results from preliminary testing are very encouraging, and further work is underway to provide a set of test problems, results for various neutrons/cycle, and the effectiveness of the statistical tests for undersampling.

It should be noted that the manual approach to undersampling – increase the neutrons/cycle and rerun, repeat as needed – could be automated. That will likely happen after more experience is gained with the undersampling tests. For now, just diagnosing the undersampling is a significant advance.

Obviously further work is needed, including:

- Apply the tests to many more criticality problems. Many more ordinary cases as well as odd-ball cases need to be tested.
- For cases that indicate the presence of undersampling, obtain plots of the differences between the fission matrix eigenfunction and the MC neutron distribution, to assess whether the tests are too stringent or not stringent enough.
- To confirm the validity of the fission matrix eigenfunction used in the tests, compare it to problems run with very large numbers of neutrons/cycle. (This is strictly validation of the fission matrix methodology, a separate task from diagnosing undersampling. But the 2 methods overlap.)
- Vary the parameters  $L$ ,  $\varepsilon_1$ , and  $\varepsilon_2$  and assess the sensitivity and reliability of the tests for a variety of cases. This could involve, for example, running high-precision MC calculations with millions of neutrons/cycle to provide a brute-force reference solution, and then comparing to results from many independent replica calculations with fewer neutrons/cycle.
- While bias in the k-effective eigenvalue due to undersampling can be quantified, the theory associated with biases in the eigenfunction shape is extremely complex. It is possible, but not likely, that extremely thorough and deep exploration of the Brissenden and Gelbard theory papers may yet provide a means for characterizing the biases in eigenfunction shape. If so, there may be some theoretical basis for selecting the parameters  $L$ ,  $\varepsilon_1$ , and  $\varepsilon_2$ , or for additional diagnostic procedures.

## Acknowledgements

This work was supported by the US-DOE-NNSA Nuclear Criticality Safety Program and by the US-DOE-NNSA Advanced Simulation & Computing Program.

## References

- 1 E.M. Gelbard and R.E. Prael, "Monte carlo Work at Argonne National Laboratory", in *Proc. NEACRP Meeting of a Monte Carlo Study Group*, ANL-75-2, Argonne National Laboratory, Argonne, IL (1974).
- 2 R.J. Brissenden and A.R. Garlick, "Biases in the Estimation of Keff and Its Error by Monte Carlo Methods," *Ann. Nucl. Energy*, Vol 13, No. 2, 63-83 (1986).
- 3 F.B. Brown, "A Review of Best Practices for Monte Carlo Criticality Calculations", ANS NCSD 2009, Hanford WA, LA-UR-09-03136 (2009)
- 4 E. Dumonteil, et al., "Particle Clustering in Monte Carlo Criticality Simulations", *Ann. Nucl. Energy*, Vol 63, 612-618 (2014)
- 5 A. Zoia, E. Dumonteil, "Neutron Clustering: Spatial Fluctuations in Multiplying Systems at the Critical Point", M&C 2017, Jeju, Korea, April 16-20 (2017)

- 6 F.B. Brown, "Investigation of Clustering in MCNP6 Monte Carlo Criticality Calculations", Int. Conf. on Transport Theory, Monterey CA, Oct 2017, LA-UR-17-29261 (2017).
- 7 C.J. Werner, et al., "MCNP6.2 Release Notes", Los Alamos National Laboratory, report LA-UR-18-20808 (2018)
- 8 T.M. Sutton, A. Mittal, "Neutron Clustering in Monte Carlo Iterated-Source Calculations", M&C 2017, Jeju, Korea, April 16-20 (2017)
- 9 T. Ueki & F.B. Brown, "Stationarity and Source Convergence in Monte Carlo Criticality Calculations", ANS Topical Meeting on Mathematics & Computation, Gatlinburg, TN April 6-11, 2003, LA-UR-02-6228 (September, 2002).
- 10 F.B. Brown, "On the Use of Shannon Entropy of the Fission Distribution for Assessing Convergence of Monte Carlo Criticality Calculations", PHYSOR-2006, Vancouver, LA-UR-06-3737 (2006).
- 11 T. Ueki, F.B. Brown, "Stationarity Modeling and Information-Based Diagnostics in Monte Carlo Criticality Calculations", *Nucl. Sci. Eng.*, **149**, 38-50 (2005).
- 12 T. Ueki, "Information Theory and Undersampling Diagnostics for Monte Carlo Simulation of Nuclear Criticality", *Nucl. Sci. Eng.*, **151**, 283-292 (2005).
- 13 T. Ueki, "On-the-Fly Diagnostics of Particle Population in Iterated-Source Monte Carlo Methods", *Nucl. Sci. Eng.*, **158**, 15-27 (2008).
- 14 F.B. Brown, S.E. Carney, B.C. Kiedrowski, W.R. Martin, "Fission Matrix Capability for MCNP, Part I - Theory", Mathematics & Computation 2013, Sun Valley, ID, LA-UR-13-20429 (2013)
- 15 S.E. Carney, F.B. Brown, B.C. Kiedrowski, W.R. Martin, "Fission Matrix Capability for MCNP, Part II - Applications", Mathematics & Computation 2013, Sun Valley, ID, report LA-UR-13-20454 (2013)
- 16 F.B. Brown, "Current Status of the Fission Matrix Method in MCNP6", in preparation (2018)
- 17 T.M. Cover & J.A. Thomas, Elements of Information Theory, 2<sup>nd</sup> Edition, Wiley (2006)