Annals of Nuclear Energy 94 (2016) 16–21

Contents lists available at ScienceDirect

Annals of Nuclear Energy

journal homepage: www.elsevier.com/locate/anucene

Monte Carlo criticality calculations accelerated by a growing neutron population

Jan Dufek [a,*], Kaur Tuttelberg [b]

[a] KTH Royal Institute of Technology, Division of Nuclear Reactor Technology, AlbaNova University Center, 10691 Stockholm, Sweden
[b] Tallinn University of Technology, Department of Electrical Power Engineering, Tallinn, Estonia

ARTICLE INFO

Article history:
Received 10 November 2015
Received in revised form 17 February 2016
Accepted 18 February 2016

Keywords:
Monte Carlo criticality
Fission source
Convergence
Bias
Convergence acceleration

ABSTRACT

We propose a fission source convergence acceleration method for Monte Carlo criticality simulation. As the efficiency of Monte Carlo criticality simulations is sensitive to the selected neutron population size, the method attempts to achieve the acceleration via on-the-fly control of the neutron population size. The neutron population size is gradually increased over successive criticality cycles so that the fission source bias amounts to a specific fraction of the total error in the cumulative fission source. An optimal setting then gives a reasonably small neutron population size, allowing for an efficient source iteration; at the same time the neutron population size is chosen large enough to ensure a sufficiently small source bias, such that does not limit accuracy of the simulation.

1. Introduction

In order to obtain reliable results, the fission source in Monte Carlo criticality calculations must be sampled from its steady-state distribution. As the steady-state fission source is generally not known, a number of inactive cycles are simulated first, during which the fission source is expected to converge from its initial (usually guessed) state to the steady state. Since no results are collected during the inactive cycles, the computing time of the inactive cycles can be considered as lost. Simulation results are combined only over a number of active cycles during which the fission source is assumed to be converged. The efficiency of Monte Carlo criticality calculations thus may be improved by methods that try to reduce the computing cost of the inactive cycles (see the works by Yamamoto and Miyoshi, 2004; Kadotani et al., 1991; Kitada and Takeda, 2001; Dufek and Gudowski, 2009; Brissenden and Garlick, 1986).

Tuttelberg and Dufek (2015) demonstrated that the efficiency of Monte Carlo criticality calculations is sensitive to the selected neutron population size–the number of neutron histories simulated at each cycle (generation), and they suggested an optimisation procedure for choosing the optimal population size. This optimisation was done with respect to the computational time allocated beforehand the simulation, and the population size was set at a fixed value common to all cycles–as is a standard practice in Monte Carlo criticality calculations.

In this paper, we demonstrate that the computing efficiency may be improved furthermore when the neutron population size, m, gradually increases over the successive cycles. The strategy of increasing the neutron population size during the criticality simulation is not new; e.g., Gast and Candelore (1974) tried a linear growth of the neutron population size at the rate of 10 neutron histories added at each cycle. The strategy has a clear benefit: it allows to cut the computing cost on cycles where the iterated fission source contains large errors inherited from the initial cycle source. The bias in fission source of the order $O(1/m)$ (Brissenden and Garlick, 1986) does not need to be kept very small when the fission source contains large errors since these are the limiting factor for achieving a good accuracy; therefore, the neutron population size (and so the computing cost per cycle) can be set relatively small at the beginning of the simulation. Nevertheless, as the errors in the fission source decay over the successive cycles, the bias in the fission source may become the dominant error unless it is reduced by increasing the neutron population size. In this paper we propose an on-the-fly control methodology for the neutron population size according to the actual convergence of the Monte Carlo fission source.

The paper is organised as follows. Section 2 describes the suggested method. Section 3 gives results of the numerical test calculations, and Section 4 summarises our conclusions.

* Corresponding author.
E-mail address: jandufek@kth.se (J. Dufek).

http://dx.doi.org/10.1016/j.anucene.2016.02.015
0306-4549/

2. Method

We would like to note that no distinction is made between inactive and active cycles in the following text as the suggested method is supposed to be applied to all inactive as well as active cycles. The neutron population size is controlled during the whole simulation. This does not eliminate the need for inactive cycles in simulations that use the method.

The fission source contains errors of statistical and systematic origins. The statistical error of the order $O(1/\sqrt{m})$ is introduced directly into the fission source by sampling the source at a limited number, m, of fission sites. The statistical error propagates into the results; however, it is reduced to $O(1/\sqrt{n \times m})$ by combining the results over n cycles.

Unlike the statistical error, the systematic errors are present in Monte Carlo calculations due to peculiarities of the specific computing procedures involved. In case of criticality Monte Carlo calculations, a systematic bias of the order $O(1/m)$ is present in the fission source (hence, in the results). Brissenden and Garlick (1986) explained and demonstrated that the source bias is caused by the normalisation of the fission source size to the required value m at each cycle. This bias does not decay over the cycles; hence, it may become the dominant error in results of large simulations where the statistical and other errors have decayed. To prevent the source bias from limiting the accuracy of results, it is often recommended to use a very large m in Monte Carlo criticality calculations.

Another source of errors is caused by a cycle-to-cycle propagation of errors originating from the guessed fission source sampled at the very first cycle; the decay of this error over the cycles is governed by the dominance ratio of the system (Ueki et al., 2003).

Tuttelberg and Dufek (2015) showed that the efficiency of the criticality calculation is affected by the selected neutron population size m. There are few factors that decide whether the calculation can benefit from setting a small or large value of m. While a large m ensures a small bias in the fission source, it restricts the number of cycles that can be simulated within a certain computing time; a small number of inactive cycles may be then insufficient for decaying the error originating from the initial fission source. On the other hand, choosing a small m increases the systematic bias in the fission source, which corrupts the results as well. Choosing the suitable value of m thus represents an optimisation problem: m should be small enough to ensure an acceptably small computing time per cycle, yet m should be large enough to limit the bias in the fission source.

Tuttelberg and Dufek (2015) derived a simplified formula for optimising the value of m, taking into account the computing time allocated for the whole simulation, the estimated dominance ratio of the system and the estimated error in the initial fission source sampled at the beginning of the simulation. The final formula was derived with an assumption that m remains fixed at all cycles, as it is common in standard Monte Carlo criticality calculations.

In this paper, we abandon the practice of keeping the neutron population size fixed over all cycles; instead, we suggest to gradually increase the neutron population size over the successive cycles. This approach has a direct impact on the efficiency of the criticality calculation. A small population size allows iterating the fission source rapidly during the initial cycles (in terms of the wall-clock time), which helps to decrease the error coming from the initial fission source. The fairly large bias associated with a small population size has no significance during these cycles since the total error in the fission source is dominated by the error coming from the source sampled at the initial cycle. As the error originating from the initial fission source decays over the successive cycles, we suggest to gradually increase the neutron population size in order to balance the source bias with other decaying errors. When the neutron population grows over the successive cycles, all sources of errors decay, which makes it possible to achieve any required accuracy.

In other words, we suggest that at any cycle the neutron population size is set as small as possible–yet without compromising the accuracy of the results by the fission source bias. This can be achieved when the source bias is, in a way, balanced to other errors in the fission source, so that neither the bias nor other types of errors dominate in the final results. Here, we must consider that the final results are tallied over a number of cycles; hence, the impact of the statistical error (that is introduced in the fission source at each cycle) on the final results is reduced by averaging over the cycles. It would, therefore, be unreasonable to balance the source bias to the statistical error in the fission source of a single cycle–unlike the statistical error, the source bias is not reduced by averaging over a number of cycles. Consequently it appears reasonable to assume that the bias should be balanced to other errors in the cumulative fission source, i.e., in the fission source combined over all cycles (and not in the fission source of a single cycle).

Therefore, in order to decide the maximal acceptable source bias (hence the minimal size of the neutron population) we suggest to estimate the total error in the cumulative fission source, and decide a certain ratio of the source bias to the total error. The neutron population size can then be controlled in such a way that this ratio is maintained at the intended level during the whole simulation.

The relative error in the cumulative fission source, ε, can be estimated using the fundamental-mode eigenvector of the fission matrix. Tuttelberg and Dufek (2014) suggest to estimate the relative error ε in the cumulative fission source $\mathbf{s}^{(n)}$ in the nth criticality cycle by $\hat{\varepsilon}$ as

$$\hat{\varepsilon} = \left\| \tilde{\mathbf{s}}^{(n)} - \tilde{\mathbf{h}}^{(n)} \right\|_1, \tag{1}$$

where $\mathbf{h}^{(n)}$ is the fundamental-mode eigenvector of the fission matrix $\mathbf{H}^{(n)}$ that is tallied over all cycles of the Monte Carlo calculation (Carter and McCormick, 1969), and $\sim$ denotes a normalisation operator defined for any vector $\mathbf{x}$ as

$$\tilde{\mathbf{x}} = \frac{\mathbf{x}}{\|\mathbf{x}\|_1}.$$

Note that $\hat{\varepsilon}$ defined by Eq. (1) may take any value from the interval $[0, 2]$, which is not convenient. Here, we therefore propose to scale the maximal value of $\hat{\varepsilon}$ to unity, which changes Eq. (1) into

$$\hat{\varepsilon} = \left\| \tilde{\mathbf{s}}^{(n)} - \tilde{\mathbf{h}}^{(n)} \right\|_1 \Big/ 2. \tag{2}$$

The fission matrix $\mathbf{H}$ can be computed during a standard Monte Carlo criticality simulation using a spatial mesh superimposed over the whole system. The (i, j)th element of $\mathbf{H}$ represents the probability that a fission neutron born in space zone j causes a subsequent birth of a fission neutron in space zone i. The fission matrix is always combined over all simulated cycles, so statistical errors in the fission matrix decay inversely to the square root of the total number of all neutron histories simulated during all cycles. The fundamental-mode eigenvalue of $\mathbf{H}$ equals the multiplication factor k_{eff}, and the corresponding eigenvector $\mathbf{h}$ equals the discretised fundamental-mode fission source.

The calculation of the fundamental-mode eigenvector of the fission matrix comes at an additional computing cost that could worsen the computing efficiency if it is performed at every cycle. Therefore, we suggest to update the neutron population size in selected cycles only, and to carry out a certain number, c, of cycles