ANL/RA/CP - 95657

STRATIFIED SOURCE-SAMPLING TECHNIQUES FOR MONTE CARLO EIGENVALUE ANALYSIS*

By

A. Mohamed** and E. M. Gelbard
Argonne National Laboratory
Reactor Analysis Division
9700 South Cass Avenue, Bldg. 208
Argonne, IL 60439
630/252-1052
amr@.anl.gov

To be presented at the

International Conference on the
Physics of Nuclear Science and Technology
Islandia, Long Island, New York
October 5-8, 1998

*Work supported by the U.S. Department of Energy, Nuclear Energy Programs, under Contract No. W-31-109-ENG-38.

**Send Correspondence to A. Mohamed.

DISCLAIMER

STRATIFIED SOURCE-SAMPLING TECHNIQUES FOR MONTE CARLO EIGENVALUE ANALYSIS

Amr Mohamed
Argonne National Laboratory
9700 South Cass Avenue
Argonne, Illinois 60439-4842
(630)252-1052

Ely M. Gelbard
Argonne National Laboratory
9700 South Cass Avenue
Argonne, Illinois 60439-4842
(630)252-4831

ABSTRACT

In 1995, at a conference on criticality safety, a special session was devoted to the Monte Carlo "Eigenvalue of the World" problem. Argonne presented a paper, at that session, in which the anomalies originally observed in that problem were reproduced in a much simplified model-problem configuration, and removed by a version of stratified source-sampling.[a] In this paper, stratified source-sampling techniques are generalized and applied to three different Eigenvalue of the World configurations which take into account real-world statistical noise sources not included in the model problem, but which differ in the amount of neutronic coupling among the constituents of each configuration. It is concluded that, in Monte Carlo eigenvalue analysis of loosely-coupled arrays, the use of stratified source-sampling reduces the probability of encountring an anomalous result over that if conventional source-sampling methods are used. However, this gain in reliability is substantially less than that observed in the model-problem results.

I. INTRODUCTION

In the course of calculating, by Monte Carlo methods, the reactivity of large, loosely-coupled arrays of reactive objects in which each object weakly interacts with all its neighbors, the Monte Carlo analyst may run into a variety of anomalous results as a consequence of the slow convergence characteristic of Monte Carlo computations of this sort.

A typical example of such a computational difficulty was documented by Elliot Whitesides in 1971,[1] in a problem which he called the "Eigenvalue of the World" problem. This name has since been applied specifically to the original Whitesides problem, but also to the class of Monte Carlo computations for arrays of similar physical characteristics. The original Eigenvalue of the World configuration was a $9 \times 9 \times 9$ array of plutonium metal spheres, each with a radius of 3.976 *cm*, arranged in a square lattice with a 60.96 *cm* pitch. The whole array was surrounded by a 30 *cm* light-water reflector, whose inner edge was located 30.48 *cm* from the centers of the outermost spheres.[2] The eigenvalue of this array was computed to be ~ 0.93. When the central sphere was replaced by another which was critical in isolation the new Monte Carlo system-eigenvalue was still ~ 0.93 with a very, misleadingly, small standard deviation.

In 1995, Gelbard and Roussel[3] proposed an improved computational method for the Whitesides problem, using a version of stratified source-sampling. This method was based on the assumption that the conventional source-selection process generates enough noise in the source so as to impede, substantially, convergence of the outer iterations. Stratified source-sampling was thus designed to weaken such fluctuations. In Eigenvalue of the World problems, convergence of power-method iterations will be slow even in deterministic calculations. Stratified source-sampling can be successful only to the extent that this already slow convergence will be slowed still further by the statistical maldistribution of starters among the various reactive objects.

In their study, stratified source-sampling was applied to a much simplified model-problem configuration in which the anomalies observed by Whitesides were reproduced. This model-problem was a configuration consisting of a set of 41 slabs. The whole slab-array was reflected at its left- and right-hand boundaries. All the

[a]This was only a presentation given at the Embedded Topical Meeting on Misapplications and Limitations of Monte Carlo Methods Directed Towards Criticality Safety Analysis, USDOE Nuclear Criticality Technology Safety Project Annual Meeting, San Diego, CA, May 17, 1995. An account of the contents of the presentation can be found in Ref. 3.

slabs were identical except for the central slab, whose k_∞ was somewhat higher than that of the others.

Real neutron transport was not treated in the model problem. Instead the differenced diffusion equation was solved by Monte Carlo, with one mesh point at the center of each mesh box. There is a very simple standard Monte Carlo random walk which may be used to solve the differenced diffusion equation.[4] It was essentially this method that was used to treat the model problem, with modifications that avoided the need for detailed Monte Carlo simulation of collisions.[3] These modifications were designed to take advantage of analytic solutions of the difference equations, but without altering, in any way, statistical properties of fission and eigenvalue estimates.

Numerical results obtained for the model-problem qualitatively indicated that, for proper convergence to the correct eigenvalue and its corresponding eigenvector, the use of stratified source-sampling dramatically reduces the number of neutron case-histories needed, per generation, over that required if conventional source-sampling methods are used. Thus, in the model problem, stratified source-sampling substantially improved the reliability of the Monte Carlo computations.

In this paper, with the goal of increasing the reliability of realistic Monte Carlo Eigenvalue of the World computations, stratified source-sampling techniques are generalized and applied to different Eigenvalue of the World configurations which take into account real-world statistical noise sources not included in the model problem, but which differ in the amount of neutronic coupling among the constituents of each configuration in order to define the conditions under which the application of stratified sampling becomes more (or less) reliable than that of conventional sampling.

II. EIGENVALUE OF THE WORLD PROBLEM

Before we discuss stratified source-sampling we need to develop an understanding of Whitesides' results. Since the original results have long since disappeared we will have to try to reconstruct them by repeating the Whitesides calculations, insofar as possible duplicating his configurations and Monte Carlo running strategy.

Whitesides' problem configuration has already been described above. A set of his Monte Carlo computations were run with 300 histories per generation and 103 generations, with the first three eigenvalues skipped in the computation of the average eigenvalue.[2] This same running strategy was used in our reruns of the Whitesides computations.

Typical results of a few of our Whitesides-problem reruns are shown in Fig. (1). Figure (1-a) shows the cumulative eigenvalue vs. generation number for the first 103 generations, skipping three in the computation of the eigenvalue average. The Monte Carlo computed eigenvalue in this case was 0.942 ± 0.00478, consistent with Whitesides' reported results.

Figure (1-b) shows a continuation of Fig. (1-a) for the subsequent 500 generations. It can be seen that the cumulative eigenvalue of the system starts to drift towards the correct eigenvalue which is generated by a converged source predominantly in the central sphere. In this case, the resulting eigenvalue tallied over 600 active generations becomes 0.99471 ± 0.00201.

Individual batch eigenvalues for the same run are depicted in Fig. (1-c). This figure illustrates the mechanism underlying the Monte Carlo performance anomaly encountered in the original Eigenvalue of the World problem. It can be seen that early eigenvalues get large contributions from off-center spheres, therefore tend to be ~ 0.93. It is not until approximately the hundredth generation that the fission source starts to concentrate in the central sphere, which only then contributes correct estimates of the eigenvalue of the system.

It should be noted here that the eigenvalue of the system with the central sphere replaced by the critical one is not detectably different from unity. This is due to the fact that once the global fission source converges to its fundamental mode, the remaining source in the other "normal" spheres becomes very small so that very few neutrons enter the central sphere from its neighbors. As a result, the effect of coupling is largely reduced and the resulting eigenvalue of the system becomes, essentially, that of the central sphere.

The above results suggest that the anomalous results which Whitesides saw were mainly due to slow convergence of the fission-source so that the number of generations considered (either skipped or tallied) is insufficient. This slow convergence behavior is typical (even for deterministic calculations) of what we may consider as a "generic" Eigenvalue of the World problem. The situation is likely to be encountered in the analysis of array configurations whenever a) the number of reactive bodies in the problem at hand is large, and b) the bodies are weakly coupled neutronically.

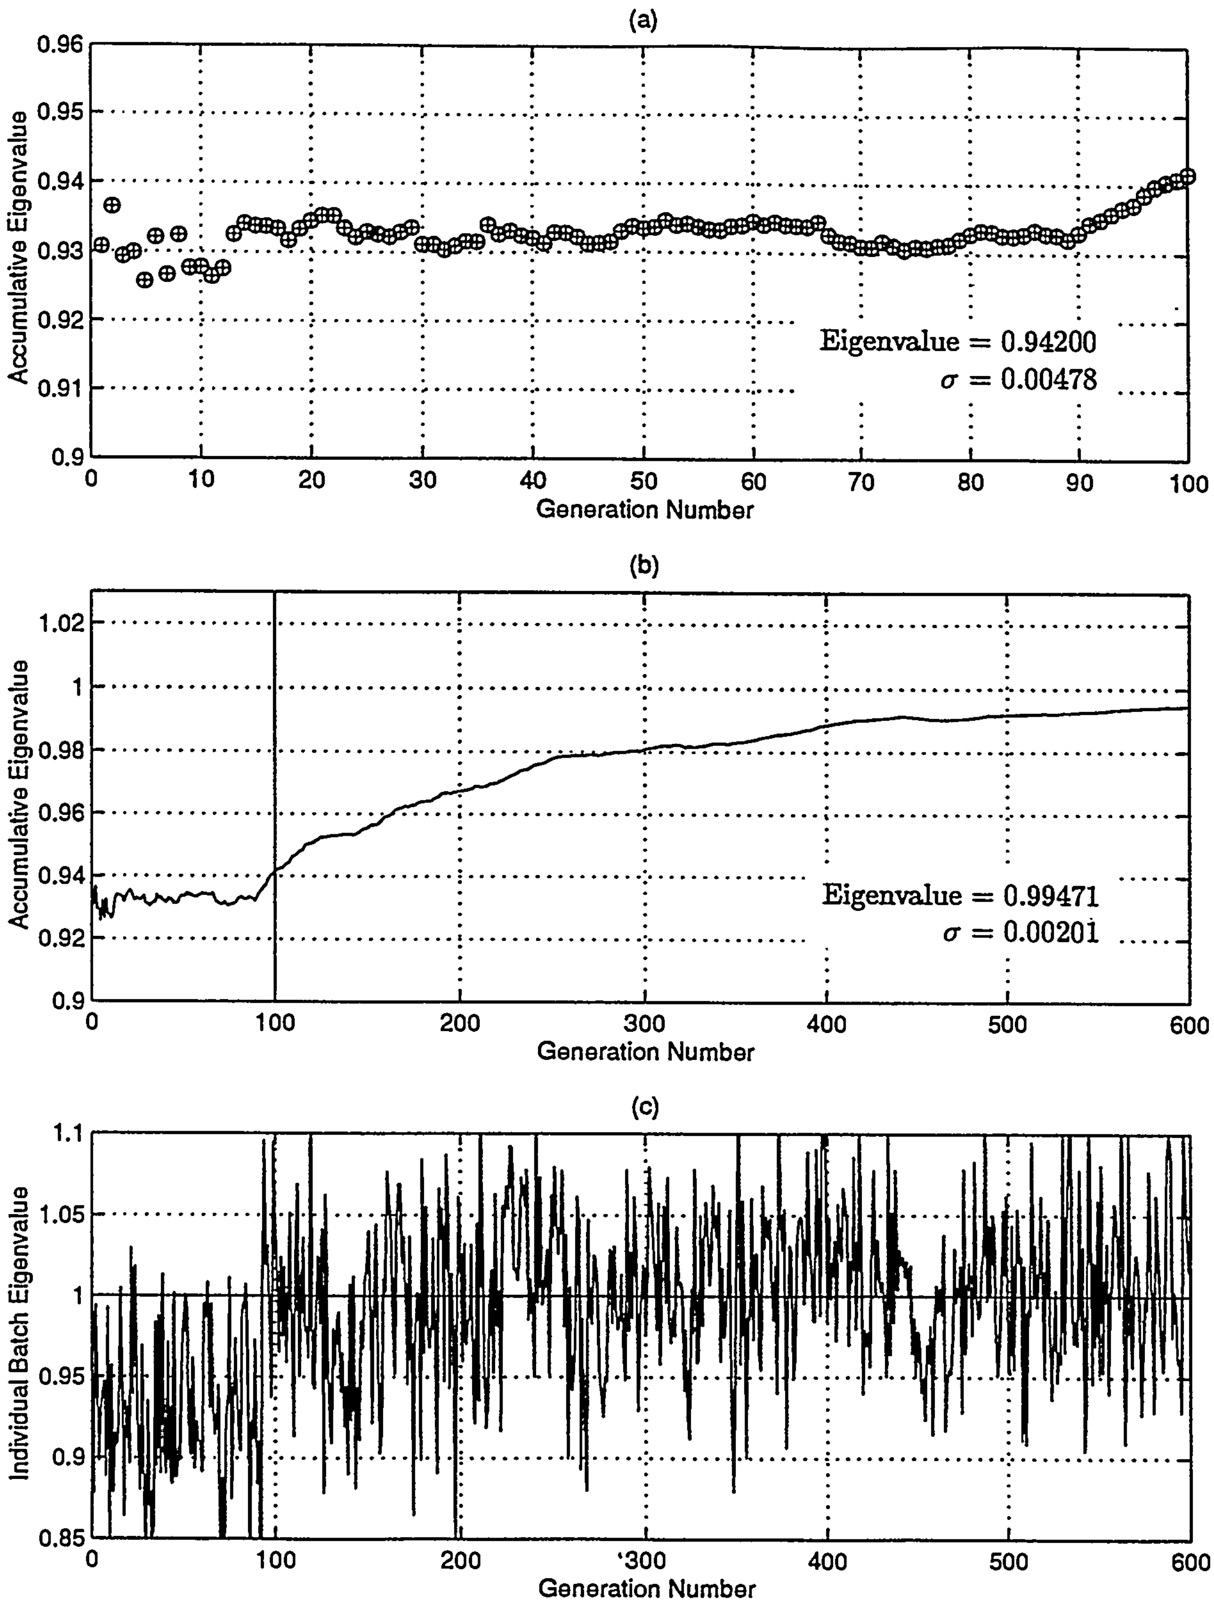

Fig. 1: Typical Whitesides Run: (a) Accumulative Eigenvalue for the First 100 Generations, (b) Accumulative Eigenvalue for the First 600 Generations, and (c) Individual Batch Eigenvalues for the First 600 Generations.

Unfortunately it is impossible to state these conditions more specifically. One can only say that anomalies characteristic of Eigenvalue of the World Monte Carlo calculations become more obvious as the number of objects increases and the coupling weakens.

III. STRATIFIED SOURCE-SAMPLING

The rationale underlying the formulation of the stratified source-sampling algorithm is best understood by analyzing the source-iteration mechanism for the model-problem in the case where the slabs are so large as to be essentially decoupled. In this limit the source iterations take on a particularly simple form. At the beginning of any generation each slab contains prospective fission sites left there at the end of the previous generation. In conventional Monte Carlo, source-sites for next-generation neutrons are sampled randomly from among these prospective sites. If the number of histories per generation is small the number of prospective sites will also, typically, be small, and there is a substantial probability that one or more slabs will be left with no starting sites for the next generation. In this case, in the absence of coupling, such slabs will permanently disappear from the Monte Carlo calculation, which will proceed just as if they did not exist. In each generation there is a finite probability that more slabs will disappear so that, eventually, only one slab will be left and the Monte Carlo calculation will seem to have converged. If this slab is not the central, most reactive slab, the Monte Carlo iterative process will completely converge to an incorrect source distribution, giving a wrong and misleadingly low eigenvalue with a very small standard deviation. Of course the behavior of the source-iteration process will be more complicated in the presence of coupling. This behavior will be discussed in Sec. IV.

The stratifed source-sampling algorithm is designed to diminish the level of noise generated by the source selection process. This is achieved, as will be seen, by stratifying the fission source in any given generation (i.e. dividing the number or weight of the neutron-starters among the different reactive objects in the problem) according to each objects fractional contribution to the total fission source as calculated in the previous generation.

In order to stratify the fission source for an initial source neutron size of N starters each with a weight W (where, typically, $W = 1$) let us define N_g^k and W_g^k to be the total number of starting neutrons and the weight of a single neutron starter in region k for generation g, respectively. We further define S_g^k to be the total weight of fission sites generated in object k at the end of generation g, then

1. At the end of generation $g-1$, we compute the regionwise probability density function p_g^k,

$$p_g^k = \frac{S_{g-1}^k}{\sum_{k=1}^K S_{g-1}^k},$$

2. We calculate the expected normalized total weight of neutron starters in region k for generation g, $\langle W_{tot,g}^k \rangle$

$$\langle W_{tot,g}^k \rangle = p_g^k N,$$

3. We then calculate the number of starters N_g^k and the weight W_g^k of each for generation g. To do so, we define a low weight cutoff threshold, $W_{cut} \ll 1$, so that
   - If $\langle W_{tot,g}^k \rangle < W_{cut}$, a game of Russian roulette is played. With probability $p = \langle W_{tot,g}^k \rangle$, we start one neutron from cell k carrying a unit weight, and with probability $1-p$ we start none.
   - If $\langle W_{tot,g}^k \rangle \geq W_{cut}$, we define $\tilde{N}$ to be the nearest integer to $\langle W_{tot,g}^k \rangle$, then
     - If $\tilde{N} = 0$, we take one starter from region k with a weight of $\langle W_{tot,g}^k \rangle$.
     - Otherwise, we take $N_g^k = \tilde{N}$ starters form region k each with weight $W_g^k = \langle W_{tot,g}^k \rangle / \tilde{N}$.

4. The locations of the N_g^k starters in each object k for generation g are then sampled from the fission sites registered at the end of generation $g-1$ in k with a probability distribution function proportional to the sites weights.

IV. NUMERICAL RESULTS

The various steps described above were "hard-wired" in the continuous-energy Monte Carlo transport code, VIM.[5] Three different variants of the Whitesides configuration were considered to test the performance of the stratified sampling algorithm vs. conventional sampling methods. These configurations were also designed to help assess the range of coupling for which the application of stratified sampling becomes more reliable than that of conventional sampling. Furthermore, since our main concern is to increase the reliability of Eigenvalue of the World type Monte Carlo calculations, we will need to adopt a criterion of reliability for all the test cases. Necessarily some large degree of arbitrariness will always enter such a criterion. Here we base our criterion on two numbers: First, the number of histories needed to give an eigenvalue accurate to within 1% in 19 out of 20 replications and, second, the number of replicas (out of 20) in which the 1% criterion is not met. Clearly an eigenvalue accurate to 1% is not always satisfactory. On the other hand the Eigenvalue of the World problem is really a criticality safety problem, and it would seem rash to allow an eigenvalue near 0.99 in such a context, for such a problem.

A. Case 1: Completely Decoupled Array

The first case considered is a completely decoupled model of the Whitesides configuration. Neutronic decoupling was achieved by placing an artificial purely absorbing material in the regions between the different spheres. Runs were submitted to search for the minimum number of neutron-starters per sphere required, by each method, to achieve a 95% reliability, i.e. to converge 19 or more out of 20 statistically independent replicas to the correct, non-anomalous result.

The anomalous behavior discussed at the beginning of Sec. III can be seen in the results shown in Fig. (2-a). The figure depicts eigenvalue estimates of 20 replicas (each run for 180 generations, skipping the first 30) for conventional and stratified sampling algorithms. It can be seen, according to the running criterion above, that for 6 starters per sphere (the minimum number required by stratified sampling to converge ≥ 19 replicas), the conventional and stratified sampling algorithms produce reliabilities of 85% and 100% respectively. The falsely-converged replicas (replicas 3, 14, and 18 in the first graph of Fig. (2-a)) estimate of the eigenvalue is ~ 0.817, which is the k-effective of a bare "normal" sphere. Furthermore, for this array configuration, conventional source-sampling requires 11 starters/sphere in order to produce the same desired 95% reliability. The "nominal" standard deviations for all replicas, i.e. those computed by the usual VIM method, are smaller than $\sim 0.3\%$, and for our purposes here can be regarded as negligible.

B. Case 2: Unreflected Array

The second configuration modeled is the Whitesides array without the water reflector. The model thus allows neutronic coupling between the different spheres in the array through direct flight of neutrons from one sphere to another. Because of the small radii of the spheres involved and the large center-to-center spacings, the eigenvalue of this system for the "all identical spheres" case was computed to be ~ 0.846 (indicating a small but significant level of coupling when compared to ~ 0.817, the single, bare sphere's eigenvalue). In addition, the eigenvalue of the system with the central critical sphere present is again ~ 1.0. The coupling in this case is negligible, again for reasons discussed in Sec. II.

Here it seems appropriate to try to understand the transition from no coupling to weak coupling. We know that, in the no coupling case, the conventional Monte Carlo iterative process will always converge to a state in which the whole fission source is confined to one sphere, though one cannot predict initially which sphere this will be. In the weak coupling case this cannot happen. If the distribution of fission neutrons among spheres is considered a "state" then (assuming that no sphere is totally decoupled from all others) there is always a nonzero probability that, in the next generation, the state will change to any other state. There is also a finite probability that the state will change to any other in a specified finite number of generations. Suppose we define a "quasiconverged state" to be one, attained at the end of a generation, in which all source neutrons are contained in a single sphere. Then, by the arguments above, there is a finite probability that a quasiconverged state will evolve into any other quasiconverged state in any specified number of generations. It follows that in a Monte Carlo computation very near the decoupled, the iterative process, as in the decoupled case, will seem to converge to a quasiconverged state; but it will eventually drift to a different one, probably one in which neutrons are confined to some sphere close to that occupied in the prior state. As the coupling increases we expect that quasiconverged states will no longer form, and that instead the overall fission source pattern will evolve more systematically into a source distribution close to the correct converged source. As the number of histories per generation increases the transition from erratic fluctuation to smooth evolution of the source should occur for weaker and weaker coupling.

None of the erratic source-behavior can be predicted by any existing theory. It will be recalled that present theories deal with "expected values", essentially averages over infinite numbers of replicas. It should also be noted that present theories dealing with Monte Carlo eigenvalue problems are correct only to leading order in $1/N_g$, where N_g is the number of histories per generation; nothing is known about the case where N_g is not "large enough". It seems likely, however, that the expected values of all Monte Carlo results will be biased when the number of histories per

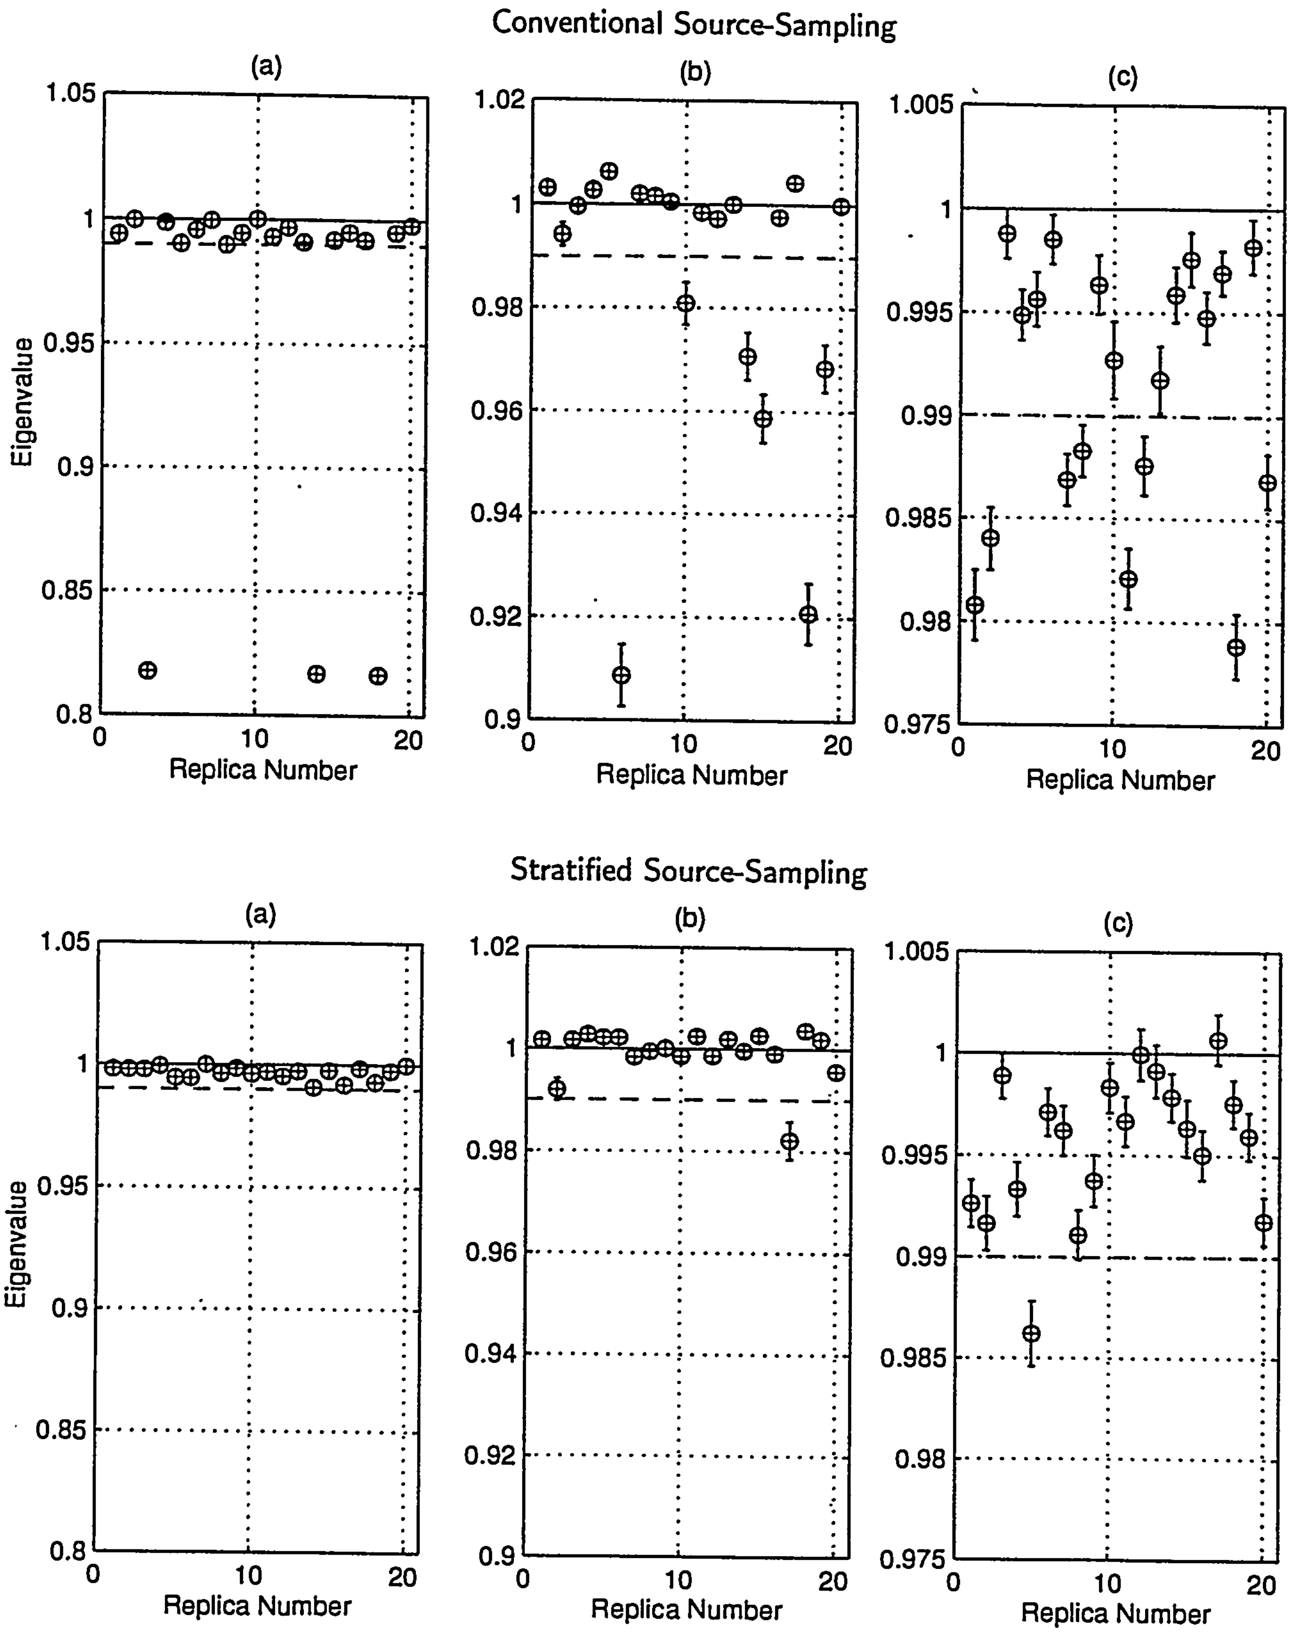

Fig. 2: Eigenvalue Estimates for 20 replicas: Conventional Source-Sampling vs. Stratified Source-Sampling for (a) Completely Decoupled Array, (b) Unreflected Array, and (c) Whitesides Configuration.

generation is small, so that averages over replicas will not be rigorously correct even as the number of generations goes to infinity. At present this bias has only been studied in simpler model problems.

Results for the case at hand are depicted in Fig. (2-b). The eigenvalue and σ estimates shown were obtained by running 3 starters per sphere for 200 generations in which the first 50 were skipped and the remaining 150 were tallied. For this running mode, the stratified source-sampling algorithm is shown to be 95% reliable, whereas the conventional one is only 70% reliable (6 starters/sphere are required for a 95% reliability). In addition, it can be seen that the maximum "apparent" standard deviation associated with any of the replicas is less than 1% (5.98×10^{-3} for replica 6 in the conventional sampling graph) which is, obviously, very small.

C. Case 3: Whitesides Configuration

In the Whitesides configuration, with the water reflector present, additional coupling is provided by virtue of neutrons reentering the spheres after being reflected in the water (the eigenvalue of the all identical spheres is ~ 0.93 and that with the critical sphere in the center of the array is, as indicated earlier, ~ 1.0). The same convergence criterion that was used in the previous cases was employed here. The results are shown in Fig. (2-c) for 6 starters per sphere. These results are based on a running mode which skipped the first 80 generations and tallied over the next 150. The results indicate that, with these Monte Carlo parameters, the reliability of stratified and conventional sampling algorithms are 95% and 60%, respectively.

There are no clear signs, here, of truly erratic behavior of the iterative process. Nevertheless it seems clear that erratic malditribution of starters among spheres must be impeding the convergence of the Monte Carlo iterations; otherwise stratified source-sampling would have no effect. To detect such an effect unambiguously one would have to study the change in convergence rate as a function of coupling strength. In fact, to improve our understanding of the Eigenvalue of the World problem, it is essential to undertake such a study in the future.

V. CONCLUSIONS

In Table (I) the Monte Carlo parameters and results for all cases are summarized. Suppose that r is the ratio of the minimum number of starters per sphere required by stratified and by conventional source-sampling to attain a "success rate" of 95%. We see that $r \sim 0.5$ in cases 1 and 2, and $r = .75$ in 3. Furthermore, from Fig. (2), we find that the "failure rate" in case 1 is about three times as great for conventional as for stratified source-sampling. This rate rises to ~ 6 in case 2 and ~ 8 in case 3. Thus the probability of anomalous results is substantially reduced in all of these computations. But here we have a puzzling result. Apparently stratified sampling develops an increasing advantage over conventional sampling as the degree of coupling increases. The problem seems to be that we have not succeeded in formulating a completely satisfactory criterion by which to judge the performance of either sampling method in Eigenvalue of the World computations. In case 1 the three "failures" of conventional sampling are pretty dramatic. In fact the computed eigenvalues are too low by about 18%, a real computational disaster. In contrast in case 3 the worst eigenvalue computed by conventional sampling is low by only about 2%. If we average over replicas, so as to estimate expected values of the computed k_{eff}'s, we get about 0.968, 0.986 and 0.991 for cases 1,2 and 3 respectively, results which show more nearly the trend we expect.

Unfortunately, these results do not necessarily indicate a firm quantitative behavior of the different parameters depicted as a function of coupling strength since the uncertainties in the numbers shown are large and can only be revealed by running a large number of sets of replicas. Nonetheless, it seems clear from our results that fluctuations in the source distribution are still important even in case 3, i.e. even in the presence of substantial coupling, and even if erratic behavior of this distribution is not immediately apparent.

The running times shown in the last column of the table (quoted for a 60 MHz Sun Sparc 20 machine) indicate that, apart from case 3, it requires more time to run a single case-history with stratified sampling than that required with conventional sampling (the reason for the "almost" exact running times in the last case is due to the fact that most of the CPU time is spent in the geometry tracking part of the history – an aspect of VIM that is the same for both algorithms). However, to achieve the same "success rate", the stratified sampling would then require less total running time as a result of requiring a smaller number of histories to be run.

Despite the work reported here there is no shortage of questions which remain to be answered. For example, many questions are suggested by the observation that stratified source-sampling was much more spectac-

Table I: Summary of Performance Evaluation Results: Conventional Source-Sampling (CSS) vs. stratified Source-Sampling (SSS)

		Generations Skipped	Generations Tallied	Starters Per Sphere	Reliability (%)	CPU/History (sec)
CASE 1	SSS	30	150	6	100	1.47E-03
	CSS	30	150	6	85	7.50E-04
		30	150	11	95	
CASE 2	SSS	50	150	3	95	1.83E-03
	CSS	50	150	3	70	1.16E-03
		50	150	6	100	
CASE 3	SSS	80	150	6	95	3.01E-02
	CSS	80	150	6	60	3.15E-02
		80	150	8	100	

ularly effective in the model problem (Sec. I) than in more realistic test computations. It is easy to see that, in realistic problems, there are many more sources of variance than in the very simple model problem, a problem designed to emphasize the effect of fluctuations in the fission-source selection process.

It remains to be seen whether more of the advantages of stratified source-sampling, as displayed in model-problem calculations, can be retained in the real world. In general, if there are very few major new variance sources in some real Eigenvalue of the World problem, as compared with a simple model problem, then it may be possible to deal with these few variance sources effectively. If there are many such major new variance sources it may be much more difficult to recover the effectiveness seen in model problem. Only a more complete study of real-world variance sources, and of appropriate variance-reducing devices, can tell us this.

ACKNOWLEDGMENTS

The authors would like to thank R. N. Blomquist for the help with the VIM code package and for many insightful discussions. B. Roussel and E. Feline contributed to the early coding stages of this work. This work was supported by the Environmental Management section of the U. S. Department of Energy through its Nuclear Criticality and Predictability Program.

REFERENCES

1. G. E. Whitesides, "A Difficulty in Computing the k-effective of the World," Trans. Am. Nucl. Soc., **14**, 680 (1971).

2. G. E. Whitesides, Private Communication.

3. E. M. Gelbard, B. Roussel, "Proposed Solution to the "k_{eff} of the World" Problem," Trans. Am. Nucl. Soc., **73**, 201 (1995).

4. Yu. A. Shreider, "The Monte Carlo Method, The Method of Statistical Trials," Pergman Press, 1966.

5. R. N. Blomquist, "VIM Continuous Energy Monte Carlo Transport Code," Proc. Int. Conf. Mathematics and Computations, Reactor Physics, and Environmental Analyses, Portland, OR, April 30 – May 4, 1995.