M&C 2017 - International Conference on Mathematics & Computational Methods Applied to Nuclear Science & Engineering, Jeju, Korea, April 16-20, 2017, on USB (2017)

Accelerating Source Convergence in Monte Carlo Criticality Calculations Using a Particle Ramp-up Technique

Amanda L. Lund, Paul K. Romano, Andrew R. Siegel

Argonne National Laboratory, Mathematics and Computer Science Division, 9700 S. Cass Ave., Lemont, IL 60439
alund@anl.gov, promano@anl.gov, siegela@mcs.anl.gov

Abstract - *We propose a procedure for accelerating convergence of the source distribution in Monte Carlo criticality calculations. The number of particles simulated during inactive generations can be significantly reduced by first converging the source using fewer particles per generation and then increasing the number of particles as the error from the initial guess decreases. We introduce a simple, automated way to adjust generation size and demonstrate its effectiveness at reducing time spent in inactive generations using a modified version of the OpenMC code. Tests on several difficult source convergence benchmark problems show a* $5-40\times$ *speedup in time taken to reach a converged source distribution.*

I. INTRODUCTION

Despite continued advances in processor performance, using the Monte Carlo (MC) method for the routine design and analysis of large nuclear reactors is still not feasible on present-day systems. Reaching the desired level of accuracy for a large, complex problem requires impractically long execution time. Techniques to reduce the execution time without sacrificing accuracy are therefore needed to make such design calculations viable.

MC neutron transport codes solve *k*-eigenvalue problems using a procedure analogous to solving a matrix eigenvalue problem by power iteration. Simulations begin with a guess of the fission neutron source distribution. Within an iteration (referred to here as a *generation*[1]), a finite number of source neutrons are tracked. Any fission neutrons produced are stored and subsequently used as the starting source for the next generation. Useful results cannot be collected until the spatial distribution of source neutrons has converged to the true fundamental eigenmode; otherwise, the solution will be biased by the initial guess. The initial generations used to converge the source—often called the *inactive* generations—are discarded. Iterations then continue so that results can be accumulated until the variance in the results has been sufficiently reduced. The rate of convergence of the source distribution depends on the ratio of higher harmonic eigenvalues to the fundamental mode eigenvalue. Looking at the ratio of the first higher harmonic eigenvalue to the fundamental eigenvalue, the *dominance ratio*, can give a simple measure of how slow or fast the source should be expected to converge; problems with a higher dominance ratio generally require more generations to converge.

The total number of particles that are discarded during the inactive generations is a product of the number of inactive generations and the number of particles per generation. The number of inactive generations necessary to converge the source distribution does not change with respect to the number of particles per generation used, and, unfortunately, practical applications typically require many particles per generation to avoid an undersampling bias [1, 2]. Thus, for high dominance ratio problems such as large light-water reactors, considerable computational time may be wasted constructing an accurate source distribution during inactive generations. Any method that can reach a converged source distribution while minimizing the number of particles in inactive generations would reduce the overall execution time. Consequently, many attempts have been made in the literature to "accelerate" source convergence, that is, to reduce the number of inactive generations required.

Much of the previous research on accelerating source convergence in MC criticality calculations draws on techniques used in deterministic methods. Yamamoto and Miyoshi [3] applied Wielandt's method to track a portion of the neutrons produced from fission within the current generation rather than banking them. This leads to reliable convergence in fewer iterations, although the computation time is not reduced since more particles are tracked per generation. Kitada and Takeda [4] effectively accelerated convergence of the source in simple 2D problems by using the fundamental mode eigenvector computed from the fission matrix to adjust the fission source distribution. Lee et al. [5] employed the coarse mesh finite difference method to update the fission source distribution and substantially reduce the number of inactive generations. She et al. [6] modified both Wielandt's method and the superhistory method [1] by adjusting the source population in inactive generations and demonstrated a significantly faster convergence rate. Carney et al. [7] further elaborated on the theory and practical work on the fission matrix. They implemented a sparse storage scheme for the fission matrix that permits a more refined mesh and accurate representation, and using it they demonstrated a substantially reduced number of generations to convergence. Pan et al. [8] presented a modified fission matrix method that reduces the instability of the original method caused by magnification of statistical error.

While the aforementioned methods accelerate convergence of the source distribution by making some modification to the source, faster convergence can be attained by simply relaxing the requirement that the number of particles per generation stay constant during the inactive generations. More precisely, one can coarsely converge the source using fewer particles per generation and then increase the number of particles as the error from the initial guess decreases. This approach can significantly reduce the total number of particles simulated during inactive generations while still reaching a valid estimate of the true source distribution. To be clear, this is not a novel

[1] The term *cycle* is also prevalent in the literature.

idea; in fact, a version of it was first proposed in 1974 by Gast and Candelore [9], albeit in a different context. They describe increasing the generation size linearly in order to avoid a bias that normally results from the use of a fixed generation size. The MCNP5 user's manual [10] also recommends manually increasing the number of particles per generation for slowly converging problems, but MCNP itself does not provide an automated way of doing this. A recent paper by Dufek and Tuttelberg [11] presents a method to automatically adjust the number of particles per generation based on an estimate of the error in the cumulative fission source obtained via comparison to the fundamental mode eigenvector of the fission matrix. While the idea is compelling, it is a considerable departure from the normal power iteration algorithm in that the adjustment is to be done during both inactive and active generations. Additionally, solving for the fundamental eigenvector of the fission matrix incurs a potentially non-trivial computational cost.

Despite the clear benefits of the *ramp-up* approach—starting with a small generation size and gradually increasing it—to our knowledge, no production MC code provides an automated way of adjusting the generation size. We propose and demonstrate a procedure for increasing the generation size during inactive generations that is shown to be both more efficient than the linear increase proposed by Gast and Candelore and considerably simpler than the automated adjustment proposed by Dufek and Tuttelberg. Because the procedure only requires a small change in the way that inactive generations are carried out, it is our hope that it will allow production MC code developers to adopt a simple, effective method for accelerating source convergence with minimal changes in both source code and user input.

II. METHOD

MC criticality calculations can be considered an implementation of the power method using MC techniques to solve the steady-state neutron transport equation. The equation governing the fission source for a k-eigenvalue problem can be expressed as

$$s = \frac{1}{k} H s, \tag{1}$$

where the fission source $s = s(\mathbf{r})$ is the concentration of neutrons at $\mathbf{r}$ and the integral fission operator $H = H(\mathbf{r}' \to \mathbf{r})$ represents the expected number of fission neutrons produced in the volume around $\mathbf{r}$ from a fission neutron born at $\mathbf{r}'$. This is an eigenvalue problem that can be solved by a MC calculation in a manner analogous to power iteration given initial guesses $s^{(0)}$ and $k^{(0)}$:

$$s^{(n+1)} = \frac{1}{k^{(n)}} H s^{(n)} + \epsilon^{(n+1)}, \tag{2}$$

where $\epsilon^{(n+1)}$ is the stochastic error term resulting from tracking a finite number of particles in each iteration. The stochastic error is $O(M^{-1/2})$, where M is the number of particles per generation.

The k-eigenvalue equation has distinct eigenvectors:

$$s_j = \frac{1}{k_j} H s_j, \quad k_0 > |k_1| \geq |k_2| \geq \dots, \tag{3}$$

where k_0 is the effective multiplication factor. The initial guess of the fission source can be expressed as a linear combination of the solution eigenvectors,

$$s^{(0)} = \sum_{j=0}^{\infty} a_j s_j, \tag{4}$$

and substituted back into Equation 2 to obtain, after some arithmetic,

$$\begin{aligned} s^{(n)} = a_0 k_0^n \prod_{i=0}^{n-1} \frac{1}{k^{(i)}} \cdot \Bigg[s_0 + \left(\frac{a_1}{a_0}\right)\left(\frac{k_1}{k_0}\right)^n s_1 \\ + \left(\frac{a_2}{a_0}\right)\left(\frac{k_2}{k_0}\right)^n s_2 + \dots \Bigg] + \epsilon^{(n)} \end{aligned} \tag{5}$$

Equation 5 illustrates how the error originating from the initial guess of the fission source diminishes with increasing number of generations. Since the first term is a constant and $1 > |k_1/k_0| > |k_2/k_0| > \dots$, the higher-order eigenmodes die off as $n \to \infty$, and the fission source converges to the fundamental mode s_0. After n iterations, the component of the error from the ith mode is reduced by $(k_i/k_0)^n$. The lower-order terms persist the longest, and the rate of convergence is governed by the ratio of the first two eigenvalues k_1/k_0, known as the dominance ratio ρ.

It is wasteful to simulate many particles per generation and reduce the stochastic error while the error stemming from the initial source guess is still large. The variance in the estimators for each generation will be small, but there will be a bias since the first iteration of particles came from a guess rather than the fundamental mode distribution. This bias will persist for a long period of computational time if M is large. If instead fewer histories are simulated per generation, iterations can be performed more quickly to reach an unbiased source, but there will be large fluctuations around the true mean due to the small M. Rather than spending histories to refine an estimate of the source that is still far from the fundamental mode or converging more quickly at the cost of high stochastic error, it would be advantageous to first approach a stationary source relatively quickly by using more generations with fewer particles per generation and later reduce the variance of the estimators by increasing the number of histories per generation.

To this end we propose the following procedure:

1. Initialize a guess of the fission source distribution with M_i initial particles.

2. Iterate over N_I inactive generations with M_i particles per generation. Each neutron in the current source is tracked, and neutrons produced from fission are stored to create the source used in the next generation. At the end of each generation the new source is rescaled to M_i particles.

3. Over $\lceil \log_2(M/M_i) \rceil$ "intermediate" generations, the expected number of fission neutrons produced in a collision is biased by a factor of two, so that each source neutron is expected to produce two fission neutrons. Neutrons produced from fission are stored, and the new guess of the source is rescaled to twice the size of the current source.

4. Once the number of particles per generation reaches M, begin active generations.

This method implicitly balances the two types of error by coarsely converging the source before ramping up the number of particles to refine the results. Its implementation requires minimal modification of the code and delivers no additional computational burden. The first guess of the source is constructed from M_i initial particles rather than the full M, but inactive generations are otherwise carried out as usual. Once the source has converged, the number of neutrons in the problem is increased in a series of intermediate generations until it reaches the final desired quantity, M. The doubling scheme used to augment the neutron population is somewhat arbitrary—the number of particles could be adjusted from M_i to M in numerous ways. The purpose of the intermediate generations is to grow the population without excessively amplifying fluctuations in the source around the mean. Adding all $(M - M_i)$ particles at once would significantly magnify the statistical error. This ramp-up method was chosen both for its simplicity and, as we will see shortly, its effectiveness on a variety of problems.

III. NUMERICAL EXPERIMENTS

To assess the performance of the particle ramp-up technique, we modified the OpenMC [12] Monte Carlo transport code to run an eigenvalue calculation according to the procedure described above. To study how well the method can accelerate convergence in high dominance ratio cases and realistic reactor problems, we tested the procedure on the OECD/NEA source convergence benchmark suite [13], which comprises some of the most challenging source convergence problems, and on the BEAVRS benchmark [14], a 3D full core model of an operating pressurized water reactor.

The challenge of determining when the multidimensional fission source distribution has converged is simplified by computing a scalar value that characterizes the source at each iteration. This quantity, called the Shannon entropy, is widely used in MC criticality calculations as a metric to assess source convergence [15]. The Shannon entropy H can be found at any generation by superimposing a rectangular mesh over the problem domain and tallying the number of source particles in each mesh bin. H is computed as

$$H = -\sum_{i=1}^{b} S_i \log_2 S_i, \tag{6}$$

where b is the total number of mesh bins and S_i is the percentage of source particles in the ith bin. The entropy can take on values between zero (if all particles are in the same mesh bin) and $\log_2 b$ (if the particles are uniformly distributed among bins). Plots of entropy vs. generation number can be used to diagnose convergence by determining the generation at which the entropy becomes stationary. In the tests performed here, the same mesh is used to compute the Shannon entropy in the reference calculation and in every stage of the ramp up calculation. The number of bins in each dimension of the mesh was chosen as $\lceil (M/20)^{1/3} \rceil$ [15].

Because the entropy is calculated from an estimate of the source distribution, composed of a finite number of particles, rather than the true distribution, it will deviate from the true entropy of the underlying distribution. The expected value of the entropy at any generation is a function of M. For $M = 1$, only one bin contains any particles, so the entropy is zero. As $M \to \infty$, the expected entropy approaches the true entropy. As a result, the entropy will converge to a different value for different choices of M, even if b is the same. This does not affect the application of the entropy as a metric to assess source convergence, as only the point at which the entropy reaches stationarity is of interest, not the true value of the entropy.

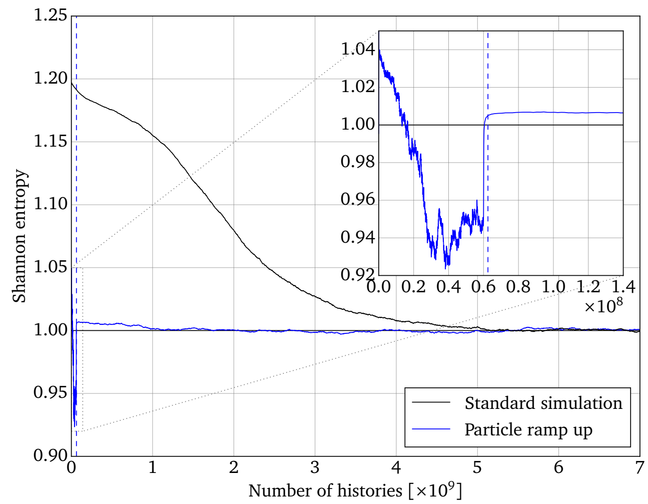

Fig. 1. Source convergence based on Shannon entropy for the OECD/NEA source convergence benchmark problem 1 with $M = 2 \times 10^6$, $M_i = 2 \times 10^4$, and $N_I = 3 \times 10^3$. The dashed vertical line marks the point at which the number of histories per generation reaches M.

The first problem used to assess the ramp-up method, the OECD/NEA benchmark, consists of test cases developed from real criticality safety problems that are used as a basis for source convergence comparisons. The poor convergence of the problems stems from high dominance ratios and undersampling of fissionable regions. Test problem 1 is a fuel storage facility in which a checkerboard array of fuel elements is surrounded on three sides by concrete and on the fourth side by water. Each fuel assembly is a 15x15 lattice of Zr-clad UO_2 enclosed by a steel wall and flooded with water. Since adjacent fuel elements are highly decoupled, adverse effects of undersampling emerge, and convergence is slow and difficult to diagnose.

Figure 1 shows the normalized Shannon entropy as a function of the number of total histories for problem 1 for both normal power iteration and the particle ramp-up method. A flat source distribution was used as the initial guess in this case and for each of the other results presented here. The dashed vertical line marks the point at which the doubling ends and the number of histories per generation reaches M. The zoomed inset show the detail of the ramp-up period. In the ramp-up method 2×10^4 initial particles and $3,000$ inactive

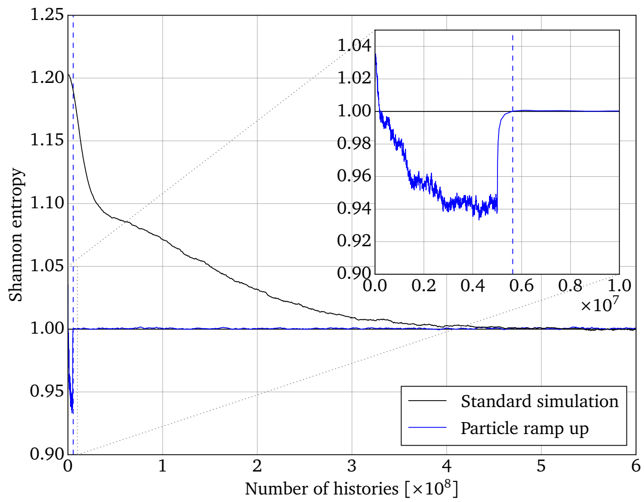

Fig. 2. Source convergence based on Shannon entropy for the OECD/NEA source convergence benchmark problem 2 case 2_3 with $M = 5 \times 10^5$, $M_i = 5 \times 10^3$, and $N_I = 1 \times 10^3$. The dashed vertical line marks the point at which the number of histories per generation reaches M.

generations were used. The standard simulation was run with 2×10^6 particles per generation. For normal power iteration, we see that 5×10^9 histories are needed to achieve stationarity. The entropy has converged by 1×10^9 total histories using the particle ramp-up method, a 5× speedup. However, note that the source has not fully converged by the point at which the ramp-up ends (after ~6×10^7 total histories); additional inactive generations are needed with the full M particles to converge.

OECD/NEA test problem 2 is a depleted fuel pin-cell array in water. The fuel rod is divided axially into nine regions, with the central region composed of high-burnup, low-multiplication fuel that decouples the two reactive ends. The problem consists of six cases with both symmetric and slightly asymmetric configurations of the fuel rod. We tested case 2_3, the most nearly symmetric burnup fuel case and worst case for source convergence. This problem has a high dominance ratio because of the length of the fuel rod, although undersampling is not a concern as in problem 1.

Figure 2 shows the normalized Shannon entropy as a function of the total number of histories for problem 2_3 for both normal power iteration and the particle ramp-up method. The standard simulation was run with 5×10^5 particles per generation. The ramp-up method uses 5×10^3 initial particles over 1,000 inactive generations. Again, we see a clear difference in the number of histories needed to reach a converged source distribution between the two methods; power iteration requires approximately 4×10^8 histories, whereas only 1×10^7 histories are needed for the particle ramp-up method.

OECD/NEA test problem 3 describes two slabs of uranyl nitrate solution decoupled by a slab of water. Twelve cases arise from varying the thickness of the water and one of the fissile slabs; we studied case 1, consisting of a 20 cm water slab and 20 cm and 30 cm fissile slabs. This problem suffers from undersampling and from slow convergence due to a

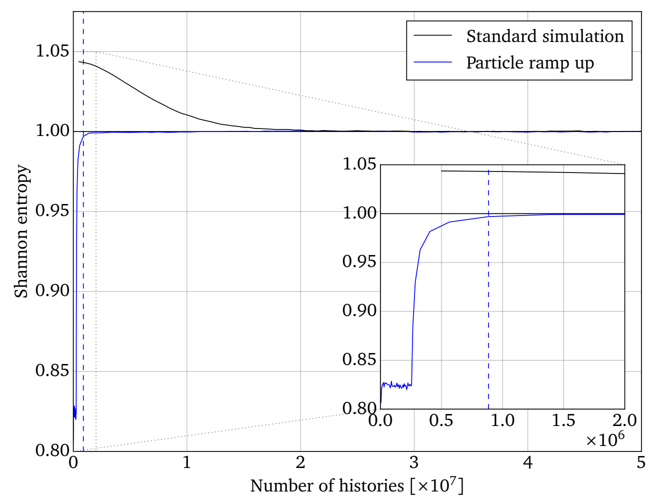

Fig. 3. Source convergence based on Shannon entropy for the OECD/NEA source convergence benchmark problem 3 case 1 with $M = 5 \times 10^5$, $M_i = 5 \times 10^3$, and $N_I = 50$. The dashed vertical line marks the point at which the number of histories per generation reaches M.

fairly high dominance ratio. Figure 3 shows the entropy for the normal power iteration (run with 5×10^5 histories) and the ramp-up method (run with 5×10^3 initial particles and 50 inactive generations) for this problem. The standard simulation converges after 2×10^7 histories, while the ramp-up technique yields a 10× speedup, converging after 2×10^6 histories.

The final OECD/NEA test problem is a 5*x*5*x*1 array of uranium spheres in air in which the center sphere is larger than the others. Unlike the previous three test cases, the dominance ratio of problem 4 is relatively low, and poor convergence arises from undersampling. The observed convergence for this problem is shown in Figure 4. The standard simulation converges after 4×10^7 histories, whereas the simulation using the ramp-up method converges after 4×10^6 histories, a 10× speedup.

To evaluate the effectiveness of this method on a detailed and relevant reactor problem, we simulated the 3D full-core BEAVRS benchmark at beginning-of-core, hot zero-power conditions. Figure 5 shows the normalized Shannon entropy as a function of histories for this problem. The ramp-up method was run with 4×10^4 initial particles over 300 inactive generations, and 2×10^6 particles per generation were used in the normal power iteration. We observe that the source has converged after 2×10^7 total histories using the ramp-up technique, whereas it takes 10× more (2×10^8) histories to reach convergence with standard power iteration. The results from each of the test cases studied here are presented in Table I.

The only new parameter introduced by the ramp-up method is the number of initial particles M_i. Convergence of the source is sensitive to the choice of M_i. If it is too low, not all regions of the domain may be sampled sufficiently, and fluctuations due to statistical error may be amplified. If it is too high, unnecessary time is spent in inactive generations. Figure 6 shows the Shannon entropy as a function of total histories for the OECD/NEA benchmark problem 2_3 for different

TABLE I. Comparison of source convergence results using standard power iteration and the ramp-up method.

	M	M_i	N_I	Histories to converge (standard simulation)	Histories to converge (ramp-up)	Speedup
OECD/NEA problem 1	2×10^6	2×10^4	3000	5×10^9	1×10^9	5×
OECD/NEA problem 2 case 2_3	5×10^5	5×10^3	1000	4×10^8	1×10^7	40×
OECD/NEA problem 3 case 1	5×10^5	5×10^3	50	2×10^7	2×10^6	10×
OECD/NEA problem 4	5×10^5	5×10^3	100	4×10^7	4×10^6	10×
BEAVRS	2×10^6	4×10^4	150	2×10^8	2×10^7	10×

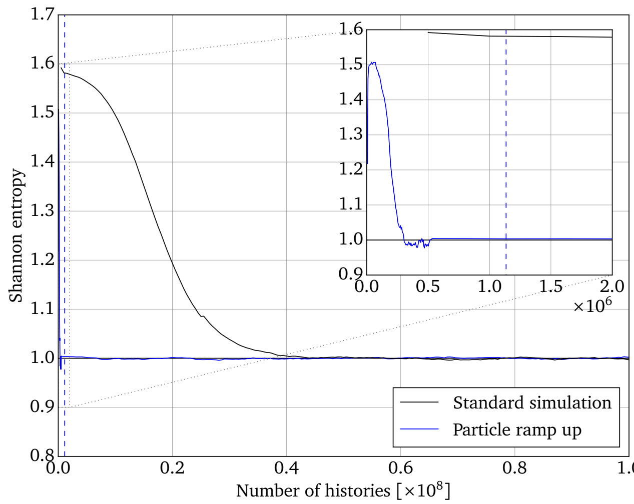

Fig. 4. Source convergence based on Shannon entropy for the OECD/NEA source convergence benchmark problem 4 with $M = 5\times10^5$, $M_i = 5\times10^3$, and $N_I = 100$. The dashed vertical line marks the point at which the number of histories per generation reaches M.

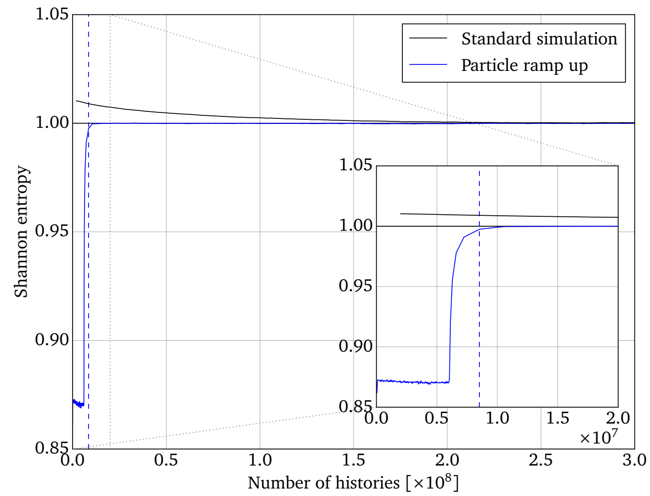

Fig. 5. Source convergence based on Shannon entropy for the 3D full-core BEAVRS benchmark at beginning-of-core, hot zero-power conditions with $M = 2\times10^6$, $M_i = 4\times10^4$, and $N_I = 150$. The dashed vertical line marks the point at which the number of histories per generation reaches M.

choices of M/M_i. For larger ratios of M/M_i (fewer initial particles) the source distribution does not immediately settle around equilibrium, but as the ratio decreases, the convergence becomes more stable. In each of the test cases studied here, we found $M/M_i \approx 100$ to be a reasonable choice that resulted in fast and stable convergence.

Next, we compare the effectiveness of this ramp-up method to a linear increase in the number of particles over inactive generations. Starting with a source made up of M_i particles, a constant M_a particles are added to the source at each inactive generation, so that at the first active generation the source consists of

$$M = M_i + M_a(N_I - 1) \tag{7}$$

neutrons. The value of M_a is constrained by the problem dependent N_I and M and by the choice of M_i as

$$M_a = \frac{M - M_i}{N_I - 1}. \tag{8}$$

Starting with large M_i (and therefore a small M_a) will not provide much benefit as the total number of inactive histories will be large. The largest possible choice of M_a, which will minimize the total number of inactive histories, is M/N_I. Figure 7 shows the observed source convergence using a linear particle increase and standard power iteration for OECD/NEA problem 2 case 2_3 with $M = 5\times10^5$ and $N_I = 1\times10^3$. The entropy for the case $M_a = M/N_I$ is shown in blue. As expected, it demonstrates better convergence behavior than the cases with smaller M_a; however, none of these cases approach the efficiency of the ramp-up method, which keeps the number of histories simulated prior to convergence small, and they do not appear to perform much better than even the standard simulation.

IV. CONCLUSIONS

Monte Carlo k-eigenvalue calculations are typically run with a constant generation size during both inactive and active generations. However, doing so results in an unnecessary number of source particle histories being simulated during inactive generations. Several works in the literature [9, 10, 11] have suggested slowly increasing the generation size during inactive generations as a way to reduce the overall execution time. In this study, a new variation on the ramp-up method was proposed. The source distribution is first converged using

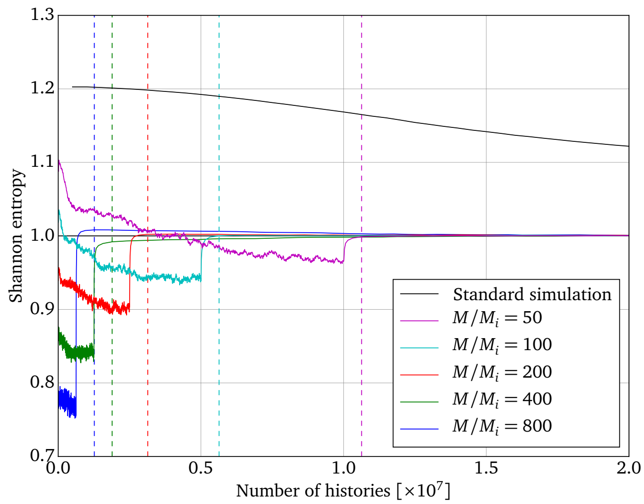

Fig. 6. Source convergence based on Shannon entropy for the OECD/NEA source convergence benchmark problem 2 case 2_3 for different choices of M/M_i with $M = 5 \times 10^5$ and $N_I = 1 \times 10^3$. The dashed vertical lines mark the point at which the number of histories per generation reaches M.

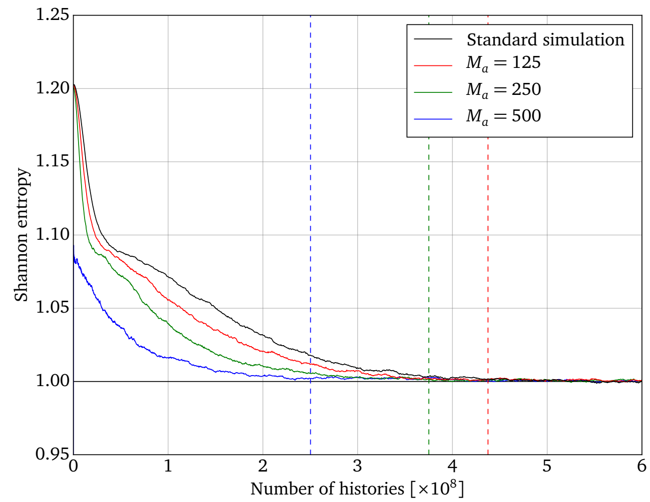

Fig. 7. Source convergence based on Shannon entropy for the OECD/NEA source convergence benchmark problem 2 case 2_3 using a linear particle ramp up with $M = 5 \times 10^5$ and $N_I = 1 \times 10^3$. The dashed vertical lines mark the point at which the number of histories per generation reaches M.

a small generation size. Once it is converged, the neutron population is doubled for as many generations as are required to reach the desired generation size for active generations. Simulations of five problems using a modified version of OpenMC demonstrated the effectiveness of the method in reducing the total number of histories spent in inactive generations and, therefore, the total amount of "wasted" time. The results also show that this method is more effective than the linear ramp-up originally proposed by Gast and Candelore [9].

In the version of the algorithm explored here, only one new parameter (the starting generation size) is introduced. Convergence of the source distribution is sensitive to the choice of this parameter: if the value is too low, the source distribution may not be completely converged by the start of active generations, but setting the value too high results in more wasted time in inactive generations. Variations of this algorithm can be conceived where the progression from a source of size M_i to a source of size M is different from what is described here. For example, rather than doubling the number of particles during intermediate generations, the population could be increased by a factor of 1.5 or otherwise. The algorithm in its present form may not be suitable for all problems—nevertheless, the results on several difficult source convergence benchmark problems are encouraging. With further improvements, the method could be made more robust.

V. ACKNOWLEDGMENTS

This material was based upon work supported by the U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research, under Contract DE-AC02-06CH11357.

REFERENCES

1. R. J. BRISSENDEN and A. R. GARLICK, "Biases in the Estimation of Keff and Its Error by Monte Carlo Methods," *Annals of Nuclear Energy*, **13**, *2*, 63–83 (1986).
2. C. M. PERFETTI and B. T. REARDEN, "Quantifying the effect of undersampling in Monte Carlo simulations using SCALE," in "PHYSOR – The Role of Reactor Physics toward a Sustainable Future," Kyoto, Japan (Sep. 28–Oct. 3 2014).
3. T. YAMAMOTO and Y. MIYOSHI, "Reliable Method for Fission Source Convergence of Monte Carlo Criticality Calculation with Wielandt's Method," *Journal of Nuclear Science and Technology*, **41**, *2*, 99–107 (2004).
4. T. KITADA and T. TAKEDA, "Effective Convergence of Fission Source Distribution in Monte Carlo Simulation," *Journal of Nuclear Science and Technology*, **38**, *5*, 324–329 (2001).
5. M. J. LEE, H. G. JOO, D. LEE, and K. SMITH, "Coarse Mesh Finite Difference Formulation for Accelerated Monte Carlo Eigenvalue Calculation," *Annals of Nuclear Energy*, **65**, 101–113 (2013).
6. D. SHE, K. WANG, and G. YU, "Asymptotic Wielandt Method and Superhistory Method for Source Convergence in Monte Carlo Criticality Calculation," *Nucl. Sci. Eng.*, **172**, *2*, 127–137 (2012).
7. S. CARNEY, F. BROWN, B. KIEDROWSKI, and W. MARTIN, "Theory and Applications of the Fission Matrix Method for Continuous-Energy Monte Carlo," *Ann. Nucl. Energy*, **73**, 423–431 (2014).
8. L. PAN, R. WANG, and S. JIANG, "Monte Carlo Fission Matrix Acceleration Method with Limited Inner Iteration," *Nucl. Sci. Eng.*, **180**, *2*, 199–208 (2015).
9. R. C. GAST and N. R. CANDELORE, "Monte Carlo Eigenfunction Strategies and Uncertainties," in "Proceedings of the NEACRP Meeting of a Monte Carlo Study Group, ANL-75-2," Argonne, Illinois (July 1–3 1974).
10. X-5 MONTE CARLO TEAM, "MCNP – A General

Monte Carlo N-Particle Transport Code, Version 5, Volume I: Overview and Theory," Tech. Rep. LA-UR-03-1987, Los Alamos National Laboratory (2008).

11. J. DUFEK and K. TUTTELBERG, "Monte Carlo Criticality Calculations Accelerated by a Growing Neutron Population," *Annals of Nuclear Energy*, **94**, 16–21 (2016).
12. P. K. ROMANO, N. E. HORELIK, B. R. HERMAN, A. G. NELSON, and B. FORGET, "OpenMC: A State-of-the-Art Monte Carlo Code for Research," *Ann. Nucl. Energy*, **82**, 90–97 (2015).
13. R. N. BLOMQUIST and A. NOURI, "The OECD/NEA Source Convergence Benchmark Program," *Trans. Am. Nucl. Soc.*, **87**, 143–145 (2002).
14. N. HORELIK, B. HERMAN, B. FORGET, and K. SMITH., "Benchmark for Evaluation and Validation of Reactor Simulations (BEAVRS), v1.0.1," in "Proceedings of the International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering," Sun Valley, Idaho (2013).
15. F. B. BROWN, "On the Use of Shannon Entropy of the Fission Distribution for Assessing Convergence of Monte Carlo Criticality Calculations," in "Proceedings of PHYSOR-2006, ANS Topical Meeting on Reactor Physics," Vancouver, BC, Canada (September 10–14 2006).