LA-UR-11-00047

Approved for public release; distribution is unlimited.

Title: Statistical Coverage Concerns in a Revised 'k-Effective of the World' Problem

Author(s): Brian C. Kiedrowski
Forrest B. Brown

Intended for: 2011 ANS Annual Meeting
June 26-30, 2011
Hollywood, FL

Los Alamos
NATIONAL LABORATORY
EST. 1943

Statistical Coverage Concerns in a Revised 'k-Effective of the World' Problem

Brian C. Kiedrowski and Forrest B. Brown

Los Alamos National Laboratory, P.O. Box 1663 MS A143, Los Alamos, NM 87545,
bckiedro@lanl.gov, fbrown@lanl.gov

INTRODUCTION

The k-Effective of the World problem proposed by Whitesides [1] in 1971 is classic in its demonstration of bias and undersampling issues in Monte Carlo eigenvalue calculations. The problem is a 9 x 9 x 9 loosely coupled array of bare plutonium spheres with the central sphere being exactly critical by itself and the others being subcritical. Using the recommended best practices of the day, (run 300 particles per batch, skip three, and run one hundred), it was impossible to get the correct result for k.

Recent work [2] shows, however, that this problem was only difficult because of computational limitations at the time, and that, obtaining a correct solution for k is easy if the current best practices [3] are followed. These best practices are to run at least 5,000 neutrons per batch and check convergence in k and other diagnostics such as the Shannon entropy before starting active batches. Further, the best practices recommend that for loosely coupled problems such as this, source neutrons be started in each individual region to ensure adequate sampling.

The current best practices were developed to address the issues of source convergence and statistical bias from stochastic renormalization. Implicit in these and the theory behind them is the assumption that the entire problem is appropriately sampled or covered. Failure to ensure this is true leads to an additional concern distinct from source convergence and statistical bias.

A revised version of the k-Effective of the World problem is created such that the current best practices can fail to accurately predict k with high likelihood because of failing to adequately cover the problem if the practitioner is not careful. The specification for the revised problem is given along with reasons why statistical coverage is an important concern. Based upon this, recommendations of additional considerations to the current best practices toward finding the correct k are detailed.

THE REVISED PROBLEM

Specifications

The k-Effective of the World by Whitesides is defined as follows: a 9 x 9 x 9 array of plutonium-239 spheres is spaced evenly in vacuum with center-to-center distances of 60 cm. The array is surrounded by a thick water reflector. The central sphere is then replaced by a plutonium-239 sphere such that it, in its bare configuration with no surrounding spheres, is exactly critical. The resulting system must, from physical considerations, be supercritical.

The Revised problem is defined to use a similar 9 x 9 x 9 array of spheres (density of 20 g/cm^3), except that their radii are 3.1 cm. The region between the spheres is flooded with water (1.0 g/cm^3) and the center-to-center spacing is reduced to 20 cm. k (all calculations use MCNP5 1.60 [4] and ENDF/B-VII.0 nuclear data [5]) for this configuration is 0.93842(8), below the typical 0.95 safety limit. The central sphere is replaced with a sphere of radius 4.1 cm with a natural cadmium shell (8.65 g/cm^3) with a 0.5 cm thickness. This sphere in an infinite water reflector has k of 0.95454(4), above the typical 0.95 safety limit.

When performing this calculation with 10,000 neutrons per cycle using a source guess that starts particles uniformly in each of the 729 spheres with 1,000 inactive cycles and 500 active cycles, there is about a 60 percent chance (as determined by 100 independent MCNP calculations) of getting an incorrect result of k (the result matches that of the uniform sphere array case) whereas the correct result for k closely matches the infinitely reflected central sphere case. Visual inspection of plots of k and the Shannon entropy (the MCNP convergence test passes somewhere between 400 and 900 cycles) show ordinary behavior for both the correct and incorrect cases.

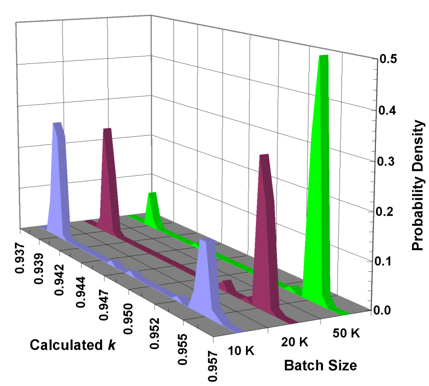

Fig. 1. Distribution of MCNP calculated k for various batch sizes.

Reasons for Problematic Behavior

The reason there is such a high probability of false convergence in k is because of failing to statistically cover the central sphere. The distribution of results for batch sizes of 10, 20, and 50 thousand is given in Fig. 1 and it appears to be a bimodal distribution with two normally distributed peaks. Inspection of the population tables in MCNP shows the sampling of the central region to be much greater for cases where k has the correct value suggesting that the issue is failing to adequately sample the central region.

In a loosely coupled problem where different regions are more and less reactive, the physics is such that the value of k is most determined by the region that is most reactive, with the coupling of the less reactive regions providing what are typically small additions to the eigenvalue. It is therefore most important in the Monte Carlo simulation to adequately sample the most reactive region. The importance of sampling the other regions depends upon how reactive those regions are relative to the most reactive and the degree of coupling between them.

In this problem, the central sphere is most reactive. Furthermore, the coupling is such that it is asymmetric: the water between the spheres and cadmium layer on the central sphere ensures that it is far easier for the central sphere to "communicate" with the ones surrounding it than it is for those to communicate with it. In the Monte Carlo simulation, this means a few things. First, if a fission chain in the central sphere is not established from the initial guess, it is difficult to establish one later because communication is weak toward the central sphere. Second, the total population of neutrons within the batch is preserved on average. Since neutrons are far more likely to leak from the central sphere and create fission in another region than the reverse occurring, the population has a tendency to migrate from the central sphere to the others should the population be too small.

Whether or not the Monte Carlo simulation will appropriately sample the central region and therefore have a high probability of getting the correct k is a function of two probabilities. The first is the chance of neutrons entering the central sphere, whether it be from the source guess or from entry from another region. In terms of the former, a uniform source guess has a probability of sampling of 1/729; therefore for a batch size of 10,000 neutrons, about fourteen neutrons, on average, will start in the central region. The chance of a neutron entering from another region is significantly smaller because the number of mean-free-paths between it and the nearest sphere is large.

The next consideration is the probability of establishing a persistent chain. A neutron within the central region may not produce progeny from fission because there is a finite probability of leakage or capture. Even with the production of progeny, those are not guaranteed to be persistent either. Only after a sufficient amount of multiplication can a chain have a great enough population such the probability of termination of that entire chain is small. Therefore, the number of neutrons must be sufficiently large such that there is a high probability of establishing at least one chain that meets this criteria.

GUIDANCE FOR PRACTITIONERS

To achieve the correct solution, the central sphere must be adequately sampled. A rigorous determination of "adequate" and how it might be determined with, at very least, a fool-proof check is a research area; the aim here is to provide practical guidance with existing capabilities.

Foremost is the statement that no software can be treated as a "black box" even if the best practices are followed. Further attention must be given to the problem and the numerics itself.

Unfortunately, there is no indication from watching convergence of k or even the Shannon entropy as implemented in MCNP; the only cases where a problem might be detected are the rare event where a persistent chain in the central sphere emerges late in the calculation. In other words, the diagnostics fail for any reasonable calculation time – eventually, a neutron will enter the central region and establish a persistent chain from chance alone; however, this appears to be a seldom occurrence.

There are three suggestions for practitioners solving loosely coupled problems in addition to the current best practices:

The first is to identify the most reactive region so to determine where to focus attention. A method of doing this is to run each region independently; a way to capture effects of other regions other than fission is to turn off fission (treat as capture) in cells outside the region being examined. Doing this for 729 distinct regions can be very onerous; however, some physical intuition is useful here in determining that there are, in reality, only two distinct regions: the array minus the central sphere, and the central sphere itself. In all cases, k of the composite system should be at least that of the most reactive region.

In MCNP, the most reliable method for determining whether the central sphere has been sampled is to examine the population tables. In all cases where k is incorrect, there is an unusually low number of neutrons (a few hundred collisions over millions of active source neutrons) in the central sphere relative to all others. Physical intuition suggests that the central sphere should have a significant fraction of neutrons. MCNP provides warnings if no source points are sampled in a fissionable cell, and these are useful; however, (for the 10,000 batch size case) this warning is issued only about two-thirds of the incorrect trials meaning that the other one-third of the cases provided absolutely no warning of an incorrect k unless the practitioner uses physical intuition and inspects the population table. Additionally, MCNP provides mesh tallies with the plotter to assist the practitioner in visualizing the coverage of the problem.

Second, some intuition of the coupling between regions is useful. Problems where there is an asymmetry in coupling that is unfavorable to the most reactive region are particularly prone to issues of statistical coverage. This can arise from spectral considerations as seen in this problem because of the cadmium shell, or from geometry if the regions are of very different sizes. In such a case, particular care needs to be given to determining sampling.

The third recommendation is to pick a source guess that will adequately sample the most reactive region. While it is impossible to say exactly what this number is with the current tools available, a good suggestion is for there to be at least several hundred within that region alone. Note that, for this problem, simply sampling a point source in the center with a batch size of 5,000 neutrons always produces the correct result (one hundred random trails used for this calculation). It is still good practice, however, to sample the other regions as well ensuring at least a few dozen start within those regions in case coupling effects are more important than initially suspected.

CONCLUSIONS & CHALLENGES

A Revised k-Effective of the World problem is presented that poses challenges for both criticality safety practitioners and methods developers. Even following the current best practices, there is still an unacceptably high probability of getting the wrong k unless great care is taken.

Suggestions have been made to practitioners with regards to checks that can be performed using existing capabilities; however, it is the onus of methods developers to ensure there are appropriate diagnostics to warn practitioners if regions of loosely coupled problems lack statistical coverage. Some work [6] has been performed in the area of undersampling, and this Revised k-Effective of the World problem can serve as a test case for the effectiveness of such methods toward statistical coverage. Additional test problems that are similarly prone should be developed, if for nothing else, to serve as part of a suite to stress methods related to statistical sampling and coverage.

ACKNOWLEDGMENTS

Funding for this work was provided by the US Department of Energy Nuclear Criticality Safety program.

REFERENCES

1. G. E. WHITESIDES, "A Difficulty in Computing the k-effective of the World," *Trans. Am. Nucl. Soc.*, **14**, No. 2, 680 (1971).
2. F. B. BROWN, "Revisiting the 'K-effective of the World' Problem", *Trans. Am. Nucl. Soc.*, **102**, No. 1, 302-304 (2010).
3. F. B. BROWN, "A Review of Best Practices for Monte Carlo Criticality Calculations", *Proc. ANS NCSD-2009*, Richland, WA, Sept 13-17 (2009).
4. X-5 MONTE CARLO TEAM, "MCNP – A General Monte Carlo N-Particle Transport Code, Volume I: Overview and Theory," LA-UR-03-1987, Los Alamos National Laboratory (2003).
5. M. B. CHADWICK, et. al. "ENDF/B-VII.0: Next Generation Evaluated Nuclear Data Library for Nuclear Science and Technology," *Nucl. Data Sheets*, **107** 2931-3060 (2006).
6. T. UEKI, "Information Theory and Undersampling Diagnostics for Monte Carlo Simulation of Nuclear Criticality," *Nucl. Sci. & Engr.*, **151** 283-292 (2005).